Cognitive Radio and Networking Research at Virginia Tech

A large research team with a wide range of expertise—from ICs and reconfigurable computing to wireless networking—works to achieve the promise of cognitive radio.

By Allen B. MacKenzie, *Senior Member IEEE*, Jeffrey H. Reed, *Fellow IEEE*,
Peter Athanas, *Senior Member IEEE*, Charles W. Bostian, *Fellow IEEE*,
R. Michael Buehrer, *Senior Member IEEE*, Luiz A. DaSilva, *Senior Member IEEE*,
Steven W. Ellingson, *Senior Member IEEE*, Y. Thomas Hou, *Senior Member IEEE*,
Michael Hsiao, *Senior Member IEEE*, Jung-Min Park, *Member IEEE*,
Cameron Patterson, *Senior Member IEEE*, Sanjay Raman, *Senior Member IEEE*, and
Claudio R. C. M. da Silva, *Member IEEE*

ABSTRACT | More than a dozen Wireless @ Virginia Tech faculty are working to address the broad research agenda of cognitive radio and cognitive networks. Our core research team spans the protocol stack from radio and reconfigurable hardware to communications theory to the networking layer. Our work includes new analysis methods and the development of new software architectures and applications, in addition to work on the core concepts and architectures underlying cognitive radios and cognitive networks. This paper describes these contributions and points towards critical future work that remains to fulfill the promise of cognitive radio. We briefly describe the history of work on cognitive radios and networks at Virginia Tech and then discuss our contributions to the core cognitive processing underlying these systems, focusing on our cognitive engine. We also describe developments that support the cognitive engine and advances in radio technology that provide the flexibility desired in a cognitive radio node. We consider securing and verifying cognitive systems and examine the challenges of expanding the cognitive paradigm up the protocol stack to optimize end-to-end network performance.

Lastly, we consider the analysis of cognitive systems using game theory and the application of cognitive techniques to problems in dynamic spectrum sharing and control of multiple-input multiple-output radios.

KEYWORDS | Automatic modulation classification; cognitive networks; cognitive radio; dynamic spectrum access; game theory; genetic algorithms; multiple-input multiple-output (MIMO); radio-frequency integrated circuit design; software defined radio; spectrum sensing

I. INTRODUCTION AND MOTIVATION

A cognitive radio (CR) is "a radio that is aware of its surroundings and adapts intelligently" [1]. This definition is, by necessity, a bit slippery. The problem is that cognition itself is an elusive quality; that which appears to be cognitive or intelligent prior to implementation is often dismissed as merely "adaptive" afterwards.

The need for CRs is motivated by many factors. Principally, though, the need for cognition is driven by the complexity of the radio systems themselves. The existence of software defined radio (SDRs) capable of implementing a near endless number of different waveforms with different modulation schemes, power levels, error control codes, carrier frequencies, etc., means that controlling the radio becomes a problem of combinatorial optimization. Such problems are often computationally hard and lend themselves to solutions based on metaheuristics—optimization methods based on simple search guided by higher level strategy. The application of such metaheuristics, which

often appear to learn and innovate, in turn, is characteristic of work in artificial intelligence.

The astute reader will note that neither our definition of CR nor the primary factor motivating the introduction of these radios explicitly mentions dynamic spectrum access (DSA). Our distinction between CR and DSA is intentional—we believe that the application of cognitive techniques, while appropriate to enable DSA, is much broader than DSA alone.

DSA, though, is important if only because of its looming presence as a "killer application" for cognitive techniques. For nearly a century, allocation of spectrum throughout the world has been based on a model of static allocation. More recently, it has been realized that this model leads to gross inefficiencies. While the entire radio spectrum from 6 kHz to 300 GHz is allocated [2], at any given point in space and time, most of the spectrum is unused (e.g., [3]–[6]). This observation has led regulatory agencies to seek more dynamic means of allocating spectrum. Such dynamic allocations might include secondary markets, spectrum commons, and licenses that enable users to access spectrum on a secondary basis. All of these techniques require radios to behave intelligently. In the case of secondary markets, radios must engage in negotiations with a spectrum broker to obtain access to spectrum appropriate to their needs. In a spectrum commons, radios must be aware of and respond to other users also using the commons. In the case of a secondary license, users must be alert for the appearance of primary users (PUs) as well as avoid other secondary users.

Extraordinary progress has been made on applying cognitive techniques to obtain seamless adaptation of radio link parameters, opportunistic use of underutilized spectrum, and increased flexibility in modulation and waveform selection to better fit the current wireless environment. Increasingly, however, there is the realization that such intelligent radios, when placed in a network, might bring about unexpected and undesirable results unless network considerations are carefully explored. We have termed a network that intelligently takes end-to-end goals into account a *cognitive network* (CN) [7].

While this paper reflects the broad scope and interdisciplinary nature of the Wireless @ Virginia Tech efforts to address the challenging problems raised in the creation of CRs and CNs, it is not a complete catalog of our related work. In particular, we have omitted significant work on the use of hidden Markov models to predict channel availability [8] and on the application of learning algorithms to the diagnosis and adjustment of cellular systems [9]. While we discuss another technique for signal recognition, we have omitted significant work on universal receiver architectures, signal detection, and recognition [10].

Although we have attempted to reference the most relevant work at other institutions, the scope of this paper is too broad to provide a comprehensive overview. In addition to pointers to specific work in individual sections, though,

we are aware of comprehensive efforts in the areas of CR and CN—spanning from the building of radio platforms to the creation of CN architectures and protocols—at some other institutions. The Information and Telecommunication Technology Center at the University of Kansas has been engaged in a broad program of CR and CN research, including the development of the Kansas University Agile Radio hardware platform [11]; in addition, they have participated in the creation of an experimental protocol stack for CR networks known as CogNet (with Rutgers University and Carnegie–Mellon University) [12]. In addition to participation in CogNet, the Wireless Information Network Laboratory at Rutgers University is involved in the development of CR hardware platforms such as WiNC2R [13]. Georgia Tech also has a substantial portfolio of work, ranging from radio-frequency (RF) chip development [14] to cooperative spectrum sensing [15] and higher layer issues such as the design of cognitive mesh networks [16]. The Berkeley Wireless Research Center at the University of California Berkeley has undertaken similarly broad efforts ranging from the development of a CR network emulator based on multiple field-programmable gate arrays (FPGA) [17] to DSA system design [18]. Finally, the Centre for Telecommunications Value-Chain Research in Ireland has a similarly broad program including both radio platform development and the development of reconfigurable radio and network architectures for cognitive systems [19].

This paper provides an overview of research related to CRs and CNs at Virginia Tech (VT). It is organized as follows. Section II briefly describes the history of work on CRs and CNs at Virginia Tech. Section III describes our work on reasoning and learning algorithms and architectures. Section IV describes work on supporting technologies such as signal recognition and radio environment maps (REMs). Section V describes our work on underlying radio platforms. Section VI addresses the problems of ensuring that CR systems are secure and, moreover, verifying that such systems will always operate within regulatory limits. Section VII presents research in the area of end-to-end network reconfiguration. Section VIII considers the application of game theory to analyze cognitive systems. Section IX considers research into specific application areas, including dynamic spectrum sharing (DSS) and multiple-input multiple-output (MIMO) systems, and Section X offers conclusions.

II. HISTORY

Early work in CR, led by Mitola, focused primarily on upper layer adaptation, in which the radio platform responded directly to anticipated user or application needs [20]. The radio seeks out the required information and provides the user with instructions or the desired service. Mitola also enhanced the observe, orient, decide, and act decision-making model taught to military officers [21] with

additional steps of "plan" and "learn" to create the "cognition loop" that has been widely used to understand and analyze the performance of cognitive processes in CRs and CNs [20]. Lastly, Mitola introduced the notions of "levels of cognition" as applied to CR; these levels allow us to assess our progress along the road to creating radios that are truly cognitive.

In many ways, work at Virginia Tech began with an opposite motivation from Mitola's work. Rather than seeking to satisfy users' goals, our work on CR began with the desire to exploit available technology. This work traces its roots to a project supported by the National Science Foundation, "Digital Government: Testbed for High-Speed 'End-to-End' Communications in Support of Comprehensive Emergency Management," which began in September 2000, led by Bostian. In the course of this project, the research team noticed the availability of paths-of-opportunity created by rough-surface scattering and hypothesized that these paths could be used to facilitate high-speed communications links in emergency scenarios. A system was needed that could autonomously identify, characterize, and equalize broadband 28 GHz channels based on signals propagating via rough surface scattering and communicate over these channels without requiring attention from a skilled operator [22].

By 2002, Bostian's Ph.D. student Rieser was focused on the concept of a cognitive engine (CE). Rather than seek to endow a specific radio with intelligence, the research team sought to create a software package that could intelligently control any radio. While Rieser was aware of Mitola's work, we believe VT was the first to separate the cognitive function from the radio platform and to focus on cognition at the physical layer.

Bostian's team embodied the cognition cycle in a CE using a multiple-objective genetic algorithm (GA) for both efficient optimization of radio configuration and as the basis for machine learning. In 2004, a patent application was filed describing this technology; this patent, titled "Cognitive Radio Engine Based on Genetic Algorithms in a Network," was issued in 2007 [23]. More recent work on this CE concept is described in [24] and in Section III.

Also in 2004, the team demonstrated a simple working prototype CR. This radio consisted of a multiple-objective genetic-algorithm-based CE paired with adaptive (but not software defined) Proxim radios. The CE controlled the radios' modulation type, modulation index, transmit power, and forward error-correction coding. The team established a video link between two of the radios and then turned on a jammer to disrupt it. The CE autonomously adjusted the radios to mitigate the effects of the jammer. It also remembered the actions it took and applied this knowledge the next time it saw the same jammer [25].

Concurrent with this work in the CR area was our work on applications of game theory to wireless systems. While much of this work does not explicitly use the term CR, game theory is a collection of tools for analyzing the interactions of rational agents. The assumption that radios will behave as rational agents presupposes the type of autonomous behavior required by the definition of CRs. We believe that [26] was among the first to make the explicit connection between game theory and networks of CRs.

In addition to the early start in CR, Virginia Tech has a history of significant research activity in the area of SDR. While CRs are not necessarily based on SDRs, it is often the complexity and capabilities of a SDR that make intelligent control of the radio necessary or desirable. Virginia Tech's work in SDR began in the early 1990s under the Defense Advanced Research Projects Agency (DARPA) GloMo program to develop reconfigurable computing, FPGA-like devices to support the functionality demanded by intense communication operations [27].

We have also examined the problem of the interactions that might arise between multiple CRs in a network. We first outlined these problems and possible approaches to solving them in [7] and further expanded these ideas in [28]. In these works, we defined a CN as a "network with a process that can perceive current network conditions, and then plan, decide, and act on those conditions. The network can learn from these adaptations and use them to make future decisions, all while taking into account end-to-end goals." Note that this definition retains the awareness, adaptability, and learning possessed by CRs but adds a focus on end-to-end goals.

Hearkening back to the origins of the VT work on CR, some of our recent work has centered on building a public safety CR, capable of recognizing any of the commonly used public safety waveforms and configuring itself to interoperate with them [29]. Important parts of this work include a universal signal detection and synchronization system [30] and techniques for rapidly configuring and reconfiguring a SDR under CE control. Full details appear in [10].

III. THE COGNITIVE CORE: ALGORITHMS FOR REASONING AND LEARNING

The CE is the intelligent system behind a CR (or a node in a CN) and combines sensing, learning, and optimization to control the radio (or network). Our recent work in CE development is described in Rondeau's dissertation [31], which describes both the theory and the prototypical implementation of the CE. It addresses cognitive components as well as particular issues related to developing algorithms for CR behavior. This work provides a theoretical foundation for developing the optimization algorithms required to design waveforms to meet particular quality of service (QoS) requirements under a given set of environmental conditions. Rondeau casts the problem as a multiple-objective optimization process that trades off objectives like bit error rate, data rate, and power consumption that measure radio performance. This provides a foundation for analyzing CR systems. This work also
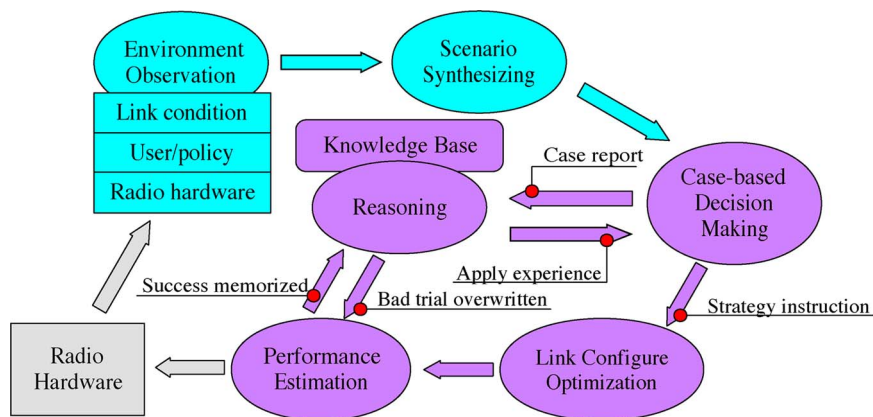
Fig. 1. *The Virginia Tech cognition cycle [31].*

provides examples of using feedback, learning, and knowledge representation in the CE.

Our conception of the cognition cycle proposed by Mitola is shown in Fig. 1 as applied to physical layer waveform adaptation. The outer loop is responsible for a reasoning process: On the basis of current observed conditions, take the best possible action through reliance on past experience and a metaheuristic optimization algorithm. The inner loop represents the learning process, through which past experiences are noted and influence the reasoning process.

Fig. 2 shows how the CE fits into the radio system as a whole. The figure includes three extrinsic domains that impact the radio.

- The user domain provides performance requirements to the radio.
- The policy domain constrains the CE to work within a given regulatory environment.
- The RF environment provides the context in which the transceiver will operate.

In addition, the figure shows a traditional communications stack on the right side. The CE operates on this stack in order to achieve the user's objectives in the policy and RF environment in which it is operating. It accomplishes these goals by controlling the stack—but is itself independent of the stack.

Fig. 3 shows a prototypical CE architecture. Central to this architecture is a cognitive controller that functions as the kernel and scheduler of the cognitive system. Other components of particular interest include:

- sensors, which collect and preprocess environmental data;
- the decision maker, which stores past experience and seeds the optimization processes;
- optimizers, which seek to develop optimized solutions to the problem currently posed by the user's requirements, the constraints of policy, and the radio environment by building on past experience.

Also of particular note in this architecture is the presence of generic interfaces between the components and the cognitive controller, allowing individual components to be modular. Current work includes further development of the architecture shown to provide a platform for international collaboration and comparative experimentation with CRs and CNs [32].

A. Genetic Algorithm Optimizer

The optimization of wireless system parameters is fundamentally a multiple-objective problem. Focusing on a single objective, such as minimizing bit error rate (BER) or maximizing signal-to-interference-and-noise ratio (SINR), in the wireless environment usually leads to poor solutions because they ignore other important objectives such as delay, spectral occupancy, power expenditure, or computational complexity. Instead, it is necessary for the optimizer


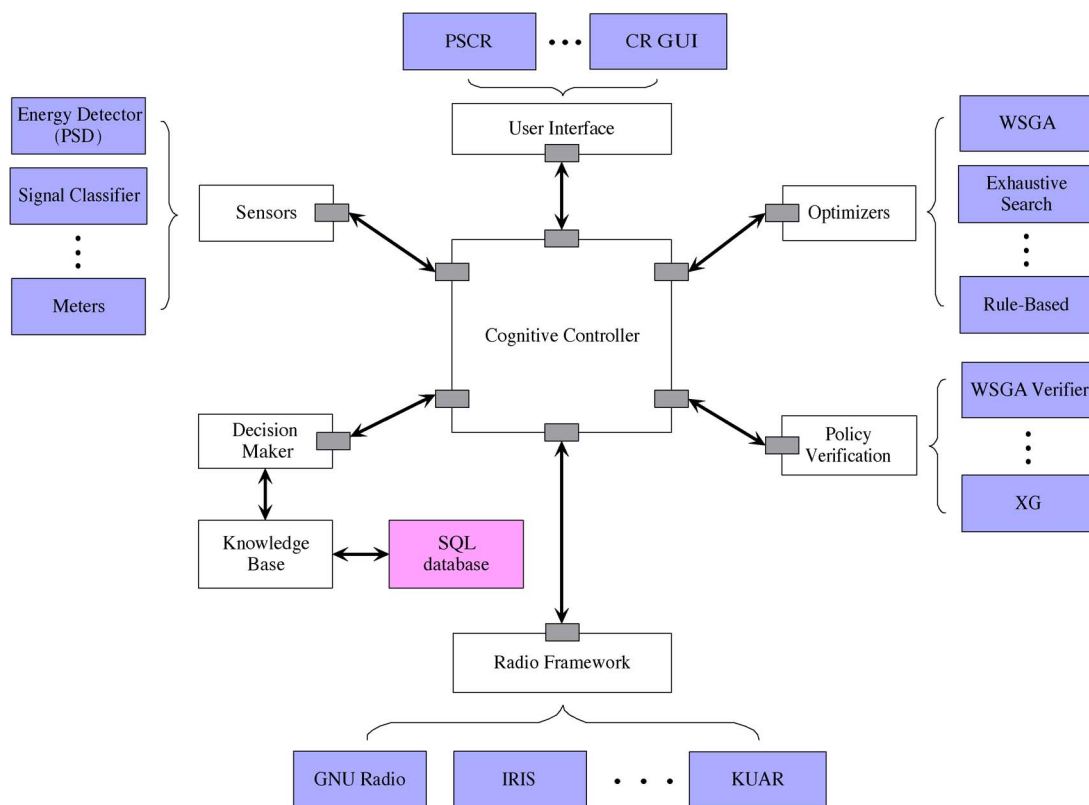**Fig. 2.** *A CR architecture [31].*

Fig. 3. *The CE architecture [31].*

in a CE to chase multiple objectives simultaneously. While it may be possible to create a single objective function by weighting individual objectives, this is quite challenging due to the difficulty of establishing appropriate weights a priori, especially when the objectives in question have wildly different units and magnitudes. Moreover, the relative rankings of these different objectives will change dramatically when considering different applications of the wireless system (e.g., file transfer versus telephony).

GAs excel in solving multiple-objective problems and are extremely flexible in representing different search spaces. While there are many ways in which a GA can compare different members of a population in a multiple-objective space, we consider Pareto ranking, where the members of the population are ranked by the number of individuals in the population to which they are Pareto superior.

Of course, the optimizer must ultimately return a single individual solution, or may need to select between individuals in the population with the same Pareto score. In so doing, we rely on a weighting of the objectives. However, the weighting is applied to normalized objectives, where each objective value is normalized against the best value that objective has received through all generations of the GA.

A key problem with using GAs as an optimizer is that they are notoriously slow, sometimes requiring thousands of generations to converge. Since the CE is working in a

real-time environment, this slow convergence is unacceptable. We solve this problem by seeding the GA's initial population with a carefully chosen set of solutions that correspond to promising areas of the search space. The choice of these candidate solutions to seed the GA's initial population is described next.

B. Case-Based Decision Maker

In addition to the problem with the convergence speed of GAs, just described, there is also a sense in which a GA is incapable of learning. While the GA may explore a search space thoroughly, when presented with a new problem it will start over rather than attempt to apply domain knowledge accumulated through past optimizations. It is this problem that the case-based decision maker attempts to solve. This concept was first described in [33].

The operation of the case-based decision maker is largely derived from case-based decision theory as developed in [34]. Essentially, the technique operates on a database of *cases*, in which each case consists of a problem faced, an action taken, and a result. When the decision maker faces a new problem, it uses a similarity function to compare that problem to those in the database, assigning a similarity value between zero and one to each case in the database. Each case in the database is further evaluated by assigning a utility, based upon the current objectives of the

CE, to the result of each case. Lastly, the cases in the database are ranked by taking the product of their similarities and utilities. The top-rated cases are then selected and the actions taken in those cases used to initialize the population of the GA (or other optimizer).

Many more details about this process and our implementation of it can be found in [31]. In summary, though, on average the case-based decision maker improves the performance of the GA. In particular, it can lead to fast attainment of high fitness solutions in few generations, which is desirable for the real-time application of the CE.

C. Cognitive Engine Experiments

Rondeau demonstrated the performance of the prototypical CE both through simulation and in over-the-air experiments with real radios. He performed a key set of over-the-air experiments during IEEE DySPAN in April 2007.

The setup for these experiments is shown in Fig. 4. To establish a link, a master CR sweeps the spectrum, determines what other radios are present, designs a waveform for the channel, and then pushes it to the other radio. It opens a link to provide a streaming audio service between the two nodes that provides a low error connection in the presence of three interfering radios operating in the 2.405 to 2.415 GHz band. Two of the interfering signals were 1-MHz-wide quadrature phase-shift keying (QPSK) signals generated by Trinity College Dublin IRIS software radios and a third was a 1-MHz-wide orthogonal frequency-division multiplexing (OFDM) signal from an Anritsu MG3700A signal generator. The signals were positioned at 2.4075 GHz (IRIS QPSK 1), 2.410 GHz (IRIS QPSK 2), and 2.4125 GHz (signal generator OFDM). When asked to design a waveform, the CE produced a 200 kbps QPSK signal with a 12 dBm transmit power.

In Fig. 4, the three interfering signals are seen at 2.4075, 2.410, and 2.4125 GHz, and the CE's waveform is located to the left of all three interferers at 2.4057 GHz. The signal on the right edge of the plot around 2.4145 GHz

Fig. 4. *Radio situation and spectral diagram showing interferers and the CR link at 2.4057 GHz.*

was not part of the experiment but a random signal that happened to be present, probably from a nearby DARPA XG demonstration.

In the early GA generations, the interference power for many of the solutions was large, but heavy selection pressure to minimize interference (the chosen objective weights emphasized minimizing interference) allowed the CE to find a spectrum free of interference quickly. The GA converged on a good solution within about 50 generations.

These results confirm that the CE can produce high-data-rate low-BER signals. The test also showed that the power spectral density sensor integrated into the cognitive node could accurately model the interference environment and that the CE optimized around the interferers.

IV. SUPPORTING TECHNOLOGIES

A. Spectrum Sensing

As previously mentioned, a key application for the development of CR systems is the promise of DSS, which could lead to more efficient spectrum usage. In order to achieve this promise, a CR system must first make sense of the radio spectrum activity in its surroundings. We define *spectrum sensing* as the combination of signal detection and modulation classification and use the general term automatic modulation classification (AMC) to denote this combined process. CRs may be required to perform AMC with no a priori knowledge of received signal characteristics. In this scenario, it is known that cyclic feature-based AMC is a possible approach with many advantages, including reduced sensitivity to noise and the ability to differentiate overlapping signals. This approach exploits the statistical characteristics of communication signals that vary periodically with time.

To take advantage of radio signal variability and allow for more reliable sensing, we present a *distributed* approach to cyclic feature-based AMC in which spectrum sensing is performed collaboratively by a network of radios. The distributed AMC system considered here is shown in Fig. 5. In this system, each of the radios consists of two stages: an AMC stage and a decision maker (DM) stage. The AMC stage processes the received signal by using features extracted from the cyclic spectrum of the signal for use in a neural network to perform pattern matching. The output of the AMC stage y_n is then used in the DM stage to determine the local decision u_n, which takes on a value in a finite alphabet. Finally, each local decision is sent to a fusion center, which uses these decisions, along with the result of its own AMC stage, to make a final global decision. More details on the technique described in this section can be found in [35], [36]. An alternative approach to the radio-level AMC problem is described in [10].
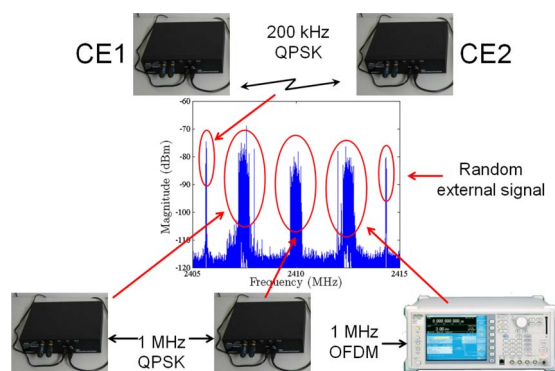
1) Radio-Level AMC Stage: The radios' AMC stage can be broken up into two main functions: feature extraction,

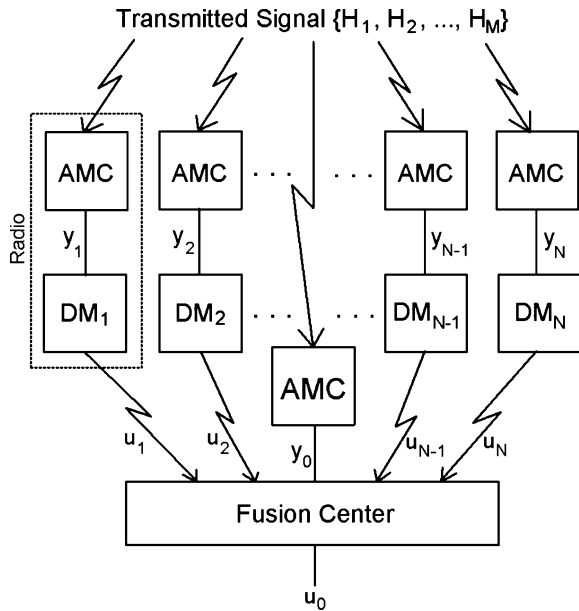**Fig. 5.** *Distributed spectrum sensing block diagram [35], [36].*

Fig. 6. *Conditional pdf for the output of the AMC stage at an E_b/N_0 of -2 dB [35], [36].*

in which the received signal's α-profile is found; and pattern matching, in which a trained feed-forward back-propagation neural network performs pattern matching on the α-profile. The α-profile, first defined for use in AMC in [37] and [38], is extracted from the cyclic spectrum of the received signal $x(t)$, defined as

$$\hat{S}_x^\alpha(f) = \lim_{T\to\infty} \lim_{\Delta t\to\infty} \frac{1}{\Delta t} \int_{-\Delta t/2}^{\Delta t/2} \frac{1}{T} U_T(t,f) V_T^*(t,f) dt \quad [39]$$

where

$$U_T(t,f) = \int_{t-T/2}^{t+T/2} x(u) e^{-j2\pi(f+\alpha/2)u} du$$

and

$$V_T^*(t,f) = \int_{t-T/2}^{t+T/2} x^*(u) e^{j2\pi(f-\alpha/2)u} du.$$

We estimate the cyclic spectrum through the use of a time-smoothing algorithm known as the fast Fourier transform (FFT) accumulation algorithm. For more details, refer to [35] and [40].
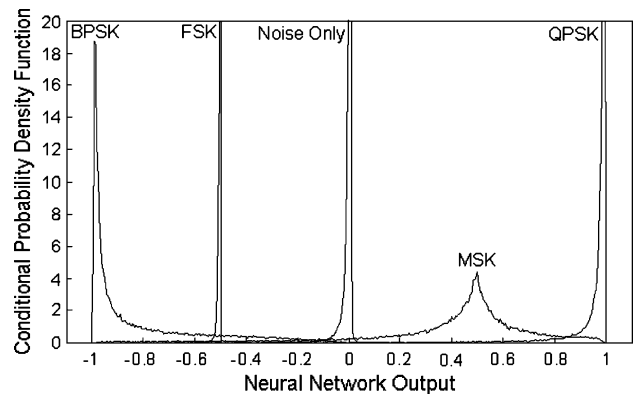
Once the cyclic spectrum has been determined, the α-profile is created by taking the maximum value along the spectral location parameter f for each spectral separation parameter α; i.e., $\text{profile}(\alpha) = \max_f[\hat{S}_x^\alpha(f)]$. This process reduces the size of the data to be used, allowing for a more computationally efficient pattern-matching algorithm without significantly reducing performance.

After the α-profile is created, a trained two-layer feed-forward back-propagation neural network is used to perform pattern matching on the profile. This neural network is trained on a set of α-profiles through the use of a Delta-Bar-Delta adaptive learning rate algorithm to give a modulation dependent output between -1 and 1.

As an example of the functionality of the radio's AMC stage, we assume a case in which there are four possible modulation schemes: binary phase-shift keying (BPSK), QPSK, frequency shift keying (FSK), and minimum shift keying (MSK), as well as the case in which no signal is present. The neural network is trained to yield the following modulation dependent outputs: -1 for BPSK, -0.5 for FSK, 0 for no signal present, 0.5 for MSK, and 1 for QPSK. In Fig. 6, the conditional probability density functions (pdf) of the output of the AMC stage y_n can be seen for an E_b/N_0 of -2 dB. It can be observed that the pdfs, conditioned on each of the possible five hypotheses, have a relatively small overlap even at this low E_b/N_0 value. From this result, it is clear that the proposed scheme has the potential to be of great use in detecting and classifying signals.

2) Distributed System Setup and Optimization: It was mentioned previously that the output of a radio's AMC stage is used by its DM stage to send a local decision, a message that takes on a value from a finite alphabet, to the fusion center. For this system, we can write person-by-person optimal decision rules for the fusion center and the DMs. These person-by-person optimal rules form a system of nonlinear coupled equations, and solving them is generally computationally hard. To avoid the "brute force" approach

Table 1 Probability of Classification for the Single Radio Case [35], [36]

	Hypothesis				
	Noise	BPSK	QPSK	FSK	MSK
Noise	0.9721	0.0020	0.0003	0.0000	0.0150
BPSK	0.0062	0.9780	0.0015	0.0067	0.0780
QPSK	0.0000	0.0000	0.9357	0.0000	0.0420
FSK	0.0001	0.0103	0.0001	0.9933	0.0022
MSK	0.0216	0.0097	0.0624	0.0000	0.8628

to solving for these rules, we use an iterative method based on the Gauss–Seidel algorithm. This algorithm, defined in detail in [41], allows for the decision rules to be solved in an efficient manner, at the expense of requiring messages to be passed between the radios and the fusion center. For more details on the optimal decision rules and the application of this technique to this problem, refer to [35] and [36].

To show the effects of performing distributed AMC over a single radio case, we expand on the AMC stage example. We assume that each of the radios in the distributed system is identical and has an AMC stage trained as discussed in the previous example, and that its outputs have empirical density functions as shown in Fig. 6. In Table 1, the results for the nondistributed case can be seen, while the results of the distributed scheme, with three radios and a fusion center, can be seen in Table 2. From these tables, it can be seen that performing AMC in a distributed manner greatly improves the detection and classification of signals over the single radio system. This can be seen by observing the average probability of classification error. In the single radio case, this error is approximately 5.2% but drops to approximately 0.2% for the distributed case with three radios and a fusion center. In the case of classifying MSK, the probability of correct classification rises from 86% for the single radio case to over 99% for the distributed case with three radios and fusion.

We have shown that performing AMC in a *distributed* manner can provide a significant increase in the probability of signal detection and correct classification over a single radio system, at the expense of requiring messages to be passed between the fusion center and the radios. The results from this work show that the proposed distributed AMC method, in the context of CR systems, can provide better analysis of the "spectral environment" by increasing the probability of signal detection and correct classification.

Table 2 Probability of Classification for the Distributed Case (Three Radios With Fusion Center) [35], [36]

	Hypothesis				
	Noise	BPSK	QPSK	FSK	MSK
Noise	0.9985	0.0000	0.0000	0.0000	0.0000
BPSK	0.0001	0.9998	0.0000	0.0008	0.0003
QPSK	0.0000	0.0000	0.9949	0.0000	0.0027
FSK	0.0000	0.0000	0.0000	0.9992	0.0000
MSK	0.0014	0.0002	0.0051	0.0000	0.9970

B. Radio Environment Maps

A distinctive characteristic of CRs and CNs is their capability of making decisions and adaptations based on past experience, on current operational conditions, and also possibly on future behavior predictions. An underlying aspect of this concept is that CRs and CNs must efficiently represent and store environmental and operational information. These resulting (individual or shared) databases enable different functionalities of the CE. In this context, we discuss a possible embodiment of such databases in the form of REMs. The application of REMs (also known as available resource maps) to CR systems was first proposed in the context of unlicensed wireless wide area networks in [42] and [43]. A detailed study of the use of REMs by different CE approaches can be found in [44] and [45].

A REM is a database that characterizes the environment in a given geographical area. The REM contains multidomain information such as spectral regulations, geographical features, and the locations and activities of radios [45]–[47]. Our work has shown that REMs are a practical, cost-efficient way to achieve a more efficient utilization of the spectrum and to help reduce harmful interference between CR systems and PUs of the spectrum.

REMs can be divided into two classes: global REMs and local REMs, which present a global view and a local view of the environment around the CR, respectively. A global REM is typically obtained from the network infrastructure, while a local REM is obtained by each radio from its own spectrum sensing and by monitoring transmissions of nearby CRs and PUs, for example. CRs use the REM's information to optimize their transmit waveform and other parameters across the protocol stack. Using the network simulation platform presented in [48], we present link- and network-level performance results of such a network.

1) Link-Level Simulations: Consider a scenario in which a CR, following a random waypoint mobility model, moves through a stationary PU network spread over a circular region. The average SINR improvement obtained by using the REM concept is shown in Fig. 7 for different interference radius to sensing radius (IR-to-SR) ratios, as a function of the CR speed. To obtain these results, we assume that the CR switches off its transmission any time it is aware that the PU network is within its interference range. As expected, the performance of local REM-based systems greatly depends on the IR-to-SR ratio. However, it is also seen that the global REM-based systems are not as dependent on this parameter; this is due to the fact that, in addition to their own sensing, radios also have access to a global REM (that contains information from multiple radios).

The performance improvement obtained by having a global REM comes at the cost of having to acquire and broadcast the contents of the global database to all CRs. To better understand the effects of such practical implementation issues, we evaluate the system performance

Fig. 7. *Performance comparison under different IR-to-SR ratios [36], [48].*

Fig. 9. *Network-level simulation scenario [36], [48].*

assuming PU mobility and information dissemination delay. The simulation scenario is similar to the previous simulation but with mobile PUs. The average SINR degradation at the PU nodes and the corresponding 95% confidence intervals are shown in Fig. 8. As expected, the simulation results indicate that the higher the speed of the PU nodes, the greater is the SINR degradation at the PU nodes due to the global REM dissemination delay. This is due to the fact that the locations of the PUs in the REM are out of date and the CR applies imperfect knowledge to adjust its transmit power.

2) Network-Level Simulations: In the network-level simulation scenario (shown in Fig. 9), 20 CRs are moving along the streets and 20 PU nodes are stationary and clustered at a street crossing. We assume that the CRs and the PUs both transmit and receive signals. In this analysis,

two typical geographical environments are considered: an open area and a dense urban area, where the two-ray ground reflection model and the Manhattan model are employed, respectively. The simulation parameters can be found in [48]. The following utility function is used to evaluate the performance of the two networks:

$$u = \frac{\text{sum throughput of both primary and CNs}}{\text{average packet delay experienced by the PUs}}.$$

Fig. 10 shows the increased network utility due to REM-enabled CRs in the context of spectrum sharing with incumbent PUs. As depicted in this figure, three different cases are considered.

1) The CRs are unaware of the topographical environment. When any PU node falls into their free-space interference range, they stop transmission.
2) The CRs estimate the path loss to the PU nodes by using the two-ray ground model and adjust their transmit power if any PU is within their interference range.
3) The REM-enabled CRs are fully aware of the radio environment and apply the Manhattan propagation model for path-loss prediction. Based on this estimate, the CRs adaptively adjust their transmit power if any PU is within their interference range.

The Manhattan propagation model differentiates the line-of-sight and non-line-of-sight conditions for appropriate path-loss prediction. The simulation results show that the high penetration loss due to the buildings in a dense urban area creates many "spectrum holes" that enable much higher


Fig. 8. *Average SINR degradation comparison under various PU moving speeds [36], [48].*

Fig. 10. *Network utility comparison when CRs adopt different adaptation schemes [36], [48].*

spectrum reuse by the REM-enabled CRs. Thus, the network utility for Case 3 is higher than that for Cases 1 and 2.

V. RADIO PLATFORM ISSUES

While SDR platforms are not strictly necessary to create CRs, the flexibility of an SDR is one motivation for pursuing the study of CR algorithms. The ability to change the functionality of the radio in response to dynamic conditions requires both intelligent algorithms and capable radio platforms. In particular, satisfying the demands of DSA requires frequency agile radio hardware. For example, mobile communications are an essential component of public safety operations, with systems in many frequency bands, including high frequency (HF) (25–30 MHz), very HF (VHF) (30–50, 138–174, and 220–222 MHz), ultra-HF (UHF) (406–512 MHz), and the 700 MHz, 800 MHz, and 4.9 GHz bands [49]. This profusion of operating frequencies and associated modes complicates interoperability and thereby impacts the effectiveness of public safety personnel [50]. Military users are in a similar situation. To effectively exploit frequency-agile CR techniques in applications such as these, it is necessary for radios to operate over large fractional bandwidths and even multiple bands simultaneously. For example, we may wish to use one (or more) RF chains to search for spectrum "white space" simultaneously with ongoing communications, or we may wish to bridge communications taking place in different bands. The ability of the radio platform to adapt dynamically and nearly instantaneously during operation is particularly valuable and novel—until recently, most work on SDR focused on static reconfigurability. Thus, research in CRs and CNs both drives and is driven by research into flexible radio platforms.

Although no operational radios capable of this kind of flexibility exist currently, single-channel radios covering

large subsets of these bands do exist. For example, the U.S. Department of Defense's Joint Tactical Radio System program has developed radios for frequencies spanning 30–512 MHz.[1] However, these radios are expensive and, ironically, have specifications that are in some cases inferior to existing commercial public safety radios [51].

This section describes efforts to address these important issues. First, we discuss a technique for real-time embedded reconfiguration of FPGAs. Then, we describe advances in RF/mixed-signal integrated circuit (IC) design that enable better and less expensive multiband radios. Lastly, we examine the application of a new single-chip complementary metal–oxide–semiconductor (CMOS) direct-conversion transceiver as the basis for a prototype simultaneous-multiband public safety radio.

A. Dynamic Assembly of Radio Structures

While the changes in functionality needed to adapt to new and unforeseen conditions can sometimes be achieved in software alone, there are many instances where general purpose processors and digital signal processors (DSPs) have insufficient computational capabilities to realize the desired functionality. FPGAs offer a potential solution to this problem by providing a means of realizing custom signal-processing pipelines within a reconfigurable fabric. The application of FPGAs to flexible radios has traditionally been hindered by the difficulty of exploiting the reconfigurability of FPGAs in an embedded environment. Designing computational structures "on-the-fly" has been intractable since powerful desktop computer-aided design (CAD) systems have been required to synthesize new computational structures. An alternative approach, though, is to synthesize the anticipated signal processing structures in advance so that they can be instanced on demand within the radio platform.

This section describes a framework that provides a means for dynamically assembling radio structures autonomously in embedded environments, eliminating the need for CAD tools at run time. At design time, all modules are wrapped, anchoring all ports at known locations and making modules self-contained. At run time, a flexible radio controller can insert or remove modules from the radio's signal-processing chain, and the framework will take care of placing the module and routing all connections between the modules and the rest of the design. Because no vendor tools are used at run time, the resulting application can be run on any platform. More details on the framework described in this section, called Wires on Demand, can be found in [52].

Partial reconfiguration (PR) has been supported in Xilinx's implementation tools with the addition of special constraints and bus macros to the modular design flow [53]. Although the PR flow has steadily improved since its introduction in 2002, substantial time and manual effort is still required to floorplan the dynamic module slots and

---

[1]See http://jtrs.army.mil/.

construct the routing between slots. All intermodule connection points must be implemented at compile time because no run-time environment is supported, precluding run-time placement and routing optimizations. The unique solution adopted by the Wires on Demand system is dynamic allocation of multiple modules within a large "sandbox" region, combined with an efficient run-time router to connect the modules.

1) Module Placement and Relocation: The design framework described here targets systems consisting of channeled datapaths connecting relatively large modules. As shown in Fig. 11, each module consists of a PR module (commonly known as an IP core) surrounded by a wrapper and compiled to a partial bitstream. The wrapper serves to connect modules and to pass through signals that need to be routed over the core. This approach permits dynamic module composition without the time and memory overheads normally associated with placement and routing algorithms.

Traditional FPGA tools tend to use time-consuming iterative approaches to tackle the difficult problem of placement. In a run-time environment, these approaches simply take too long. Previous work has shown that some performance may be sacrificed to reduce run-time costs while still achieving good results [54]. To reduce time and memory requirements of the placement process, placement occurs at the module level rather than at the gate level. This reduces the size of the problem from placing many thousands of cells to placing tens of blocks. Many previous approaches for run-time FPGA placement take a naive view of architecture, treating placement as a purely geometric problem or ignoring features such as block RAM that are not a part of the homogeneous logic fabric. Important issues such as timing and routeability are also ignored.

The main goal of the datapath placement approach is to promote neighbor connections and reduce routing delays between blocks by minimizing the lengths of the connecting wires. The first step is to arrange the modules along a one-dimensional axis in the data flow ordering. This one-dimensional ordering is then mapped to the FPGA. The precise placement of modules depends on the resources that they require, such as multipliers and block RAM. These extra resources are heterogeneous blocks separate from the regular reconfigurable fabric and are a limiting factor in determining the placement for a given module. The datapath is primarily routed in the vertical direction, with horizontal jogs to allow the datapath to fold around the FPGA. A timing estimation may also be performed to see if the design will meet its timing requirements along the critical path. Space is allocated to each needed module, and routing channels are reserved for data connections between modules. Using this approach, we have demonstrated dynamic assembly of radios in an untethered reconfigurable radio platform; see [52] for details.

2) Dynamic Module Library Preparation: The dynamic module library is composed of preprocessed partial reconfigurable models (IP cores) stored in the form of partial bitstreams. Before compilation, blocks are encased in wrapper structures that provide routing anchor points for block ports.

A module interface template describes the wrapper structure required by a particular functional block. Information in the template includes the port names and ordering, preferred block dimensions, dataflow direction, and routing options (such as the number of pass-through connections). Functional block preprocessing takes as input the module's port declarations and interface templates, and produces VHDL and constraints for a wrapped module. Mainstream CAD tools are then invoked to generate one or more bitstreams for the module. Defining similar interface templates for a set of modules promotes port alignment when the modules are connected.

Software has inspired high-level abstractions such as extending operating systems to manage FPGA configuration, but the low-level details are often left undefined. Careful partitioning is required between costly algorithms used to create efficient module implementations and run-time tasks. Efficient, algorithm-tailored module communication is as important as optimized modules to reap the benefits of reconfigurable computing. The framework summarized here acknowledges this by focusing on flexible, rapid, and efficient module composition.

B. Radio-Frequency and Mixed-Signal Integrated Circuit Design

Another aspect of CR research under active investigation at Virginia Tech is the area of RF/mixed-signal IC design.


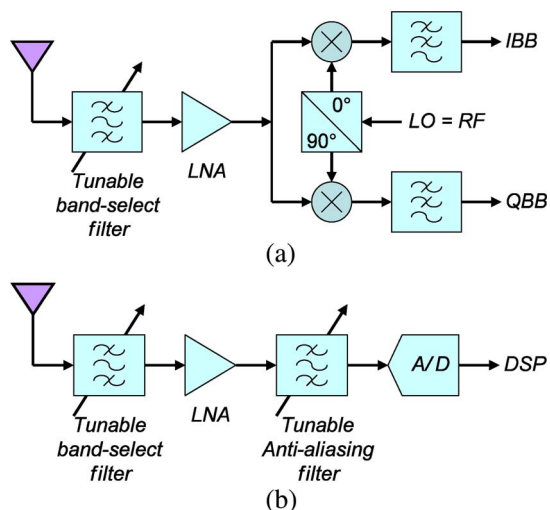**Fig. 11.** *Components of a dynamically assembled design.*

Fig. 12. *Candidate receiver architectures for CR. (a) Zero-IF (direct conversion) receiver architecture. (b) RF bandpass sampling receiver architecture.*

Receiver architectures currently under consideration for CR include 1) zero intermediate frequency (IF) (direct conversion) receivers (Fig. 12(a)) and 2) RF bandpass sampling receivers (Fig. 12(b)) [55]. Zero-IF receivers down-convert all desired frequencies directly to baseband. This is a mature approach for commercial wireless communications receivers, but for CR this puts extremely challenging requirements on the frequency synthesizer for local oscillator (LO) generation. In addition, appropriate, tunable front-end filtering is required to avoid blockers and meet dynamic range requirements. RF bandpass sampling receivers move the analog-to-digital converter (ADC) to RF, which enables extremely flexible signal reception. However, for the wideband, multimode CRs envisioned, ADC dynamic range and sample rate requirements will require large amounts of power for the foreseeable future. ADC clock jitter also impacts the dynamic range, and tunable RF filters are still required to mitigate some interference.

To reduce the ADC/DSP power consumption requirements of RF bandpass sampling architectures for CR, some of the signal-processing functions may be shifted to the analog domain. For example, we are applying analog signal-processing techniques for OFDM signal reception in ultra-wide-band (UWB) receivers [56]. To reduce the information conversion burden on the receiver ADCs, the multiplication-intensive FFT function is relocated ahead of the ADC into the discrete time analog domain. A prototype analog/mixed-signal FFT processor IC implementing serial-to-parallel/sample-and-hold and discrete analog multiplier-based FFT functions has been successfully demonstrated in CMOS technology. The IC can demodulate OFDM symbol streams at 1 GS/s with a linearity equivalent to a 9-bit ADC. The processor consumes 25 mW of power, a reduction of more than an order of magnitude compared

to an equivalent ADC followed by digital FFT function. The underlying basis for this is that analog multiplication is much more power efficient than digital multiplication. Current work involves extending the demonstrated FFT processor design to spectral sensing applications. Similar work at Georgia Tech and Samsung employs an analog wavelet transform block ahead of the ADC to perform spectral sensing in wireless regional area network (WRAN) applications [57]. These efforts represent a trend in the use of analog signal processing to relieve the power and complexity burdens of high speed in emerging CR systems.

Meanwhile, in the case of direct conversion receiver architectures, a significant challenge lies in the generation of all required LO frequencies with acceptable phase noise/spur performance. An attractive approach is multiband LO generation through multiplication or division from a "golden" phase-locked loop (PLL)/synthesizer or synthesizers (Fig. 13). This approach offers fast band switching (< 10 ns) with no PLL settling time issues by selecting amongst different mixer/divider outputs and allows the synthesized source to be optimized for power consumption, phase noise, and spurs at a fixed center frequency. By switching in different combinations of mixers and variable-ratio dividers, a wide range of LO center frequencies can be generated. A similar approach was recently taken to build an 800 MHz–5 GHz programmable SDR receiver [58]. Our recent work has included the design of UWB transmitter circuits in RF CMOS technology with multiband LO generation based on a wide-band PLL design and incorporating ultraprecise nanosecond pulse generation circuitry [59]. We are also currently developing highly linear, low-power resistive MOS single-sideband mixers with improved spur performance (< −40 dBc) for incorporation into multiband LO generation architectures [60]. Spurs have been cited as a key cause of false alarms in PU detection; hence, reducing them may be critical to CR performance.

An additional aspect of LO generation requirements for CR direct conversion receivers is the need for ultraprecise I/Q signal paths. To this end, we have developed new approaches to I/Q phase and amplitude error correction in


Fig. 13. *Diagram of multiband frequency generation using single fixed-frequency PLL, dividers and single-sideband mixers.*

integrated LC quadrature voltage-controlled oscillators (QVCOs). For example, we have demonstrated a 5 GHz range phase-tunable QVCO based on differential G_m tuning in CMOS technology [61]. An ultrawide ($\sim 30°$) I/Q phase balance tuning range was achieved with minimal impact on amplitude. While this amount of I/Q phase error is unlikely for the QVCO itself at this technology node, the excess phase correction capability can be used to correct for errors introduced by other components in the receive chain.

C. Building Multiband Radios

Through the application of technologies like those just described, a new generation of single-chip CMOS direct-conversion transceivers, which cover an astounding range of frequencies with performance and bandwidth sufficient for almost any wireless application in the tuning range, may soon become available [62]. One such chip is Motorola's recently announced 90-nm CMOS SDR RFIC [63]. We have been collaborating with Motorola since January 2007 to use this chip as the basis for a prototype simultaneous-multiband public safety radio (see Fig. 14).

An aspect of the multiband radio problem that remains essentially unsolved is antennas for such radios, specifically, how they can be integrated into the design without degrading performance or leading to objectionable sizes or shapes. For example, existing military handheld transceivers achieve this over 30–512 MHz by allowing the antenna-front end interface to be somewhat lossy, degrading sensitivity and the efficiency of final-stage transmit amplifiers [51]. The latter also tends to increase power consumption. The focus of our current research is to develop techniques for the design of multiband mobile radios using the same type of monopole antennas currently in common use, with performance comparable to existing single- and dual-band radios.

Fig. 14. *Prototype multiband radio using the Motorola RFIC. This design uses a four-band (138–174, 220–222, 406–512, and 764–900 MHz) antenna-transceiver RF multiplexer.*

This research also has application in vehicle installations, where the objective is mitigation of the need to install multiple antennas (the dreaded "porcupine effect").

It is certainly possible to design compact antennas that perform well on two bands simultaneously, and in certain cases the number of simultaneous bands can be increased to three or even four (e.g., [64]). Antennas that perform well in receive-only applications over large bandwidths are also often possible, as sensitivity is typically not a primary goal for such radios. However, developing antennas that perform well over many frequencies distributed over many bands is a daunting task, especially for transmit operation where low voltage standing wave ratio (VSWR) is important. Perhaps the most difficult aspect of the problem is dealing with frequencies less than 200 MHz, for which practical antennas become very short relative to wavelength. For example, a 20-cm linear antenna—about as long as most users will accept—is only about 2% of a wavelength at 30 MHz, the low end of the VHF-low band. In this case, the theoretical best possible bandwidth [65] for a 2 : 1 VSWR is less than 10 kHz. Thus, any impedance match between antenna and transceiver that is efficient from a power transfer point-of-view will have unacceptably narrow bandwidth and furthermore will need to be tuned as the channel changes. Attempts to bypass this limitation by designing the antenna to be well matched over the bandwidths of interest invariably result in lossy interfaces or awkward antenna shapes and sizes.

If we are prevented by physics from making acceptable antennas with better multiband characteristics, then we are forced to consider ways to modify the front end of the transceiver to achieve this. One approach is to use multiple transceivers operating in parallel (for example, the Motorola chip has five receivers and three transmitters), each of which can be directly connected to an off-chip filter bank. In this case, RF multiplexers (e.g., a diplexer if two bands, a triplexer if three bands, and so on) become attractive. A multiplexer separates the antenna output into appropriate frequency ranges, thereby providing selectivity for subsequent direct conversion tuning. Although the design of RF multiplexers is an old problem, the existing techniques [66] are focused on the problem of interfacing single-ended devices with roughly constant impedance to other single-ended devices with roughly constant impedance. In contrast, the impedance of compact antennas operating over large fractional bandwidths varies from extremely capacitive with very high Q (hence inherently narrow-band) at low frequencies to wildly variable at higher frequencies as various disparate current modes become more or less important with varying frequency.

We are developing multiband front ends consisting of multiplexers with the desired current mode conversion characteristics that simultaneously yield better overall performance—not necessarily from an impedance matching perspective, but rather from a *receive sensitivity* and *transmit efficiency* perspective. As an example of our ability to achieve

the former, we have demonstrated front-end designs that are capable of expanding the effective (external noise-limited) sensitivity of dipole antennas from about 10% to about 25% using a front-end "co-design" strategy [67], [68]. We are seeking additional improvements using "non-Foster" matching [69], which is a high-risk but high-payoff technique in which matching circuits are developed using active devices that achieve "nonphysical" impedance characteristics—e.g., negative capacitance—which result in dramatically improved impedance matching.

VI. SECURITY AND VERIFICATION OF COGNITIVE SYSTEMS

A. Security

In order to achieve successful deployment of CRs and CNs, these systems require robust security mechanisms to resist misuse. The emergence of DSS and CR raises new security implications that have not been studied previously. We are investigating several security issues within the context of CRs and CNs. In Fig. 15, we classify some of those security issues into two categories: spectrum access-related threats and radio software security threats. The former can be further classified into spectrum sensing-related threats and spectrum sharing-related threats. In this section, we briefly describe the security threats in each category and discuss some possible countermeasures. For examples of other work on security in CR and SDR systems, see [70] and [71].

1) Spectrum Sensing-Related Security Threats: A CR engaged in DSS needs to carry out spectrum sensing for the purpose of identifying fallow spectrum bands—i.e., spectrum "white spaces." Here, we focus our discussions on one problem that poses a threat to the spectrum sensing process—the primary user emulation (PUE) attack [72].

In the DSS paradigm, CRs opportunistically utilize fallow licensed bands after identifying them via spectrum sensing. These secondary users are permitted to operate in licensed bands only on a noninterference basis. A secondary user must constantly monitor for the presence of PU signals in the current operating band and candidate bands. If a secondary user detects the presence of a PU in the current band, it must immediately switch to another band. On the other hand, if the secondary user detects the presence of another secondary user, it invokes a coexistence mechanism to share spectrum resources. In a PUE attack, a malicious secondary user attempts to gain priority over other secondary users by transmitting signals that emulate the characteristics of a PU's signals. The potential impact of a PUE depends on the legitimate secondary users' ability to distinguish attacker's signals from actual PU signals while conducting spectrum sensing.

Energy detection is one of the simplest methods for spectrum sensing. In energy detection, a detector infers the existence of a PU based on the measured signal energy level. Obviously, energy detection is unable to distinguish primary signals and secondary signals and thus particularly vulnerable to PUE attacks. Another spectrum sensing approach, signal feature detection, uses more advanced techniques such as cyclostationary feature detection or matched filter detection to detect specific characteristics of PUs [73]. However, relying solely on signal features may not be sufficient to reliably distinguish PUs' signals from those of an attacker; hence, these spectrum sensors may also be vulnerable.

Detecting instances of a PUE attack is a challenging problem. When the PUs are stationary, though, a localization-aided countermeasure may be effective in detecting PUE attacks. For instance, consider PUE attacks in the context of an IEEE 802.22 network. In 802.22, broadcast towers are primary transmitters; hence, transmitter location and transmission power are fixed. If the location and approximate transmit power of a transmitter can be determined based on its signal, then a detector operating in an 802.22 network can distinguish between a primary transmitter and those of an adversarial secondary node. More detailed discussions of spectrum sensing-related security threats can be found in [72].

2) Spectrum Sharing-Related Security Threats: Spectrum sharing, or coexistence, is an important attribute of CR networks. Typically, CR networks support two types of coexistence: *incumbent coexistence* (i.e., coexistence between primary and secondary networks) and *self-coexistence* (i.e., coexistence between overlapping, cochannel secondary networks). It is possible for an adversary to exploit the vulnerabilities in the coexistence mechanisms to attack CR networks. To facilitate our discussion, we use IEEE 802.22 [74] as an example, but the security issues that we discuss are relevant to other types of CR networks as well.

Fig. 16 illustrates two overlapping 802.22 WRANs, and depicts an attack that exploits 802.22's self-coexistence mechanism. To achieve self-coexistence, 802.22 WRANs rely on intercell beacons to exchange spectrum utilization information and use the on-demand spectrum contention
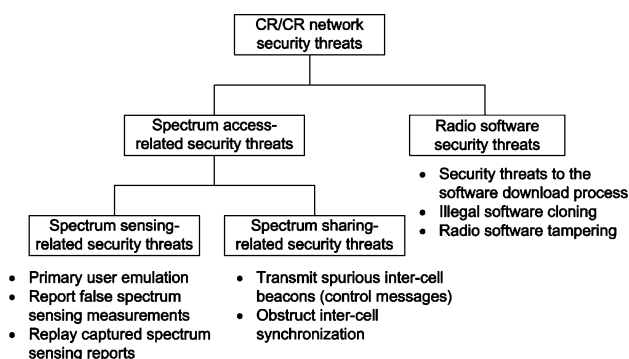

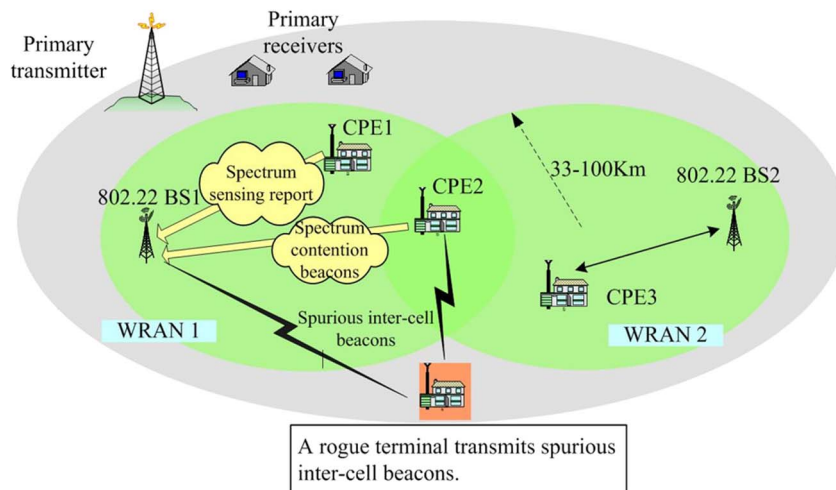**Fig. 15.** *Classification of CR and CN security threats.*

Fig. 16. *Security threats to CR network self-coexistence.*

algorithm to handle spectrum contention issues [74]. The contention process enables a cell to acquire better or more channels to support admitted workloads. Because direct inter–base station communication is not always possible, a base station collects a neighboring cell's spectrum utilization information by receiving reports from customer premises equipment (CPE) devices under its control that overhear the neighboring cell's beacons. Although these beacons provide important self-coexistence information, they are not protected by authentication mechanisms. This implies that a rogue terminal can send modified or forged intercell beacons to obstruct the spectrum contention process of a targeted cell. We coin the term *spurious beacon attack* to describe this threat, which can have a significant impact on the performance of the targeted cell by invoking unnecessary spectrum contention processes.

To address the aforementioned vulnerability, intercell control messages need to be protected using cryptographic solutions. This implies that an intercell key management system is needed, but operating such a key management system can be complex because contending cells may be managed by different operators. The existence of a common backhaul infrastructure among competing operators cannot be assumed, further complicating intercell key distribution. Although 802.22's *security sublayer* includes a key management protocol, this protocol only handles intracell keys and has no provisions to support the management of intercell keys.

3) Security Threats to Radio Software: Radio software for a CR has unique properties that distinguish it from conventional software. Because of the intrinsic operating characteristics of CRs, software running on them is likely to have the following attributes.

- *Complex and modular architecture.* Most of the existing SDR systems, including distributed object computing software radio architecture, software communications architecture (SCA), GNU radio, and Vanu software radio, adopt a complex, distributed, object-oriented software framework to promote modularity.
- *Reconfigurability.* CRs may be required to make frequent configuration changes, and radio software must support such changes.
- *Real-time requirements.* Radio systems have stringent real-time requirements. Hence, software execution timing must be tightly controlled.

Without proper software protection mechanisms in place, CRs are vulnerable to a host of attacks. We classify these threats into three broad categories: security threats to the software download process, illegal software cloning, and unauthorized software tampering. Although the first two problems have attracted attention from the research community (e.g., [75] and [76]), there is little existing research on the third problem in the context of CR software. The threat posed by the third problem is especially serious because adversaries may attempt to manipulate radio software to gain operational advantages (e.g., transmit at power higher than the authorized limit) or launch attacks against PU networks. The prospect of unauthorized changes to SDR/CR operating characteristics (e.g., power, frequency, and modulation) is a major concern for regulators and developers.

In recent years, a number of technical approaches have been proposed to protect software intellectual property and thwart exploitation of software vulnerabilities. Software obfuscation and tamper resistance are particularly appropriate for combating unauthorized software tampering. Obfuscation transforms a program into a functionally equivalent program that is more difficult to understand, thus thwarting reverse engineering. On the other hand, a tamper resistance scheme detects and/or prevents integrity

violations of the original software. Typically, adversaries carry out reverse engineering prior to modifying the software since modification requires at least a partial understanding of the target software. Therefore, obfuscation and tamper resistance should be considered together.

Existing techniques for tamper resistance and obfuscation, though, do not adequately address the problems of protecting radio software because they do not take into account the distinguishing features of radio software. For instance, the stringent real-time requirements of radio systems prohibit the use of code encryption for software obfuscation. Obviously, the aforementioned features of radio software need to be considered when designing tamper resistance mechanisms for radio software.

To ensure the security and reliability of CR software, the implementation of the aforementioned security solutions is not enough. Along with the security mechanisms, potential vulnerabilities within the radio software need to be identified and fixed through systematic testing and verification. We address the issues of testing and verification in the next section.

B. Verification

While SDR and CRs are becoming a reality, the advantages of such radios could be offset by a lack of security and reliability in the underlying software that serves as the command and control for the radio system. Guaranteeing security and functional correctness of the embedded software is critical to the success of their deployment. A subtle bug in the software could render the radio vulnerable to attacks that might not only crash the system but also could produce disastrous results, such as unwanted transmissions in bands that are critical to public safety. These potential problems appear as a major roadblock to the acceptance of CR (and SDR) technologies by regulatory agencies.

Many security vulnerabilities (such as those discussed in the prior subsection) are the direct result of poorly tested and verified code. For example, malicious inputs can exploit vulnerable software unless the software has guards against such inputs. Another example of vulnerable software is software that allows a hacker to exploit buffer overflows by carefully crafting a message that overwrites the allowable buffer space. Without aggressive testing and verification of the software, such security vulnerabilities may leave a CR as an easy target for an attacker.

To check CR code for correctness, we have applied a verification approach using aggressive program slicing and a proof-based abstraction-refinement strategy. In this approach, a high-level model is constructed from the static single assignment representation of the program. Program slicing is performed to reduce the initial model. Aggressive abstractions are further applied to reduce the verification cost. An example of this model construction is illustrated in Fig. 17. In particular, an underapproximation model is constructed in which every free variable is


1. $x_2 = x_1 + y_1$;

2. if $(x_2 \mathrel{!=} 1)$

3. $x_3 = 2$;

4. $\#x_4 = Phi \langle x_2, x_3 \rangle$

5. assert$(x_4 \mathrel{<=} 3)$;

(a)   (b)

Fig. 17. *(a) Example program code and (b) associated constructed model.*

assigned an encoding size that is typically much smaller than the bitwidth of the original data type. This model is rigorously and formally verified. Whenever the verification engine cannot draw conclusions from this underapproximate model, the model must be refined by increasing the value ranges of the variables. We have also explored overapproximation to allow for certain internal variables to be free variables. Such an abstraction reduces the constraints between internal variables and may allow the problem to be more easily handled by the prover. However, when no conclusion is obtained, refinement is needed. For CR applications, several orders of magnitude speedups were possible in proving the correctness of the CE program when compared with conventional formal techniques. These techniques are described in more detail in [77].

Furthermore, in order to enhance the reliability and to meet the stringent requirement that radio systems only transmit legal waveforms, we have designed, tested, and formally verified a software mask verifier to guarantee the legality of the output from the embedded software system with respect to regulatory requirements. This mask verifier is used to check that every waveform to be transmitted lies within legal ranges. However, the correctness of the mask verifier is critical, as it serves as the guard between the CE and the outside world. To verify the mask verifier, a hybrid testing and verification framework including both unit testing and formal verification was used to validate the correctness of the software mask. The design and verification of the mask verifier is described in [78].

VII. COGNITIVE NETWORKING

As previously noted, CRs are expected to deliver seamless adaptation, opportunistic use of underutilized spectrum, and increased flexibility in modulation and waveform selection to better fit the current wireless environment. Such intelligent radios, when placed in a network, however, may bring

about unexpected and undesirable results (e.g., adaptation cycles and local optimizations that do not translate into end-to-end performance improvements) unless network considerations are carefully explored.

A CN is a network that is capable of intelligently optimizing the end-to-end performance of a network. Fundamental aspects of this optimization include learning and reasoning. Some of the possible techniques for distributed learning and reasoning in a CN environment are examined in [79].

An architecture for CNs is shown in Fig. 18. Adaptations performed by individual nodes are driven by end-to-end goals such as creating a connected topology, maximizing network throughput, or maximizing spectral efficiency of the network. Each individual radio, however, can only partially impact such network-wide objectives. The radio can more directly impact local performance objectives, such as SINR or channel capacity. A cognitive specification language must be developed to translate end-to-end goals into objective functions the radios can understand. Individual radios running the cognitive elements shown in the figure have control over parameters such as frequency of operation, transmit power, waveform selection, and multiple-antenna operation. The CE is responsible for determining how parameters should be set to accomplish the specified network goals. An application programming interface (API) allows the CE to dynamically interact with the radio, setting parameters through the appropriate device drivers or radio operating system.

Dynamic spectrum access is one of the primary motivations for CRs. A CN may build on these DSA capabilities to form a more efficient or robust network. For instance, one application of CNs is spectrum-aware topology control, where a self-organizing network seeks to form a connected topology while maximizing spectrum reuse or minimizing interference with detected primary spectrum users.
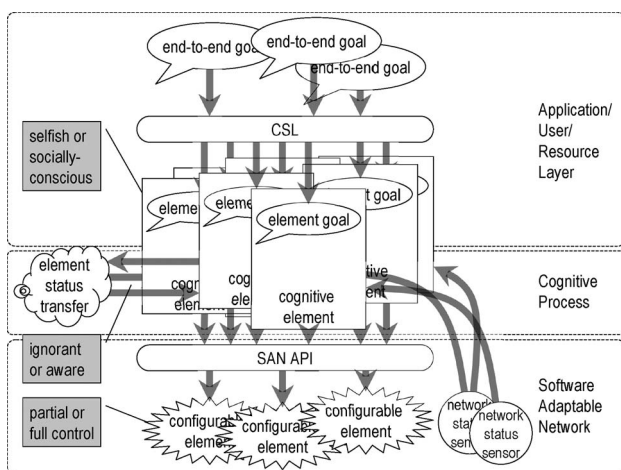

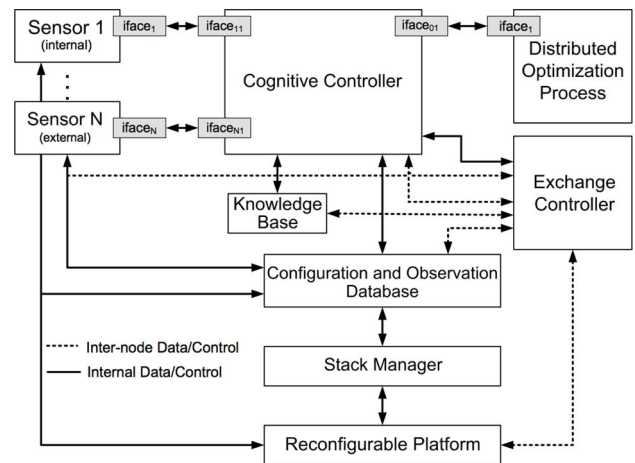
Fig. 18. *A CN architecture [28].*

Fig. 19. *CN node architecture [81].*

We have applied the concept of CNs to the problem of spectrum-aware topology control, considering network nodes that are capable of dynamically changing transmit power and frequency of operation. In our approach, the cognitive element in each radio is responsible for two core processes: one that selects interference-free transmission channels and another that seeks to maximize the lifetime of the topology by minimizing maximum transmission power. The cognitive elements also adapt to radios joining and leaving the network. Using a game theoretical model, we were able to show that these adaptations converge, across the network, to a stable and efficient state. We have also shown that individual nodes can effectively perform these adaptations, which chase network-wide goals, without having full information about the network state. This is critical, since in many cases it is infeasible to collect and propagate network-wide information (such as current topology) and radios may need to adapt with only local information about network conditions in their immediate neighborhood. Details are available in [80].

In other work, we have explored the feasibility of using distributed reasoning based on an island GA for solving a unique DSA channel allocation problem [81]. In the process, we have created a cognitive node architecture, shown in Fig. 19, that builds on our work with CEs for CR. To that work, we add a reconfigurable platform with a network stack manager based on [19] that allows the CE to reconfigure the stack for different network conditions. In addition, we believe that the complexity and overhead introduced by the multinode nature of the CN problem necessitates the addition of a database, which we denote the *configuration and observation database* to maintain information about network status and node configuration as developed from both locally observable information and data reported by other nodes. In addition, we have added an *exchange controller*, which offloads communication and management overhead from the CE. The CE sets policies

for exchanging observations with other nodes and the exchange controller executes and enforces those policies. For example, if the CE requests that updated routing table information from each of its one-hop neighbors be obtained every 10 s, the exchange controller will fulfill this request and store the results in the configuration and observation database.

We have applied this architecture to solve a channel allocation problem for DSA. Specifically, we have defined a unique DSA problem in which a multichannel ad hoc network seeks to assign channels to each of its links. When conflicting assignments are made, the links must time-share access to the channel via a medium access control protocol, such as carrier sense multiple access. The objective of the CN, then, is to maximize the sum of the throughput over all links in the network. This goal is achieved by generating a channel assignment plan with few interfering links. Through the application of the architecture described above with an island GA-based CE, we have been able to find channel assignments that achieve sum throughput within 1% of the optimal for network sizes up to 100 nodes. While this initial work, described in [81], focused on a scenario in which all nodes had complete information, our early investigations with an imperfect information version of the algorithm also show promise.

The end result of the cognitive process undertaken by autonomous nodes in the network will depend, among other factors, on whether node behavior is selfish or altruistic, on how much information about the overall state of the network is available at decision time, and on how much control the CE has on the radio operation parameters. It is critical to quantify the impact of each of these design decisions. For instance, one would like to assess the distance between a globally optimal solution and one obtained through adaptations made by CRs that possess only local information about the network. We quantify that impact, which we call the price of ignorance, for an example CN application in [82]. Similarly, we can (and do, for the same application) assess the price of selfishness and the price of partial control. Each of these may significantly affect the state to which adaptations in a CN will ultimately converge.

Research on CNs has made great strides in the past two years, building on the advances in CRs as well as on past work on cross-layer optimization. New architectures have been proposed: in addition to our reference architecture, discussed above, [19] and [83] are other examples. Some analytical models have been developed to analyze CNs, for example, using classical optimization theory or game theory. And, of course, simulation results have been reported for various CN solutions. Many challenges, however, remain. Among those, we highlight the following.

- *Cognitive specification languages.* Expressive languages must be developed to represent network objectives, to allow the mapping of such objectives into cognitive element goals, and to describe network element capabilities (network knobs) and sensed parameters (network meters).
- *Standard APIs.* Standard APIs must emerge to allow the development and reusability of CEs, controlling CRs with diverse adaptation capabilities.
- *Experimentation platforms and testbeds.* While analysis and simulation are logical first steps in the development of CNs, it is critical that experimental network testbeds using CR platforms be deployed to test the effectiveness of CN solutions in real wireless environments.
- *CE development.* This is an open issue for both CRs and CNs, and it must consider the effectiveness of different machine learning techniques, the processing and storage limitations of CR platforms, and the adaptation speed requirements of wireless environments and applications.

VIII. GAME THEORETIC ANALYSIS OF COGNITIVE RADIOS AND COGNITIVE NETWORKS

Game theory is a set of analytical tools used for analyzing the interactions of autonomous agents. Since a major aim of work in CRs and CNs is to endow radio nodes with the ability to behave autonomously, it is not surprising that game theory is a useful tool for analyzing the interactions of such nodes. The use of game theory to analyze the interactions of nodes in a network actually predates the study of CRs and CNs by many years, dating back to at least the 1970s [84]. We recently surveyed the applications of these techniques to the analysis of wireless ad hoc networks, including applications to power control, waveform adaptation, medium access control, routing, and packet forwarding [85].

A noncooperative game is a mathematical object consisting of a set of players, a set of actions available to each player, and a utility function for each player to express the player's preferences over all action tuples. The solution of such a game, called a *Nash equilibrium* (NE), is an action tuple such that no player can benefit by unilaterally deviating from the specified action. Often, a game will have many NEs, and there is no guarantee that any of them will be Pareto optimal. Moreover, in general there is no guarantee that a simple distributed algorithm will converge to a NE at all. Hence, work in applying game theory to problems in wireless networks usually focuses on 1) showing the existence (and, in some cases, uniqueness) of Nash equilibria for a given game model, 2) demonstrating an algorithm that will converge to a Nash equilibrium, and 3) either showing that the Nash equilibrium (or equilibria) is reasonably efficient or providing a mechanism to entice players to move to a more efficient operating point. An introduction to game theory focusing on wireless applications is available in [86]. Here, we review two applications of the theory in CR-like environments.

A. Analysis of Power and Coding Adaptations

A key aspect of waveform adaptation is the selection of a basic adaptation criterion. In game theory, the objective function that reflects an individual user's preferences is known as a *utility function*. Clearly, the effective per-user throughput should be a component of the utility function of each node. However, the attempt to maximize throughput at any cost will likely result in each node consuming maximum power, since throughput increases monotonically with power. Such an objective will also create excessive interference, leading to performance degradation for all nodes in the network, since throughput decreases monotonically with interference. Hence, nodes must employ a utility function that reflects multiple objectives: nodes wish to maximize throughput while minimizing power consumption.

For this application of game theory, we choose the following utility function:

$$u_i(\mathbf{p}, r_i) = L_i(\gamma_i(\mathbf{p}), r_i) - c_i(p_i).$$

Here, L_i represents node i's obtained throughput with an SINR of γ_i and a rate r_i and c_i represents the cost to node i of transmitting at power p_i. We model the throughput using a sigmoid function

$$L_i(\gamma_i, r_i) = \frac{\alpha r_i}{1 + e^{-\lambda r_i(\gamma_i - \delta r_i)}}.$$

We model the cost function as

$$c_i(p_i) = K p_i^q$$

where K and q are positive constants. In this paper, we set $K = 1$ and $q = 2$. We use $q > 1$ to reflect that players have a greater aversion to high power levels than would be reflected by a linear cost function. For a more detailed explanation, see [87] and [88].

In our simulations, we have modeled a general packet radio service (GPRS) network, since GPRS supports link adaptation. In GPRS, the modulation scheme is Gaussian minimum-shift keying, but four options for code rate are specified. The details of the simulations are described in [87] and [88].

A key result of the work is that by using game theory, the convergence of the link adaptation game (LAG) algorithm can be demonstrated. In [88], it is specifically shown that 1) the LAG defined according to the utility function described above has at least one NE and 2) the distributed LAG algorithm described in [88] (which is omitted here due to space constraints) can be shown to converge to an NE using an algorithmic mapping approach. Of course,
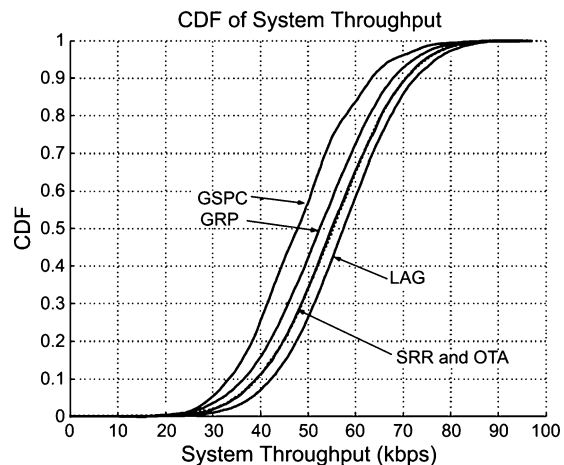

Fig. 20. *Performance comparison of GSPC-GRR, GRP, SRR, OTA, and LAG using cdf of system throughput (kbps) for the interference-limited scenario (SNR = 100 dB). Note that the performances of SRR and OTA are nearly identical [88].*

there is no guarantee that the NEs represent desirable states. To examine this, the algorithm was compared to four well-known approaches: optimal target assignment (OTA), stepwise rate removals (SRR), generalized selective power control with gradual rate removals (GSPC-GRR), and greedy rate packing (GRP). These techniques are described in detail in [88], but of primary importance is that OTA and SRR are centralized whereas GRP and GSPC-GRR are (like the proposed approach) distributed.

We refer to the sum of per-user throughputs as the *system throughput*. To obtain an unbiased performance comparison between the techniques, we conducted several simulations with random mobile locations within a seven-cell configuration. In addition, the path gain on each link also included a log-normal shadowing component. In Fig. 20, we plot the empirical cumulative distribution function (cdf) of the downlink system throughput for GSPC-GRR, GRP, SRR, OTA, and LAG. The improvement in system throughput achieved by LAG is clearly evident. The nearest competitors are OTA and SRR, which are centralized algorithms, while LAG is distributed.

B. Interference Avoidance

The previous section assumed that all users were in the same frequency band. Spatial separation guaranteed reasonable SINR values; thus frequency band adaptation was unnecessary. In many envisioned scenarios, nodes occupy the same geographical area and can cause significant interference to each other. In this case, allowing the frequency band or waveform to adapt is necessary for good network performance. Some of our recent work in [89]–[92] considers adapting the transmit waveform over a number of orthogonal signaling dimensions. Note that this is a generalization of simply choosing a single transmit frequency band, which is equivalent to DSA.

In this paper, the utility function of the kth node is defined as

$$u_k = -s_k^T X_k s_k = -s_k^T \left(\frac{\gamma_k R_k}{g_{kk}^2} + p_k \sum_{j \neq k} \frac{s_j s_j^T g_{kj}^2 \gamma_j}{g_{jj}^2} \right) s_k$$

where s_k is the transmit waveform in vector notation, p_k is the received power, g_{kj} is the channel gain from the kth node to the jth node, X_k is the matrix of interference waveforms, and γ_k is the target SINR, all of the kth user.

The basic algorithm uses a two-step process that chooses s_k as the inferior eigenvector of X_k and then chooses the power level p_k to satisfy the SINR requirement γ_k. The simple implementation of this approach requires communication between all nodes to convey the transmit waveforms and powers. However, we have investigated reduced feedback approaches (termed gradient iterations) that require minimal (as little as one bit) feedback per iteration. The work also considered a two-stage approach where the transmit power levels were readjusted to avoid overachieving SINR targets.

We have used game theory to show that the chosen utility function guarantees convergence to a NE. Moreover, simulations show that the achieved equilibria are desirable points [92]. As expected, an eigeniteration approach provides faster convergence, at the cost of considerably greater feedback than a gradient iteration approach.

IX. APPLICATIONS

A. Dynamic Spectrum Access

Of the many applications of CR, the most recognized is DSA. DSA approaches can be broken down into two basic categories: open access and hierarchical access [93], [94]. Open access puts all users on equal footing, provided that users obey specific rules, similar to current unlicensed band usage. Hierarchical access, on the other hand, views the spectrum as having a primary, licensed user and secondary sharing users. Many algorithms have been proposed for maximizing the capacity of both approaches.

Our work differs from existing work in that we do not focus on specific algorithms for spectrum sharing, but rather on the fundamental differences between hierarchical approaches: namely, spectral overlay based on interference avoidance and spectrum underlay based on interference averaging.

1) Comparison of Underlay and Overlay Approaches to DSA: In addition to choosing a vacant spectrum band on the basis of the spectrum sensing described in the previous section, other approaches to sharing spectrum include spread spectrum and band notching. This section describes our work investigating the fundamental behavior of three spectrum sharing approaches: interference averaging (i.e., spread spectrum), interference avoidance (IA) (i.e., choosing a vacant band), and interference averaging with IA included (e.g., band notching, adaptive hopping, etc.).

We investigate these approaches from two perspectives: 1) the impact of DSA on existing (i.e., legacy) radios that cannot adapt their frequency band and 2) the capacity of DSA networks in terms of sum rate. In this summary, we will concentrate on the impact of DSA on existing radios, although the relative performance of various spectrum sharing approaches carries over to network capacity as well.

The impact of DSA on legacy systems is considered through the metric of *outage probability,* the probability that DSA radios cause the legacy receiver to experience an SINR below a desired threshold. For sensing-based DSA transmitters, a single DSA radio can cause an outage at the legacy radio when sensing errors occur or when the receiver is hidden. Additionally, the sum interference from multiple DSA radios can cause an outage even if none of the DSA radios alone causes an outage.

In examining the outage probability of a legacy receiver due to DSA, we consider three basic approaches.

1) An overlay approach in which a node senses the environment and choses a band that appears to be unoccupied. This is the approach to DSA most often associated with CR and can be termed IA.

2) An underlay approach where a node simply spreads its signal over the entire available bandwidth. This is the classic spread-spectrum approach. Although generally effective, it proves to be difficult in near–far scenarios.

3) An underlay approach with IA where the node spreads its signal over the entire available bandwidth while avoiding bands where it senses transmissions. This approach combines interference averaging with IA by using either band notching (e.g., OFDM-UWB) or adaptive frequency hopping.

Under the assumption of perfect sensing, the exact distributions of the interference seen by the legacy receiver in the presence of each DSA approach are intractable. However, it is possible to calculate the cumulants of the interference in each case. Using a log-normal model for the interference and matching the cumulants of the distributions to the parameters of the log-normal distribution, we can approximate the outage probability of all three approaches.

Examining the cumulants of the interference for IA-based overlay as compared to spread spectrum–based underlay, it is found that the ability of IA to avoid nearby interferers has a dramatic impact on both the first and second cumulants. However, interference averaging (spread spectrum) reduces the power of all interferers, not just the nearby interferers. This particularly improves the second cumulant of the interference. Overall, the IA approach is superior in terms of outage probability, but the analysis shows that interference averaging does provide
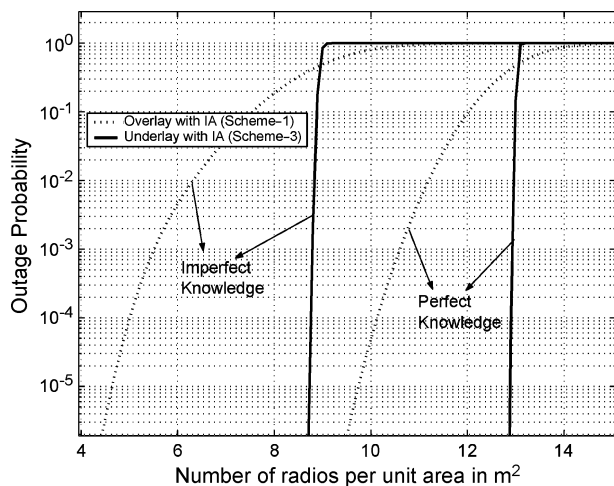
Fig. 21. *Outage probability of legacy receiver in the presence of DSA radios with overlay or underlay spectrum sharing techniques ($N_b = 512$) [95], [96].*

some desirable benefits. Thus, an approach that can combine the two techniques could harness both types of improvement, leading to the third approach.

To demonstrate the advantage of combining the two techniques as compared to IA alone, the outage probability of the legacy receiver under the first (overlay) and third (underlay with IA) approaches is presented in Fig. 21. First, we can see that with perfect sensing, underlay with IA provides a 20% improvement in DSA node density. Also shown is the outage probability with imperfect sensing. The improvement of underlay with IA increases to 30% when imperfect sensing is done at the receiver and 50% when imperfect sensing is done at the transmitter. Underlay with IA minimizes mistakes by averaging them over all available bands. When log-normal shadowing is included, the advantages to the approach are magnified. See [95] and [96] for details.

The general message of this work is that while IA is clearly desirable, it is beneficial to include wide-band transmissions that average signal power over multiple bands to minimize the impact of sensing errors and cumulative interference.

2) Spectrum Sharing in a CN: We have also studied optimal DSS in a (multihop) CN setting. In such a network, each node has a set of spectrum bands available for use; these bands may be further divided into subbands for transmission and reception. A set of source–destination pairs, each having certain rate requirements, representing the QoS requirements of user sessions in the network, is given. We study the problem of performing spectrum allocation, scheduling, and multihop multipath routing in order to minimize the required network-wide spectrum usage.

To formulate the problem mathematically, we model behaviors and constraints from multiple layers focusing on

spectrum sharing and (uneven) subband division, scheduling and interference modeling, and multipath routing. We formulate an optimization problem with the objective of minimizing the required network-wide spectrum usage for a given set of rate requirements. Since such a problem can be characterized as a mixed-integer nonlinear program (MINLP) (a class of problems which are, in general, NP-hard), we develop an approximation algorithm to find good (but possibly suboptimal) solutions.

Our approximation algorithm to the MINLP is based on a novel sequential fixing (SF) solution procedure where the determination of integer variables is performed iteratively through a series of linear programs. Once the integer variables have been fixed, other variables in the optimization problem can be solved with a linear program. We then compare the solution that is obtained by this SF procedure with a lower bound on the objective obtained through linearization and relaxing the integer variables.

The efficacy of the algorithm is shown through its application to a 20-node network containing five active sessions, each with a rate requirement. We assume that there are five total bands available to the network, with a subset of these five bands available at each node. We compare the normalized cost obtained by the SF algorithm to the lower bound cost for 100 data sets. The average normalized cost was 1.04 with a standard deviation of 0.07. Hence, the results obtained through the SF algorithm are close to the lower bound, suggesting that: 1) the lower bound is tight and 2) the solution obtained by the SF is close to the optimum. More details on this problem formulation and its solution can be found in [97].

B. Cognitive Multiple Input Multiple Output (MIMO)

A goal of a CR is to determine the most effective mode of transmission to achieve its objectives. These modes might include setting knobs such as transmit power, frequency band, modulation type, and coding scheme. The proliferation of multiple antenna nodes opens up another dimension to the problem. Multiple antennas can be used to a) improve SNR, b) provide diversity, c) introduce an extra signaling dimension (i.e., spacial multiplexing), and d) mitigate interference. The "optimal" use of the available antenna resources at the transmitter and receiver clearly depends on the channel and the nodes' objectives. Additionally, the exploitation of multiple antennas requires the existence of multiple channels in the node radios. These multiple channels can either be used in the same frequency band (i.e., for exploiting the multiple antennas) or be used to simply transmit over multiple frequency bands, ignoring the potential benefits of the multiple antennas. Additional details on the work described in this section can be found in [98].

The challenges of exploiting multiple antennas in a CR system are manifold. First, by adding the spatial dimension

to the available transmission techniques, one increases the size of an already immense parameter space. Secondly, the fact that the multiple antennas could be used for transmission of multiple, parallel single-antenna signals increases the number of possible techniques even more. Thirdly, the performance of various MIMO techniques depends heavily on the long-term and short-term matrix channel as well as the level and quality of the information available concerning the channel. Given these facts, deriving an adaptation technique for choosing the best transmission scheme becomes problematic. Thus, a role for cognition is to determine the best mapping between the set of observable parameters and the available transmission schemes. We use the term *cognitive MIMO* to refer to techniques that consider these issues.

Adaptive MIMO uses the concept of adaptive modulation and coding and adds the MIMO scheme as an extra degree of freedom. It falls short of cognitive MIMO in that it does not *learn* the best technique, but rather assumes that the best mapping is fixed based on certain channel metrics. This adaptation can be thought of as the inner cognition loop of a cognitive MIMO system (see Fig. 22). The key to adaptive MIMO is developing metrics that can predict the performance of the jointly configured system by relating the MIMO scheme performance to measurable channel parameters. The channel must be measured and parameterized in order to quantify diversity mechanisms and their rate of change, thus determining suitable schemes and their relative performance.

As an example, consider two MIMO schemes with channel knowledge available at the transmitter: 1) a fully adaptive MIMO scheme that can adapt MIMO technique, modulation scheme, and code rate and 2) a spatial multiplexing scheme that uses adaptive modulation and coding. The spectral efficiency of the two techniques for an example channel is shown in Fig. 23. In the figure, we choose the modulation/coding scheme (and MIMO technique for adaptive MIMO) so that spectral efficiency is maximized while still maintaining a bit error rate below 10^{-3}. Fig. 23(a) shows the SNR and eigenspread of a 4×4 matrix channel versus time, while Fig. 23(b) presents the achieved spectral efficiency of the two schemes.

Adaptive MIMO provides extra options via the spatial dimension, which allows spectral efficiency to be im-
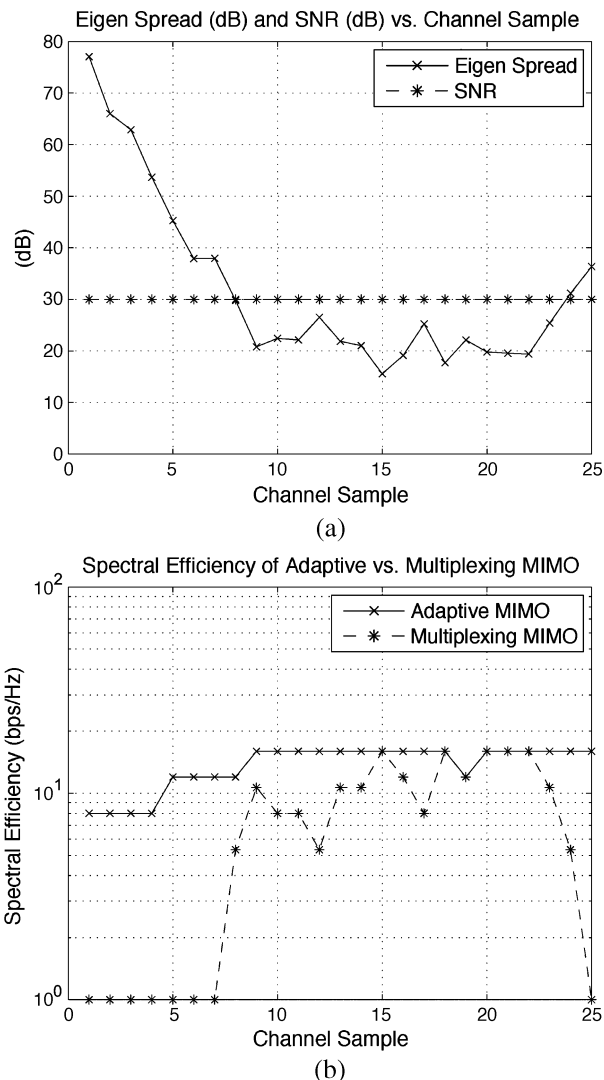

(a)

(b)

Fig. 23. *(a) Example SNR and eigenspread and (b) spectral efficiency of a time-varying channel of adaptive MIMO and spatial multiplexing with adaptive modulation/coding (target error rate of 0.1%, perfect transmitter knowledge assumed).*

proved substantially over the case where only modulation and coding can be adapted. This is particularly true when the SNR or eigenspread is insufficient to support spatial multiplexing. In these cases, the spatial dimension should be used to achieve array gain (e.g., beamforming) and spectral efficiency should be obtained via the modulation technique. On the other hand, when the eigenspread allows, we can increase spectral efficiency by using spatial multiplexing. Adaptive MIMO provides the flexibility to accomplish this.

The main assumption in the above example is that the transmitter, the receiver, or both know the channel. Using this information, a decision can be made as to the appropriate MIMO technique. Several approaches have been proposed in the literature to adapt the MIMO

Fig. 22. *Operation of a cognitive MIMO system [98].*

technique, including adapting modulation and transmit power on each eigenmode of the channel (assuming channel state information is available at the transmitter) [99]; adapting modulation, coding, and MIMO scheme using measurements at the receiver and a performance lookup table [100], [101]; and adapting the modulation and number of antennas used in a spatial multiplexing scheme [102], [103]. Although these approaches assume various levels of knowledge of the channel state, they each assume a known relationship between the channel measurement and the optimal MIMO scheme. However, in reality, channel information will be noisy and time-limited. This means that the mapping between what is measured and the optimal MIMO technique is no longer easily determined. In fact, when the channel is time-varying with an unknown distribution, predicting performance can be highly problematic. Thus, we need a technique to learn which measurements are correlated to performance and how these measurements should be mapped to the appropriate MIMO approach. A technique that includes such learning will meet our definition of *cognitive*.

The cognition cycle begins with the radio recording observations about the performance of a used MIMO/modulation/coding scheme and the associated channel conditions. These observations are used by the CE to update its knowledge base. The CE will then run its reasoning process with consideration of the new data. The CE may then choose to modify the rules used by the adaptive MIMO scheme.

Consider the case where we parameterize the channel by SNR and maximum antenna correlation. This formulation makes each measurement pair a search problem that can be solved easily by employing a CE based on a GA, such as that described in Section III. The CE operates as follows.

1) The channel conditions are observed.
2) Those schemes that were used successfully under similar conditions are used to initialize the GA's population.
3) The GA returns the best scheme based on its search.
4) The new scheme is implemented and stored in the knowledge base.

The performance of the CE was evaluated by varying the SNR over the range 0–39 dB and the neighboring antenna correlation over the range 0.1–0.9. The performance under those channel conditions was evaluated by means of simulation. The spectral efficiency of the mapping (versus SNR) determined by the CE is shown in Fig. 24 for a maximum antenna correlation of 0.1. Compared to the optimal case (found by brute force), the CE did not always learn the best mapping, but the performance was extremely close. This is due to the fact that the GA exploits the knowledge of the performance of schemes at nearby measurement pairs.


Fig. 24. *Spectral efficiency versus SNR using the optimal rule and the learned relationship (maximum pairwise correlation of 0.1) [98].*

C. Parallel SISO

Another means of utilizing multiple antennas in a CR is to transmit multiple parallel information streams in different frequency bands, a technique termed parallel single input single output (SISO) [104]. Although parallel SISO reduces spectral efficiency, it can improve performance, an important consideration in constant rate applications.

To understand this better, consider a system with N_f available frequency bands. If the transceiver has N available antennas (at both ends of the link), it can use all of the N antennas in one of the N_f bands with a traditional MIMO approach (e.g., beamforming, spatial multiplexing, diversity). Such an approach provides selection diversity of order N_f but also provides whatever benefits afforded by the MIMO technique (SNR gain, rate gain, or diversity gain).

On the other hand, a parallel SISO approach chooses to transmit N parallel signals in N separate frequency bands. On first blush, this appears to be a losing proposition since there is no diversity benefit provided to any of the parallel links. However, selection diversity is available since only N of the N_f available bands are used. Additionally, the constellation size can be reduced since only $1/N$ of the information is transmitted per band.

An example is shown in Fig. 25, where $N = 2$, $N_f = 4$, and the rate is 4.5 bps. Two approaches to single-band MIMO are shown: antenna selection and space-time block coding. The single-band MIMO approaches assume the use of 64-QAM and rate 3/4 convolutional coding (with a constraint length of eight), while the multiband parallel SISO approach uses 8-PSK with the same convolutional code in each band. Note that the parallel SISO approach provides 1–3 dB of performance improvement, albeit at a loss in spectral efficiency. The use of multiple antennas in

Fig. 25. *Performance comparison of three approaches to multiple antenna use in DSA systems: multiband parallel SISO, single-band space-time block coding, and single-band antenna selection diversity.*

systems with DSA remains an open research area. Some initial results can be found in [104].

X. CONCLUSIONS AND FUTURE RESEARCH

By describing the efforts of a team of CR and CN researchers with Wireless @ Virginia Tech, we have demonstrated the breadth of approaches required to address the challenges of CRs and CNs. To achieve the promise of CRs and CNs, researchers with a wide range of expertise must work together to create complex, coordinated systems. In particular, a research team must address the application-specific artificial intelligence challenges at the heart of a cognitive system without neglecting the myriad of radio hardware, algorithmic, networking, and analysis challenges raised by the creation of autonomous radio and network nodes. Beyond the technical challenges, impacts of policy considerations and the needs of radio system users must also be kept close at hand.

Our research to date demonstrates such a team-based coordinated-system approach. Our work in creating a CE at the core of our radio systems is supplemented by work on supporting technologies such as spectrum sensing and REMs, which provide the core with environmental awareness. Work on reconfigurable, multiband, multimode radio hardware provides the CE with greater ability to actuate its environment, a key characteristic of autonomous systems. By using game theory to understand the interactions of these autonomous agents and networking to design protocols to facilitate these interactions, we begin to address the interactions of these new cognitive devices with each other and with the world around them. Through early investigations of security and verification, we help to assure the acceptance of the new technology by both end users and regulators.

While much has been accomplished towards the creation of CRs and CNs, there is still much to do. Standardized interfaces between the pieces of a cognitive system will help to facilitate the further development of cognitive systems, much as language facilitates clearer thought. The cross-layer nature of CRs and CNs makes the development of simulation and emulation tools difficult; as a result, such tools are extremely immature. A lack of progress in analytical methods makes it difficult to clearly understand the engineering tradeoffs faced by the designer of a CR or CN. Robust, affordable, and flexible SDR and software adaptable network platforms are desperately needed to assure continued research progress.

We plan to address some of these issues and support the continued integration of work within our team through the creation of a CR and CN testbed. Preliminary plans call for placing nodes throughout a new building currently under construction on the Virginia Tech campus. Each node will consist of a computer, a universal software radio peripheral (USRP), and an RF front-end built around the Motorola chip described in Section V-C. Nodes will be remotely reconfigurable and able to run software built around any software architecture that supports the USRP, including GNU radio and the Open Source SCA Implementation-Embedded (OSSIE). We believe that this testbed will be an important key to the continued development of CRs and CNs by Wireless @ Virginia Tech. ∎

REFERENCES

[1] J. Mitola, "Cognitive radio: Model-based competence for software radios," Ph.D. dissertation, Dept. of Teleinformatics, KTH, 1999.

[2] "Telecommunication," 2007. [Online]. Available: http://www.access.gpo.gov/cgi-bin/cfrassemble.cgi?title=200747

[3] FCC Spectrum Policy Task Force, "Report of the Spectrum Efficiency Working Group,"

Nov. 2002. [Online]. Available: http://www.fcc.gov/sptf/reports.html

[4] F. H. Sanders and V. S. Lawrence, "Broadband spectrum survey at Denver, Colorado," NTIA Rep. 95-321, Sep. 1995.

[5] S. W. Ellingson, "Spectral occupancy at VHF: Implications for frequency-agile cognitive radios," in *Proc. IEEE Veh. Technol. Conf.*, Sep. 2005, vol. 2, pp. 1379–1382.

[6] M. A. McHenry and D. McCloskey, "Multi-band, multi-location spectrum occupancy measurements," in *Proc. Int. Symp. Adv. Radio Technol.*, Boulder, CO, Mar. 2006.

[7] R. W. Thomas, L. A. DaSilva, and A. B. MacKenzie, "Cognitive networks," in *Proc. 1st IEEE Int. Symp. New Frontiers Dyn. Spectrum Access Netw. (DySpan)*, Nov. 2005, pp. 352–360.

[8] I. A. Akbar, "Statistical analysis of wireless systems using Markov models," Ph.D. dissertation, Virginia Polytechnic Inst. and State Univ., Blacksburg, Jan. 2007.

[9] M. Mohammad, "Cellular diagnostic systems using hidden Markov models," Ph.D. dissertation, Virginia Polytechnic Inst. and State Univ., Blacksburg, Oct. 2006.

[10] B. Le, "Building a cognitive radio: From architecture definition to prototype implementation," Ph.D. dissertation, Virginia Polytechnic Inst. and State Univ., Blacksburg, 2007.

[11] G. J. Minden, J. B. Evans, L. Searl, D. DePardo, V. R. Petty, R. Rajbanshi, T. Newman, Q. Chen, F. Weidling, J. Guffey, D. Datla, B. Barker, M. Peck, B. Cordill, A. M. Wyglinsky, and A. Agah, "KUAR: A flexible software-defined radio development platform," in *Proc. IEEE Int. Symp. New Frontiers Dyn. Spectrum Access Netw. (DySpan)*, 2007, pp. 428–439.

[12] D. Raychaudhuri, N. B. Mandayam, J. B. Evans, B. J. Ewy, S. Seshan, and P. Steenkiste, "Cognet: An architectural foundation for experimental cognitive radio networks within the future internet," in *Proc. 1st ACM/IEEE Int. Workshop Mobility Evolv. Internet Architect. (MobiArch '06)*, New York, 2006, pp. 11–16.

[13] Z. Miljanic, I. Seskar, K. Le, and D. Raychaudhuri, "The WINLAB network centric cognitive radio hardware platform," in *Proc. Int. Conf. Cogn. Radio Oriented Wireless Netw. Commun. (CrownCom)*, Orlando, FL, Aug. 2007, pp. 155–160.

[14] R. Mukhopadhyay, Y. Park, P. Sen, N. Srirattana, J. Lee, C.-H. Lee, S. Nuttinck, A. Joseph, J. D. Cressler, and J. Laskar, "Reconfigurable RFICs in Si-based technologies for a compact intelligent RF front-end," *IEEE Trans. Microwave Theory Tech.*, vol. 53, pp. 81–93, Jan. 2005.

[15] G. Ganesan and Y. G. Li, "Cooperative spectrum sensing in cognitive radio networks," in *Proc. IEEE Dyn. Spectrum Access Netw. (DySPAN)*, Baltimore, MD, Nov. 2005, pp. 137–143.

[16] K. R. Chowdhury and I. F. Akyildiz, "Cognitive wireless mesh networks with dynamic spectrum access," *IEEE J. Sel. Areas Commun.*, vol. 26, pp. 168–181, Jan. 2008.

[17] H. So, A. Tkachenko, and R. W. Brodersen, "A unified hardware/software runtime environment for FPGA based reconfigurable computers using BORPH," in *Proc. Int. Conf. Hardware-Software Codesign Syst. Synth.*, Seoul, Korea, Oct. 2006.

[18] S. M. Mishra, A. Sahai, and R. W. Brodersen, "Cooperative sensing among cognitive radios," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Istanbul, Jun. 2006.

[19] P. Sutton, L. Doyle, and K. Nolan, "A reconfigurable platform for cognitive

networks," in *Proc. 1st Int. Conf. Cogn. Radio Oriented Wireless Netw. Commun. (CROWNCOM 2006)*, Jun. 2006, pp. 8–10.

[20] J. Mitola, III and G. Q. Macguire, Jr., "Cognitive radio: Making software radios more personal," *IEEE Pers. Commun.*, vol. 6, pp. 13–18, Aug. 1999.

[21] J. Boyd, *Patterns of conflict,* Dec. 1987, briefing. [Online]. Available: http://www.d-n-i.net/second_level/boyd_military.htm

[22] C. W. Bostian, S. F. Midkiff, W. M. Kurgan, L. W. Carstensen, D. G. Sweeney, and T. M. Gallagher, "Broadband communications for disaster response," *Space Commun.*, vol. 18, no. 3–4, pp. 167–177, 2002.

[23] C. J. Rieser, T. W. Rondeau, C. Bostian, W. R. Cyre, and T. M. Gallagher, "Cognitive radio engine based on genetic algorithms in a network," U.S. Patent 7 289 972, Oct. 2007.

[24] B. Le, T. W. Rondeau, and C. W. Bostian, "Cognitive radio realities," *Wireless Commun. Mobile Comput.*, vol. 7, no. 9, pp. 1037–1048, Nov. 2007.

[25] T. W. Rondeau, B. Le, C. J. Rieser, and C. W. Bostian, "Cognitive radios with genetic algorithms: Intelligent control of software defined radios," in *Proc. Software Defined Radio Tech. Conf.*, Phoenix, AZ, Nov. 2004.

[26] J. Neel, R. M. Buehrer, J. H. Reed, and R. P. Gilles, "Game theoretic analysis of a network of cognitive radios," in *Proc. Midwest Symp. Circuits Syst.*, Aug. 2002, vol. 3, pp. 409–412.

[27] S. Srikanteswara, J. H. Reed, P. Athanas, and R. Boyle, "A soft radio architecture for reconfigurable platforms," *IEEE Commun. Mag.*, pp. 140–147, Feb. 2000.

[28] R. W. Thomas, D. H. Friend, L. A. DaSilva, and A. B. MacKenzie, "Cognitive networks: Adaptation and learning to achieve end-to-end performance objectives," *IEEE Commun. Mag.*, pp. 51–57, Dec. 2006.

[29] B. Le, F. A. G. Rodriguez, Q. Chen, B. Li, F. Ge, M. Y. ElNainay, T. W. Rondeau, and C. W. Bostian, "A public safety cognitive radio," in *Proc. Software Defined Radio Tech. Conf.*, Denver, CO, Nov. 2007.

[30] Y. Wang, Q. Chen, B. Le, and C. W. Bostian, "Universal synchronization design for cognitive radios," in *Proc. Software Defined Radio Tech. Conf.*, Denver, CO, Nov. 2007.

[31] T. W. Rondeau, "Application of artificial intelligence to wireless communications," Ph.D. dissertation, Virginia Polytechnic Inst. and State Univ., Blacksburg, 2007.

[32] T. W. Rondeau, A. B. MacKenzie, C. W. Bostian, K. E. Nolan, L. E. Doyle, C. Doerr, D. Grunwald, G. Minden, J. Evans, and D. Raychaudhuri, "International collaboration for a cognitive radio testbed," in *Proc. Software Defined Radio Tech. Conf.*, 2007.

[33] T. W. Rondeau, B. Le, D. M. Maldonado, D. Scaperoth, A. B. MacKenzie, and C. W. Bostian, "Optimization, learning, and decision making in a cognitive engine," in *Proc. Software Defined Radio Tech. Conf.*, Orlando, FL, 2006.

[34] I. Gilboa and D. Schmeidler, *A Theory of Case-Based Decisions.* Cambridge, U.K.: Cambridge Univ. Press, 2001.

[35] W. C. Headley, J. D. Reed, and C. R. C. M. da Silva, "Distributed cyclic spectrum feature-based modulation classification," in *Proc. IEEE Wireless Commun. Netw. Conf.*, Las Vegas, NV, 2008, pp. 1200–1204.

[36] C. R. C. M. da Silva, W. C. Headley, J. D. Reed, and Y. Zhao, "The application of distributed spectrum sensing and available resource maps to cognitive radio systems," in *Proc. Inf. Theory Applicat. Workshop*, La Jolla, CA, 2008.

[37] A. Fehske, J. Gaeddert, and J. H. Reed, "A new approach to signal classification using spectral correlation and neural networks," in *Proc. IEEE Dyn. Spectrum Access Netw. (DySPAN)*, Baltimore, MD, 2007, pp. 144–150.

[38] K. Kim, I. A. Akbar, K. K. Bae, J. Um, C. M. Spooner, and J. H. Reed, "Cyclostationary approaches to signal detection and classification in cognitive radios," in *Proc. IEEE Dyn. Spectrum Access Netw. (DySPAN)*, Dublin, Ireland, 2007, pp. 212–215.

[39] W. A. Gardner, *Statistical Spectral Analysis: A Nonprobabilistic Theory.* Englewood Cliffs, NJ: Prentice-Hall, 1987.

[40] W. A. Gardner, *Cyclostationarity in Commun. and Signal Processing.* Piscataway, NJ: IEEE Press, 1994.

[41] Z. B. Tang, K. R. Pattipati, and D. L. Kleinman, "A distributed m-ary hypothesis testing problem with correlated observations," in *Proc. 28th Conf. Decision Contr.*, 1989, pp. 562–568.

[42] A. Batra, W. Krenik, and C. Panasik, "Cognitive radios for unlicensed WANs," presented at the BWRC Cognitive Radio Workshop, 2004.

[43] W. Krenik and A. Batra, "Cognitive radio techniques for wide area networks," in *Proc. Conf. Design Automation*, Anaheim, CA, 2005, pp. 409–412.

[44] Y. Zhao, J. Gaeddert, L. Morales, K. Bae, J.-S. Um, and J. H. Reed, "Development of radio environment map enabled case- and knowledge-based learning algorithms for IEEE 802.22 WRAN cognitive engines," in *Proc. Int. Conf. Cogn. Radio Oriented Wireless Netw. Commun. (CrownCom)*, Orlando, FL, 2007, pp. 44–49.

[45] Y. Zhao, L. Morales, J. Gaeddert, K. K. Bae, J. Um, and J. H. Reed, "Applying radio environment maps to cognitive wireless regional area networks," in *Proc. IEEE Dyn. Spectrum Access Netw. (DySPAN)*, Dublin, Ireland, 2007, pp. 115–118.

[46] Y. Zhao, B. Le, and J. H. Reed, "Network support: The radio environment map," in *Cognitive Radio Technology*, B. A. Fette, Ed. Amsterdam, The Netherlands: Elsevier, 2006.

[47] Y. Zhao, J. Gaeddert, K. K. Bae, and J. H. Reed, "Radio environment map-enabled situation-aware cognitive radio learning algorithms," in *Proc. Software Defined Radio Tech. Conf.*, Orlando, FL, 2006.

[48] Y. Zhao, D. Raymond, C. R. C. M. da Silva, J. H. Reed, and S. F. Midkiff, "Performance evaluation of radio environment map-enabled cognitive spectrum-sharing networks," in *Proc. IEEE Military Commun. Conf.*, Orlando, FL, 2007.

[49] U.S. Department of Homeland Security SAFECOM Program, "Statement of requirements for public safety wireless communications & interoperability, ver. 1.1," Jan. 2006.

[50] U.S. Department of Justice, National Task Force for Interoperability, "Why can't we talk? Working together to bridge the communications gap," Nat. Inst. of Justice Rep., Feb. 2005.

[51] S. W. Ellingson. (2006, Jun.). "A comparison of some existing radios with implications

for public safety interoperability," Virginia Tech, Tech. Rep. 4. [Online]. Available: http://www.ece.vt.edu/swe/chamrad

[52] P. Athanas, J. Bowen, T. Dunham, C. Patterson, J. Rice, M. Shelburne, J. Suris, M. Bucciero, and J. Graf, "Wires on demand: Run-time communication synthesis for reconfigurable computing," in *Proc. Int. Conf. Field Program. Logic Applicat. (FPL)*, Amsterdam, The Netherlands, Aug. 2007, pp. 513–516.

[53] *UG208: Early Access Partial Reconfiguration User Guide*, ver. 1.1, Xilinx, Mar. 2006.

[54] R. Lysecky and F. Vahid, "A configurable logic architecture for dynamic hardware/ software partitioning," in *Proc. Conf. Design, Automat. Test Eur. (DATE '04)*, Washington, DC, 2004, pp. 480–485.

[55] A. Bourdoux, J. Craninckx, A. Dejonghe, and L. van der Perre, "Receiver architectures for software-defined radios in mobile terminals: The path to cognitive radios," in *Proc. IEEE Radio Wireless Symp.*, Jan. 2007, pp. 535–538.

[56] M. Lehne and S. Raman, "A prototype analog/mixed-signal fast Fourier transform processor IC for OFDM receivers," in *Proc. IEEE Radio Wireless Symp.*, 2008.

[57] J. Park, T. Song, J. Hur, S.-M. Lin, J. Choi, K. Kim, J. Lee, K. Lim, C.-H. Lee, H. Kim, and J. Laskar, "A fully-integrated UHF receiver with multi-resolution spectrum sensing (MRSS) functionality for IEEE 802.22 cognitive radio applications," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, Feb. 2008, pp. 526–527, 633.

[58] R. Bagheri, A. Mirzaei, S. Chehrazi, M. Heidari, M. Lee, M. Mikhemar, W. Tang, and A. Abidi, "An 800 MHz to 5 GHz software-defined radio receiver in 90 nm CMOS," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, Feb. 2006, pp. 1932–1933, 1941.

[59] Z. Zhao and S. Raman, "A 0.18 μm RF CMOS ultra wideband transmitter front end RFIC," in *IEEE RFIC Symp. Dig.*, San Francisco, Jun. 2006, pp. 506–509.

[60] Z. Zhao, "Silicon-based RFIC transmitter front ends for ultra-wideband communications and sensor applications," Ph.D. dissertation, Virginia Polytechnic Inst. and State Univ., Mar. 2007.

[61] I. Chamas and S. Raman, "A 5 GHz I/Q phase-tunable CMOS LC quadrature VCO (PT-QVCO) for analog phase calibrated receiver architectures," in *IEEE Topical Meeting Silicon Monolithic Circuits RF Syst. Dig.*, Long Beach, CA, Jan. 2007, pp. 269–272.

[62] J. Ryynanen, S. Lindfors, K. Stadius, and K. Halonen, "Integrated circuits for multi-band multimode receivers," *IEEE Circuits Syst. Mag.*, vol. 6, no. 2, pp. 5–16, 2006.

[63] G. Cafaro, T. Gradishar, J. Heck, S. Machan, G. Nagaraj, S. Olson, R. Salvi, B. Stengel, and B. Ziemer, "A 100 MHz–2.5 GHz direct conversion CMOS transceiver for SDR applications," in *IEEE Radio Freq. Integr. Circuits (RFIC) Symp.*, Jun. 2007, pp. 189–192.

[64] R. Li, B. Pan, J. Laskar, and M. M. Tentzeris, "A compact broadband planar antenna for GPS, DCS-1800, IMT-2000, and WLAN applications," *IEEE Antennas Wireless Propagat. Lett.*, vol. 6, pp. 25–27, 2007.

[65] R. M. Fano, "Theoretical limitations on the broadband matching of arbitrary impedances," *J. Franklin Inst.*, vol. 249, Jan.–Feb. 1950.

[66] G. A. Matthaei, L. Yong, and E. M. T. Jones, *Microwave Filters, Impedance-Matching*

Networks, and Coupling Structures. Norwood, MA: Artech House, 1980.

[67] S. W. Ellingson, J. H. Simonetti, and C. Patterson, "Design and evaluation of an active antenna for a 29–47 MHz radio telescope array," *IEEE Trans. Antennas Propagat.*, vol. 55, pp. 826–831, Mar. 2007.

[68] S. M. S. Hasan and S. W. Ellingson, "Multiband antenna-receiver integration using an RF multiplexer with sensitivity-constrained design," in *Proc. IEEE Int. Symp. Antennas Propagat.*, 2008.

[69] S. E. Sussman-Fort, "Matching network design using non-Foster impedances," *Int. J. RF Microw. Computer-Aided Eng.*, pp. 135–142, Mar. 2006.

[70] T. C. Clancy and N. Goergen, "Security in cognitive radio networks: Threats and mitigation," in *Proc. Int. Conf. Cogn. Radio Oriented Wireless Netw. Commun. (CrownCom)*, May 2008.

[71] P. Flanigan, V. Welch, and M. Pant, "Dynamic policy enforcement for software defined radio," in *Proc. Annu. Simul. Symp.*, 2005.

[72] R. Chen, J.-M. Park, and J. H. Reed, "Defense against primary user emulation attacks in cognitive radio networks," *IEEE J. Sel. Areas Commun.*, vol. 26, pp. 25–37, Jan. 2008.

[73] I. F. Akyildiz, W.-Y. Lee, M. C. Vuran, and S. Mohanty, "Next generation/dynamic access/cognitive radio wireless networks: A survey," *Elsevier Comput. Netw. J.*, vol. 50, pp. 2127–2159, Sep. 2006.

[74] *IEEE P802.22/D0.1 Draft Standard for Wireless Regional Areas Networks Part 22: Cognitive Wireless RAN Medium Aceess Control (MAC) and Physical Layer (PHY) Specifications: Policies and Procedures for Operation in the TV Bands*, IEEE doc. 22-06-0067-00-0000_ P802-22_D0.1, IEEE 802.22 Working Group on Wireless Regional Area Networks, May 2006.

[75] L. B. Michael, M. J. Mihaljevic, S. Haruyama, and R. Kohno, "A framework for secure download for software-defined radio," *IEEE Commun. Mag.*, vol. 40, pp. 88–96, Jul. 2002.

[76] A. Brawerman and J. A. Copeland, "Towards a fraud-prevention framework for software defined radio mobile devices," *EURASIP J. Wireless Commun. Netw.*, vol. 2005, no. 3, pp. 401–412, 2005.

[77] N. He and M. S. Hsiao, "Bounded model checking of embedded software in wireless cognitive radio systems," in *Proc. IEEE Int. Conf. Comput. Design*, Oct. 2007, pp. 19–24.

[78] X. Cheng, N. He, and M. S. Hsiao, "Hybrid testing and verification techniques for a cognitive radio system," in *Proc. Int. Conf. Software Eng. Applicat.*, Nov. 2007.

[79] D. H. Friend, R. W. Thomas, A. B. MacKenzie, and L. A. DaSilva, "Distributed learning and reasoning in cognitive networks," in *Cognitive Networks: Towards Self-Aware Networks*, Q. Mahmoud, Ed.   New York: Wiley, 2007.

[80] R. W. Thomas, R. S. Komali, L. A. DaSilva, and A. B. MacKenzie, "Joint power and channel minimization in topology control: A cognitive network approach," in *Proc. IEEE Int. Conf. Commun.*, Glasgow, Scotland, Jun. 24–28, 2007, pp. 6538–6543.

[81] D. H. Friend, M. Y. ElNainay, Y. Shi, and A. B. MacKenzie, "Architecture and performance of an island genetic algorithm-based cognitive network," in *Proc. IEEE Consum. Commun. Netw. Conf.*,

Las Vegas, NV, Jan. 10–12, 2008, pp. 993–997.

[82] R. W. Thomas, L. A. DaSilva, K. N. Wood, and M. V. Marathe, "Critical design decisions for cognitive networks," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Glasgow, Scotland, Jun. 2007, pp. 3993–3998.

[83] P. Mähönen, M. Petrova, J. Riihijärvi, and M. Wellens, "Cognitive wireless networks: Your network just became a teenager," in *Proc. IEEE INFOCOM*, 2006.

[84] V. Nesteruk, "Basic methods for applying game theory in radio engineering and communications," *Telecommun. Radio Eng. Part 2 (Radio Eng.)*, vol. 31, no. 7, pp. 57–66, Jul. 1976.

[85] V. Srivastava, J. Neel, A. B. MacKenzie, R. Menon, L. A. DaSilva, J. E. Hicks, J. H. Reed, and R. P. Gilles, "Using game theory to analyze wireless ad hoc networks," *IEEE Commun. Surveys*, vol. 7, no. 4, pp. 46–56, 2005.

[86] A. B. MacKenzie and L. A. DaSilva, *Game Theory for Wireless Engineers.*   San Rafael, CA: Morgan and Claypool, 2006.

[87] S. Ginde, J. Neel, and R. M. Buehrer, "Game theoretic analysis of joint link adaptation and distributed power control in GPRS," in *Proc. Fall 2003 Veh. Technol. Conf.*, 2003.

[88] S. V. Ginde, A. B. MacKenzie, R. M. Buehrer, and R. S. Komali, "A game-theoretic analysis of link adaptation in cellular radio networks," *IEEE Trans. Veh. Technol.*, vol. 57, no. 5, pp. 3108–3120, Sep. 2008.

[89] J. E. Hicks, A. B. MacKenzie, J. A. Neel, and J. H. Reed, "A game theory perspective on interference avoidance," in *Proc. IEEE Globecom*, 2004, vol. 1, pp. 257–261.

[90] R. Menon, A. B. MacKenzie, R. M. Buehrer, and J. H. Reed, "A game-theoretic framework for interference avoidance in ad hoc networks," in *Proc. IEEE Global Conf. Commun.*, 2006.

[91] R. Menon, A. B. MacKenzie, R. M. Buehrer, and J. H. Reed, "Joint power control and waveform adaptation for distributed networks," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Washington, DC, Nov. 2007, pp. 694–699.

[92] R. Menon, A. B. MacKenzie, J. E. Hicks, R. M. Buehrer, and J. H. Reed, "A game-theoretic framework for interference avoidance," *IEEE Trans. Commun.*, vol. 57, no. 4, Apr. 2009, to appear.

[93] Q. Zhao and A. Swami, "A survey of dynamic spectrum access: Signal processing and networking perspective," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Apr. 2007, vol. 4, pp. 1349–1352.

[94] O. Ileri and N. B. Mandayam, "Dynamic spectrum access models: Toward an engineering perspective in the spectrum debate," *IEEE Commun. Mag.*, vol. 46, pp. 153–160, Jan. 2008.

[95] R. Menon, R. M. Buehrer, and J. H. Reed, "Outage probability based comparison of underlay and overlay spectrum sharing techniques," in *Proc. IEEE Dyn. Spectrum Access Netw. (DySPAN)*, Nov. 2005, pp. 101–109.

[96] R. Menon, R. M. Buehrer, and J. H. Reed, "On the impact of dynamic spectrum sharing techniques on legacy radio systems," *IEEE Trans. Wireless Commun.*, vol. 7, no. 11, pp. 4198–4207, Nov. 2008.

[97] Y. T. Hou, Y. Shi, and H. D. Sherali, "Optimal spectrum sharing for multihop software defined radio networks," in *Proc.*

IEEE INFOCOM, Anchorage, AK, May 2007, pp. 1–9.

[98] H. Volos, C. I. Phelps, and R. M. Buehrer, "Initial design of a cognitive engine for MIMO systems," in *Proc. Software Defined Radio Tech. Conf.*, Denver, CO, Nov. 2007.

[99] Z. Zhou, B. Vucetic, M. Dohler, and Y. Li, "MIMO systems with adaptive modulation," *IEEE Trans. Veh. Technol.*, vol. 54, pp. 1828–1842, Sep. 2005.

[100] S. Catreux, V. Erceg, D. Gesbert, and R. W. Heath, "Adaptive modulation and MIMO coding for broadband wireless data networks," *IEEE Commun. Mag.*, vol. 40, pp. 108–115, Jun. 2002.

[101] A. Forenza, M. R. McKay, A. Pandharipande, R. W. Heath, and J. B. Collings, "Adaptive MIMO transmission for exploiting the capacity of spatially correlated channels," *IEEE Trans. Veh. Technol.*, vol. 56, pp. 619–630, Mar. 2007.

[102] R. Narasimhan, "Spatial multiplexing with transmit antenna and constellation selection for correlated MIMO fading channels," *IEEE Trans. Signal Process.*, vol. 51, pp. 2829–2838, Nov. 2003.

[103] H. Zhuang, L. Dai, S. Zhou, and Y. Yao, "Low complexity per-antenna rate and power control approach for closed-loop V-BLAST," *IEEE Trans. Commun.*, vol. 51, pp. 1783–1787, Nov. 2003.

[104] C. I. Phelps and R. M. Buehrer, "Dynamic spectrum access for multi-antenna systems," in *Proc. IEEE Int. Symp. New Frontiers in Dyn. Spectrum Access Netw.*

ABOUT THE AUTHORS

Allen B. MacKenzie (Senior Member, IEEE) has been an Assistant Professor in Virginia Tech's Bradley Department of Electrical and Computer Engineering since 2003. He joined Virginia Tech after receiving his Ph.D. from Cornell University and his B.Eng. from Vanderbilt University, both in Electrical Engineering. Dr. MacKenzie's research focuses on wireless communications systems and networks. His current research interests include cognitive radio and cognitive network algorithms, architectures, and protocols and the analysis of such systems and networks using game theory. In addition to the IEEE, Dr. MacKenzie is a member of the ASEE and the ACM. In 2006, he received the Dean's Award for Outstanding New Assistant Professor in the College of Engineering at Virginia Tech.

Jeffrey H. Reed (Fellow, IEEE) is the Willis G. Worcester Professor in the Bradley Department of Electrical and Computer Engineering. From June 2000 to June 2002, Dr. Reed served as Director of the Mobile and Portable Radio Research Group (MPRG). He currently serves as Director of the newly formed umbrella wireless organization Wireless @ Virginia Tech.

Dr. Reed's area of expertise is in software radios, cognitive radio, wireless networks and communications signal processing. His book, Software Radio: A Modern Approach to Radio Design was published by Prentice Hall in May 2002. His latest book, An Introduction to Ultra Wideband Communication Systems was published by Prentice Hall in April 2005.

Dr. Reed received the College of Engineering Award for Excellence in Research in 2001. In 2005, Dr. Reed became Fellow to the IEEE for contributions to software radio and communications signal processing and for leadership in engineering education.

Peter Athanas (Senior Member, IEEE) received the B.S. degree from The University of Toledo, Toledo, OH, and the M.S. degree from Rensselaer Polytechnic Institute, Troy, NY, both in electrical engineering. He received the Sc.M. degree in applied mathematics and the Ph.D. degree in electrical engineering from Brown University, Providence, RI.

His doctoral work focused on configurable computing architectures and compilers. He is a Professor in the Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University (Virginia Tech), Blacksburg. His research interests include high-performance embedded computing, configurable computing, VLSI, and signal processing. He currently is Director of the Virginia Tech Configurable Computing Laboratory and Site Director of the NSF Center for Reconfigurable High Performance Computing.

Charles W. Bostian (Fellow, IEEE) received the B.S., M.S., and Ph.D. degrees from North Carolina State University at Raleigh in 1963, 1964, and 1967, respectively.

He is Alumni Distinguished Professor of Electrical and Computer Engineering at Virginia Polytechnic Institute and State University (Virginia Tech), Blacksburg, where he has been a Faculty Member since 1969. Prior to joining the university, he was a U.S. Army Officer and worked briefly for Corning Glassworks. Since 1993, he has been Director of the Virginia Tech Center for Wireless Telecommunications. He is also an active member of Wireless @ Virginia Tech.

R. Michael Buehrer (Senior Member, IEEE) joined Virginia Polytechnic Institute and State University (Virginia Tech), Blacksburg, from Bell Labs as an Assistant Professor with the Bradley Department of Electrical Engineering in 2001. He is currently an Associate Professor and is part of Wireless @ Virginia Tech, a comprehensive research group focusing on wireless communications. His current research interests include dynamic spectrum sharing, cognitive radio, multiple-input multiple-output communications, intelligent antenna techniques, position location networks, ultra-wide-band, spread spectrum, interference avoidance, and propagation modeling. His work has been supported by the National Science Foundation, Defense Advanced Research Projects Agency, Office of Naval Research, and several industrial sponsors.

Prof. Buehrer received the Bell Laboratories President's Silver Award for his work on intelligent antenna systems. In 2003, he was named Outstanding New Assistant Professor by the College of Engineering, Virginia Tech.

Luiz A. DaSilva (Senior Member, IEEE) received the Ph.D. degree in electrical engineering from the University of Kansas, Lawrence.

He was with IBM for six years. He joined the Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University (Virginia Tech), Blacksburg, in 1998, where he is now an Associate Professor. His current research involves experimental and game-theoretic analysis of mobile ad hoc networks, and analysis and protocol design for cognitive networks. His research has been supported by the National Science Foundation, Defense Advanced Research Projects Agency, Office of Naval Research, and Intel, among others. He is a member of the Wireless @ Virginia Tech research group.

Prof. DaSilva is a member of ASEE and ACM. In 2006, he was named a College of Engineering Faculty Fellow at Virginia Tech.

Steven W. Ellingson (Senior Member, IEEE) received the B.S. degree in electrical and computer engineering from Clarkson University, Potsdam, NY, in 1987 and the M.S. and Ph.D. degrees in electrical engineering from The Ohio State University, Columbus, in 1989 and 2000, respectively.

From 1989 to 1993, he served on active duty with the U.S. Army. From 1993 to 1995, he was a Senior Consultant with Booz-Allen Hamilton, McLean, VA. From 1995 to 1997, he was a Senior Systems Engineer with Raytheon E-Systems, Falls Church, VA. From 1997 to 2003, he was a Research Scientist with The Ohio State University ElectroScience Laboratory. Since 2003, he has been an Assistant Professor in the Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University, Blacksburg. His research interests include antennas and propagation, applied signal processing, and instrumentation.

Jung-Min Park (Member, IEEE) received the Ph.D. degree in electrical and computer engineering from Purdue University, West Lafayette, IN, in 2003.

His research interests include security issues in cognitive radio systems/networks, network attack traceback, and cryptographic algorithms and protocols. Currently, he is carrying out several sponsored projects that involve investigating various security issues in cognitive radio and conventional radio networks. His current research sponsors include the National Science Foundation (NSF), SANS Institute, Samsung Electronics, and SCA Technica.

Dr. Park received an NSF CAREER award and an AT&T Leadership award.

Y. Thomas Hou (Senior Member, IEEE) received the B.E. degree from the City College of New York in 1991, the M.S. degree from Columbia University, New York, in 1993, and the Ph.D. degree from Polytechnic Institute of New York University, in 1998, all in electrical engineering.

From 1997 to 2002, he was a Researcher with Fujitsu Laboratories of America, Sunnyvale, CA. Since August 2002, he has been with Virginia Polytechnic Institute and State University, Blacksburg, where he is now an Associate Professor. His current research interests are radio resource (spectrum) management and networking for cognitive radio wireless networks, optimization and algorithm design for wireless ad hoc and sensor networks, and video communications over dynamic ad hoc networks.

Prof. Hou received an Office of Naval Research Young Investigator Award (2003) and a National Science Foundation CAREER Award (2004).

Cameron Patterson (Senior Member, IEEE) received the B.Sc. (honors) and M.Sc. degrees from the University of Manitoba, Winnipeg, MB, Canada, in 1980 and 1983, respectively, and the Ph.D. degree from the University of Calgary, Calgary, AB, Canada, in 1992, all in computer science.

From 1992 to 1994, he held an NSERC Industrial Research Fellowship at the Alberta Microelectronic Centre, Canada. During this time he designed ASICs, developed algorithms to partition circuits over multiple field-programmable gate arrays (FPGAs), and lectured on VLSI design at the University of Calgary. From 1994 to 2003, he was a Researcher with Xilinx, the leading manufacturer of FPGAs. Since 2003, he has been an Associate Professor in the Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University, Blacksburg, where his research and teaching interests include reconfigurable, high-performance, and secure computing, and run-time algorithms for electronic design automation.

Michael Hsiao (Senior Member, IEEE) received the B.S. degree in computer engineering and the M.S. and Ph.D. degrees from the University of Illinois at Urbana-Champaign in 1992, 1993, and 1997, respectively.

Since 2006, he has been a Professor in the Electrical and Computer Engineering Department, Virginia Polytechnic Institute and State University, Blacksburg. He has published more than 160 refereed journal and conference papers and has served on the editorial boards of several journals. His current research focuses on testing, verification, and diagnosis of digital systems.

Dr. Hsiao received the National Science Foundation CAREER Award. He was recognized for the most influential papers in the first ten years (1998–2007) of the Design Automation and Test Conference in Europe. He received the Digital Equipment Corporation Fellowship and McDonnell Douglas Scholarship. He has served on the Program Committee for more than 20 IEEE international conferences and workshops.

Sanjay Raman (Senior Member, IEEE) received the bachelor's degree (with highest honors) from the Georgia Institute of Technology, Atlanta, in 1987 and the M.S. and Ph.D. degrees from The University of Michigan, Ann Arbor, in 1993 and 1998, respectively, all in electrical engineering.

Since 1998, he has been a Member of the Faculty of the Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University (Virginia Tech), Blacksburg, where he is now an Associate Professor. His research interests include RF/microwave/millimeter-wave integrated circuits and antennas; high-speed/mixed-signal ICs and SoC integration; interconnects and packaging; RF microelectromechaincal/nanoelectromechanical (MEMS/NEMS) and MEMS sensors; micromachining, solid-state technology, and nanotechnology; and integrated wireless communications and sensor microsystems. He is currently a Program Manager with the Microsystems Technology Office, Defense Advanced Research Projects Agency, Arlington, VA.

Claudio R. C. M. da Silva (Member, IEEE) received the B.S. and M.S. degrees from the State University of Campinas, Campinas, Brazil, in 1999 and 2001, respectively, and the Ph.D. degree from the University of California, San Diego, in 2005, all in electrical engineering.

He is currently an Assistant Professor in the Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University, Blacksburg. His research interests include cognitive radio systems and ultra-wide-band technology for communications and positioning systems.

Dr. da Silva received the Best Student Paper Award at the 2003 IEEE Conference on Ultra-Wide Band Systems and Technologies. He received a graduate student fellowship from the California Institute for Telecommunications and Information Technology for 2001–2002.