Workshop Report
NSF Workshop on “Beyond Cognitive Radios”
Urbana, Illinois, June 13-14, 2011

1. Workshop Overview

The concept of "Software Radio" was introduced in the early 1990s. Since then there has been significant research and development activity on software-defined radios (SDR). An important feature of the SDR vision is its ability to integrate legacy systems with newer and more advanced radio and access technologies. The work on SDR led to the realization that opportunistic spectrum utilization can lead to a dramatic improvement in the capacity of practical wireless networks, and the concept of "Cognitive Radio" was born.

In recent years, the scarcity of radio spectrum and the resulting cost of provisioning wireless services have become an impediment to the growth of the wireless industry. As the services offered by the wireless communications industry continue to expand, so does the need for higher data rates both, instantaneous (to support real-time communication), and aggregated (to support the increasing pool of customers). Indeed, the use of the cellular spectrum has increased impressively in recent years, due to the introduction of new mobile wireless devices, and a large number of applications supported on such devices. Keeping these challenges in mind, a June 2010 Presidential Memorandum issued a call to "create and implement a plan to facilitate research, development, experimentation, and testing by researchers to explore innovative spectrum-sharing technologies".

The workshop on "Beyond Cognitive Radio" was organized with the goal of examining the state of the art in cognitive radios, and identifying research challenges that need to be addressed in future research on cognitive radio systems.

The attendees at the workshop included representatives from academia, industry, and the government. The workshop was held at the Siebel Center for Computer Science at the University of Illinois at Urbana-Champaign on June 13-14, 2011.

2. Workshop Organization

The workshop consisted of an opening session, four sessions focused on four technical areas, namely (a) Physical layer, (b) Protocols and architectures, (c) Fundamental limits, and (d) Testbeds, and an “open” session.

Welcoming remarks in the opening session at the start of the workshop were followed up by a presentation by Andrew Clegg of NSF on “Enhancing Access to the Radio Spectrum (EARS)”.

The “open” session was held at the end of the first day of the workshop, and was moderated by Zygmunt Haas of NSF. The open session provided the attendees an opportunity to offer input on a wide range of topics.

Each of the remaining four sessions was co-organized and moderated by two session chairs, and included multiple presenters, as listed below.

- P**hysical Layer**
  - *Chair*: Venu Veeravalli (University of Illinois at Urbana-Champaign)
  - *Co-chair*: Yingbin Liang (Syracuse University)
  - *Scheduled Speakers*: Natasha Devroye (University of Illinois-Chicago), Michael Pursley (Clemson University), Venkatesh Saligrama (Boston University), Qing Zhao (University of California-Davis)

- **Protocols and Architectures**
  - *Chair*: Alhussein Abouzeid (RPI)
  - *Co-chair*: Sachin Katti (Stanford)
  - *Scheduled Speakers*: Simon Haykin (McMaster University), Sachin Katti (Stanford), Joseph Mitola (Stevens Tech), Ram Ramanathan (BBN Technologies).

- **Fundamental Limits**
  - Chair: Lang Tong (Cornell)
  - Co-chair: Allen MacKenzie (Virginia Tech)
  - Scheduled Speakers: Randall Berry (Northwestern University), R. Chandramouli (Stevens Tech), Narayan Mandayam (Rutgers University), Pramod Viswanath (University of Illinois at Urbana-Champaign)

- **Testbeds**
  - *Chair*: Jennifer Bernhard (University of Illinois at Urbana-Champaign)
  - *Co-chair*: Romit Roy Choudhury (Duke University)
  - *Scheduled Speakers*: Suman Banerjee (University of Wisconsin), Ranveer Chandra (Microsoft Research), Kapil Dandekar (Drexel), Dirk Grunwald (University of Colorado), Jeffrey Reed (Virginia Tech)

Presentation summaries were provided by some of the presenters above; these summaries are included in this report. At the time of writing this report, the slides for some of the presentations are available at the workshop website at http://www.icws.illinois.edu/bcr.html.

Four breakout sessions were organized on the second day of the workshop with themes corresponding to the above four technical sessions. Discussions in the breakout sessions benefitted from the material presented in the above sessions.

3. Workshop Outcomes

This section provides a description of the outcomes from the workshop, in the form of future research challenges identified in the context of *future cognitive radio (FCR)* systems.

3.1 **Fundamental Limits**

As noted above, the workshop devoted a session of presentations and a breakout session to issues and challenges in understanding fundamental limits of future cognitive radio networks. The presentations by the panelists and follow up discussions centered on the following four topic areas with coexistence as an overarching theme: (i) architectures for future cognitive networks; (ii) incentive mechanisms to facilitate cognitive access and sharing of network resources; (iii) the role of information theory in the design of cognitive systems; and (iv) the roles of learning and cognition.

Architecture: The theme of architecture pervaded the presentations and discussions. Critical issues include the role of layering and hierarchy, centralized vs. decentralized protocol architectures, the coexistence of a variety of wireless networks, and the cooperation and competition among cognitive users. To this end, fundamental theoretical and analytical tools are needed to provide insights into the appropriate architectures for cognitive radio networks. There is also a critical need to provide seamless interoperability between systems with disparate capabilities. The architectural design of cognitive radio networks should include where feasible the ability to interface with the cloud and internet architecture. Security and privacy concerns should also be part of design consideration; characterizing vulnerabilities of cognitive radio networks is essential.

Economic Incentives and Game Theory: The potential roles of economic and game theory were discussed at the workshop. In particular, there is a need to re-conceptualize what is being bought and sold in a cognitive radio network and what mechanisms may be applied to these transactions. Game theoretic and economic analyses should further provide insights into the appropriate design of incentive and enforcement mechanisms, and should enable the anticipation and mitigation of system vulnerabilities.

Information Theory: There is a compelling need to examine many of the architectural and system design issues from information theoretic perspectives. Information theory can be a valuable tool for constructing proper architectural layering of cognitive network. Feedback is expected to play a critical role in the design of cognitive networks. To that end, information theory provides an essential framework for the design of appropriate feedback systems with insights into the required amount, type, and resources for feedback operations.

Learning and cognition: The theory of learning and cognition is one of the major pillars of cognitive network design. While learning theory has been developed extensively in different communities, significant research is required to tackle challenging problems in the design of cognitive networks. Specifically, there is a lack of understanding of the required features of cognition in a real communication networks. Theory and tools are limited for incorporating cognition and learning into network protocol design. There is a need for the optimal use of network resources for exploration and exploitation of spectrum opportunities across layers of network design.

3.2 Beyond Cognitive Radio: Lower Layer Protocols

The discussion at the physical layer session involved four presentations on physical layer aspects of FCR systems and an open discussion session about the related research challenges. We summarize this session as follows.

In FCR systems it should be possible for the radios to more efficiently and reliably learn and respond to diverse environments. FCR systems are therefore applicable to a wider range of networks, including networks with and without infrastructures. More generally, FCR may also be employed for networks with open communication resources for all users, and with certain rules in place for users sharing these resources. In this case, all users may need to sense and avoid interference, without an a priori distinction between primary and secondary users. Towards understanding, designing, and implementing such FCR systems, there are research challenges at lower layers (including the physical layer and the medium access control (MAC) layers) that need to be investigated. These research topics include: (1) how to model FCR networks; (2) how to define accurate performance measures for FCR networks; and (3) how to design efficient cognition schemes for FCR. We summarize the major issues and open problems on these topics as follows.

(1) Modeling of FCR Networks: A basic research problem is to find valid models for FCR networks. Two lower layer models were proposed at the workshop: one captures FCR systems at the physical layer via information-theoretic models, and the other captures channel state detection and access at the MAC layer via a partially observed Markov decision process (POMDP).

Existing information-theoretic models mainly characterize interference avoiding behavior (spectrum interleaving), interference controlling behavior (spectrum underlay), and interference mitigating behavior (spectrum overlay) in cognitive networks. Based on these models, fundamental communication limits on transmission rates can be characterized, which provide information-theoretic benchmarks for these networks. However, these models and approaches are well studied only for small networks. It is still an open problem to generalize these models and approaches for large FCR networks. Other open questions include: (i) how to apply information-theoretic insights to improve interference control and mitigation; (ii) how to include the cognitive information of nodes in the models; and (iii) how to characterize impact of such information on communication rates.

MAC layer models characterize dynamic spectrum availability via POMDPs. These models may also include multiple interacting cognitive users, and may have unknown model parameters. Based on these models, the actions of the users such as learning model parameters and accessing available channels can be determined. Although these models may be mathematically convenient, they may not fully capture practical reality of networks and need to be further generalized to capture realistic features of FCR networks. For example, in some recent work non-Markov models have been proposed to capture long-range dependency of channel availability. Research issues that need to be further addressed in the future include: (i) how to take into account geographical locations of radios in the model to improve channel selection and access efficiency; (ii) how to model node mobility; and (iii) how to model spatially heterogeneous spectrum availability across a network.

(2) Performance Metrics for FCR Networks: Another important topic that needs further investigation is to determine appropriate metrics for measuring the quality of information cognition and for characterizing the performance of FCR networks. This topic is closely related to modeling of FCRs. For information-theoretic models, the capacity region that represents the best possible rate distribution among radios over a network is a fundamental measure of the network performance. Open problems for future research include how to derive the network throughput based on the capacity region under certain QoS constraints, and how to generalize the capacity region for small networks to large networks. For MAC models, if the model parameters are unknown and need to be learned, it is not clear and hence is of interest to study what an appropriate measure is for the quality of learning. From the viewpoint of other higher networking layers, one possible performance measure is based on resource consumption. The idea is that a transmission that prevents N radios from receiving in a band of width B Hz for an average of *T* seconds consumes *T BN* units of network resources. In this way, “bits/sec/Hz/sqm” may be introduced as a measure of network resource consumption. Moreover, the success rate of packet delivery and the rate of link connection failure can also be used as the performance measures for transmissions of a FCR. An interesting open question to be addressed in the future is how to translate these higher layer metrics to lower layer metrics, or how to combine these higher layer metrics and lower layer metrics so that a cross-layer design may be implemented.

(3) Efficient Cognition for FCR: The salient feature of the cognitive radio technology, which differentiates it from WiFi and cellular networks, is cognition. Cognition of FCR is expected to be more powerful. In addition to sensing the availability of channels to avoid causing interference to primary transmissions, a FCR should also gather information about its environment to enhance its access of network resources and improve the overall performance of the network. In many scenarios, the environment may change dynamically over time and hence must be learned online by the radios. For example, signal propagation conditions of wireless transmissions are in general time-varying. Learning such channel state information helps substantially in adapting transmission schemes to achieve better transmission rates. Moreover, interference caused by other users in coexisting networks may also change over time. Being aware of such information helps a radio to judiciously select channels to access. Such information about interference also helps in interference management over the networks. Traffic models of coexisting networks could also be learned online for efficient allocation of sensing resources. It is also beneficial to learn the power levels and modulation schemes that other users are using, which help to adapt the cognitive radio’s power and modulation schemes for better access of the channel.

In order for a FCR to achieve efficient and accurate information cognition and fully exploit such information, the following issues must be addressed in future research:
(i) How much information should a FCR provide toward another radio’s needs in order to communicate with it efficiently and reliably?
(ii) How should we tradeoff robustness and overhead in cooperative sensing?
(iii) How should a FCR respond to increases in propagation loss and interference without increasing interference to unintended receivers?
(iv) How should a FCR adapt its use of communication resources based on learned information about the environment?

3.3 Protocols and Architectures

This workshop examined challenges and research directions beyond the current cognitive radio technology. So far, the cognitive radio concept has been explored in primarily one dimension: the possibility of sensing the presence of users in a shared wireless channel has given rise to the idea of cognition that permits a hierarchical differentiation between primary (licensed) and secondary (unlicensed/opportunistic) users.

Even within that limited view of cognition, and despite two decades since the idea of software defined radio and cognitive radio had emerged, there are several challenges that remain unaddressed and questions that remain unanswered. These include: methods for architectural design that integrate the technology pieces (reconfigurable antennas, measurement databases, distributed networking protocols, security

/authentication) into a flexible and efficient whole; effective spectrum management approaches that are evolutionary in allowing a mix of fixed and dynamic spectrum access methods, business and usage models, and realistic regulatory rules; architectural designs that integrate economic, societal and environmental factors; and, quantifiable advances in spectrum efficiency demonstrable at large-scale.

It is important to note that the major challenge that spurred significant interest in new approaches for spectrum sharing, including cognitive radio, is not going away. Spectrum scarcity and the efficient use of wireless spectrum will remain a fundamental challenge. The wireless spectrum is a limited national resource, while conservative predictions of wireless data traffic estimate a 20x increase in data traffic in the next 5 years. Therefore, it is important to investigate methods to achieve efficient utilization of the valuable spectrum resource.

However, moving beyond the limited view of cognition opens up a number of largely unexplored and challenging directions. We believe the resource challenge is not only demonstrated along the dimension of spectrum as a limited resource, but also in energy consumption and sustainability. Information and communication systems in general and wireless networks in particular have often been designed until now without regard for energy efficiency and sustainability. It is imperative that energy efficiency be a prime driver in the expanding development of communication networks as energy supplies are limited and as the reliance of the World on Information Systems is increasing. We believe that joint consideration of cognitive networks and energy consumption has not been explored adequately yet, and will become an important topic in the upcoming decade.

Energy dissipation in communication networks can occur during transmission, processing, storage, and device operation. The materials used in the electronic components, the microprocessor layout, the embedded software, the hardware and antennas, the communication signals, the network protocols, the logic core of the control systems, the applications themselves, and the interactions among all these entities affect the amount of dissipated energy. Therefore, an energy optimal design calls for joint optimization across all layers of the protocol stack. However, although existing methods can be used to minimize the energy consumption at each layer, there is no unified scientific base for assessing the energy limits for the overall communication process. A mathematical basis for a systematic energy efficient network design is necessary to enable a rigorous analysis of the impact of different cognitive network architectures (e.g., cooperative vs. individual spectrum sensing schemes, centralized vs. distributed resource allocation methods, etc.) on the overall energy and spectral efficiencies.

Another major challenge is the problem of deriving network protocols comprising flow control, routing, scheduling and physical resources management, which can provide Quality of Service (QoS) guarantees and ensure network stability with minimum energy consumption and under a wide range of service demands. Equally important is how this integrates with future Internet designs and broadband Internet access. In particular, network protocols that optimize the performance not only across layers (i.e., cross-layer), but also over time, are needed. The optimization of cognitive network operations should be carried out not only across layers but also over time because the network dynamics are usually temporally correlated and also interacting with network operations. The ability to optimize across both time and layers is particularly essential to the design and evaluation of cognitive networks at various protocol layers since cognition requires optimization (and learning) over time. Approaches are needed that can integrate the different primary performance concerns and time-scales at different protocol layers (e.g., physical and link layers versus network layer).

3.4 Testbeds

When developing future of cognitive radio and related systems, numerous new technologies, protocols, and policies will need to be developed. These will need to be evaluated using testbeds for validation, experimentation, and research. Thus, testbed development should be an important thrust of future research activity.

When developing the future testbeds, the following questions should be addressed:

- What traits should research testbeds have to make them most useful to a large number of researchers?
- What traits should research testbeds have to make them most useful over time?
- How constrained should research testbeds be to the current state (or near-future state) of system design or practice?
- How transformable should testbeds be to new architectures, protocols, layer interconnections/dependencies, and new hardware and software capabilities?
- How can testbeds be used to reproduce and/or anticipate real-world operational scenarios, including coexistence with scientific or public safety uses of spectrum?

The following desirable properties for cognitive radio testbeds were identified.

Verifiable operation: The experimental results from the testbeds should be verifiably correct. This is critical to enable validation before fielding new technologies, and for performing scientifically reproducible tests.

Upgradability: The testbeds should be designed such that they can easily accommodate newly developed technologies. Considering the high cost of testbed development, such upgradability is important to maximize the benefit of the investment, and to lengthen the lifecycle duration of the testbeds.

Broad applicability: The testbeds should be able to evaluate a wide range of technologies such as new antenna technologies (including both frequency and spatial control), tunable radio components, signal processing techniques, spectrum sharing protocols, and technological implementation of policy/access rules. Traditionally, individuals in these different disciplines tend to work independently, and it will be important to bring diverse research experts together in developing a suitable cognitive radio testbed.

Community resource for large-scale evaluations: Given the high cost of testbed deployment, the testbeds should be shared resource, available to the broader research community. The shared resource model allows development of more complex testbeds that can support experimental evaluation on a larger scale. Such a testbed will allow the research community to evaluate bold new research ideas, and also provide an eco-system for sharing resources (including software) and a platform for cross-validation.

Measurements of channel usage: In addition to testbeds that evaluate FCR technologies, there is also a significant need to gain a much better understand of the characteristics of the current spectrum usage patterns – spatially and temporally – to enable design of cognitive radios that will be better able to efficiently select and operate in the most appropriate spectral bands. To achieve this goal, a broad effort is needed to measure the detailed spectrum usage in a wide variety of environments, including rural, urban, suburban, in various geographies around the country. These measurements should yield temporal and spatial usage data sets that can be used to assess the benefits of future cognitive radio system approaches. There is also a need to develop a standardized interface to access the large amount of usage data. Real-time access to the data will be helpful for evaluating research ideas, as well as for classroom educational purposes.

Support for education: While the testbeds will support fundamental research, there is also a need for using the testbeds for educational purposes. For instance, laboratory assignments may be developed to make use of the testbeds. Cognitive radio networking requires contributions from many different domains, and such laboratory exercises will help in educating future researchers in these different domains. Additionally, the use of the testbeds on a larger scale will help identify its shortcomings, which can be mitigated in future testbed designs.

In addition to the above properties for testbeds, the workshop attendees also identified two other needs:

- The cost for deploying and maintaining a testbed is often (far) greater than the cost of the equipment itself. For instance, equipment repair, license renewals, upgrades, and access management, all incur non-trivial costs that may be difficult to predict at the time of writing a proposal for funding for the testbed. Moreover, the maintenance of the system often warrants dedicated individuals who would oversee regular software upgrades, policy enforcement, coordination across groups, inventory management, security, and other control operations. These needs may change substantially over the life of the testbed. Therefore, the funding agencies should make it possible easier to incrementally request additional funding for additional costs that may arise over the lifecycle of a testbed. This ability is also expected to increase the lifecycle duration for the testbeds.

- Due to the diverse needs of the research community, it is to be expected that, in addition to a few large testbeds, many small testbeds will also need to be built. Testbed design and development requires substantial expertise in many diverse fields. Such expertise is not likely to be available in many universities. To ease the task of testbed deployment, it may be beneficial to provide support for a team of “testbed experts” who can share their expertise with the various teams that may be working to build new testbeds.

Summaries of Presentations

Some of the presenters provided summaries of their presentations, which are included below.

Natasha Devroye

The role of information theory in designing innovative spectrum sharing technologies of the future lies in its ability to act as a benchmark, as a form of prospector, and provide insight and guidance as to where the potential gains lie in innovative spectrum sharing techniques. While several forms of spectrum sharing techniques based on primary and secondary users have been modeled information theoretically, one of the main challenges moving forward lies in going beyond this primary/secondary spectrum sharing paradigm. Indeed, how to accurately model future cognitive networks in ways that both reflect reality but are not too hampered by current technological limits and policy constraints (e.g. primary and secondary classes only) and that go beyond those, is one of the key challenges. Should layered or hierarchical networks be used, where layers may be distinguished perhaps based on their quality of service class or the ability of the nodes (or their side information such as codebook knowledge), or should networks be modeled as flat in terms of the nodes' abilities and constraints? How should we capture cognition in both ad hoc networks, infrastructure-based networks, and networks which are a combination of the two? In addition, how to take the results obtained by information theory for relatively small networks and extend those to larger networks is of key importance in analyzing where the gains lie and how they scale. In this line, how to take the dynamic nature of a network into account, rather than analyzing networks snapshot by snapshot as information theory currently does, is yet another long term challenge. This points at the need for a dynamic or ``adaptive information theory.'' This will likely be coupled with the integration of information theory into other layers of the classical network stack. Another challenge in the usage of information theory lies in predicting gains of cognitive networks is the modeling of non-traditional constraints information theoretically, which will likely result in the need for new non-classical information theoretic tools.

Challenges in enabling and exploiting cognitive radios
Ram Ramanathan

Beyond the problems studied by cognitive radio researchers over the past decade, there are a number of challenges faced in actually engineering and deploying a cognitive radio based network. These include scalability of sensing with increasing number of channels, sensing in a CSMA/CA regime where synchronized “quiet time” to allow sensing is not readily available, accommodating cross-channel interference in selection of frequencies to use, and accommodating the wide range in communication range over frequencies.

We need to go beyond the current dichotomy between channel access, which deals with interference, collisions etc. and spectrum-adaptive radios which deal with primaries. A truly cognitive radio should take a unified approach to these, treating primaries, interference, noise etc. in a unified, adaptive manner, simply moving to the right frequency, code, and antenna that is the “clearest” communication channel.

We need to incorporate advances in Artificial Intelligence to utilize cognitive radios more effectively. We have reached the limit of where human engineers can predict and design for all contingencies, such as which radios will be used, where they will be etc. We have a dynamic, distributed, heterogeneous, partially-observable, multi-objective optimization problem, for which AI techniques such as distributed planning and optimization, and machine learning are perfectly suited. With such adaptation, a cognitive network trained in a desert can learn to perform well under water.

Beyond the problems and solutions for cognitive radios lies the question: how can we best exploit cognitive radios for better network performance? Some exciting opportunities lie in the area of Cognitive Radio Ad Hoc Networks (CRAHNs). Cognitive radios allow such networks to access a large number of orthogonal frequencies, which can be used to multiplex end-to-end data over multiple orthogonal paths for end-to-end burst throughput well in excess of an individual radio's data rate. The ability to transmit on one frequency while receiving on another enable cut through routing, dramatically reducing latencies. Finally, the assignment of frequencies results in topology which is now a controllable parameter rather than a given, and allows the network layer to tailor topology to the needs. These capabilities are enabled by cognitive radio, but need a re-think of MANET architectures to fully exploit them.

Future Adaptive Protocols for Cognitive Radios
Michael B. Pursley

New protocol suites will be required for future cognitive radio networks, including those that will have dynamic access to radio spectrum. A framework based on resource consumption is provided for the selection of bandwidth, session duration, transmission power level, modulation technique, and error-control coding method in cognitive radio networks with time-varying communication environments. We discuss the need for protocols that adapt modulation and coding rather than increase radiated power in response to deteriorating channel conditions (e.g., increasing fading or interference). Information-theoretic bounds and practical performance results are presented to give quantitate assessments of the compensation that can be achieved from adaptation of the modulation and error-control coding. We examine methods for the cognitive radio's real-time determination of channel conditions, and we show how the use of simple receiver statistics can avoid the need for channel measurements and channel parameter estimates. Such receiver statistics provide control information for adaptive protocols and they also permit spectrum sensing while communicating. We describe the need for protocol design and development methods that employ direct generation of receiver statistics. These methods avoid time-consuming embedded simulations of the physical layer in performance evaluations of adaptive transmission, media access control, routing, and other cross-layer protocols for future cognitive radio networks.

Testbeds and Antennas "Beyond Cognitive Radio"
Kapil R. Dandekar

As cognitive radio systems evolve, there is a need for hardware testbeds and antenna systems to keep pace. Cognitive radio is traditionally defined to be biologically inspired algorithms for radios to sense and adapt to surroundings. Transceivers, including antennas, must be developed that can provide additional degrees of freedom that can be exploited by cognitive sensing and adaptation techniques. Leveraging recent advances in metamaterial antennas, reconfigurable antennas, body worn antennas, and flexible electronics, we have developed transceiver systems that can unobtrusively implement pattern, polarization, and frequency degrees of freedom for cognitive radios. We have also developed transceivers to apply cognitive inspired algorithms to communication modalities beyond radio, including free space optical and ultrasonic communication. Furthermore, we are developing a ultrawideband software defined radio platform for the community with GHz sampling so that next generation cognitive radios can be prototyped and field-tested.

Putting cognition into CR networks
R. Chandramouli

We first note that dynamic spectrum access (DSA) networking is just one application enabled by cognitive radios. Therefore, moving beyond cognitive radio enabled DSA, information sensing and adaptation can be applied at different layers of the network stack. Spectrum and interference awareness at the PHY/MAC layer, continually measuring the wireless route states and adapting at the IP layer, spectrum agile radio aware transport protocols, radio policy awareness, knowing the competition at the wireless service provider level are all examples of moving beyond DSA.

Learning theory plays an important role in the cognitive process. Some key challenges in (stochastic) learning based cognitive radio networking include: (a) **multi-time scales** (e.g., spectrum sensing layer vs. application layer adaptive error control coding) across the stack; (b) **uncertainties** -- complete knowledge, partial information, etc.; (c) **type of learning** -- ergodic vs. non-ergodic learning; (d) number of bits of (implicit or explicit) feedback required to adequately represent the cognitive radio states at the various layers; and (e) self-learning for interacting network of cognitive radios.

Therefore, in summary, research in learning enabled cognitive radios cuts across social science (e.g., human behavior), wireless networking, economics/game theory and testbed experimentation.

Dynamic Channel Selection for Spectrum Exploration and Exploitation
Qing Zhao

Dynamic channel selection is one of the main design issues in cognitive radio networks with opportunistic spectrum access. With the dynamic and stochastic spectrum availability, the design of channel selection strategies often lead to a complex problem of optimization over time. The presence of multiple interacting cognitive users and the absence of an accurate spectrum occupancy model further complicates the problem. In this talk, we examine the technical challenges in dynamic channel selection and explore approaches based on sequential decision theory and stochastic online learning.

Venkatesh Saligrama

The advent of frequency-agile radios holds the potential for improving the utilization of spectrum by allowing wireless systems to dynamically adapt their spectral footprint based on the local conditions. Whether this is done using market mechanisms or opportunistic approaches, the gains result from shifting some responsibility for avoiding harmful interference from the static regulatory layer to layers that can adapt at runtime. However, this leaves open the major problem of how to enforce/incentivize compliance. If there is only one system then there is a fear of being caught. However when there are many coexisting systems there is an incentive to cheat and hide in the crowd. Hence there is a need for the traceability of culprits. Rather than heavy-handedly specifying a specific identity beacon and requiring every primary user to be able to decode it, we propose a light-handed approach in which cheating users only need to be identifiable by their pattern of interference. The problem breaks into two parts: discrimination, where we distinguish between natural fading and culpable interference and identification, where we decide which potential users are actually liable.

Light-handed regulation is interpreted as making unambiguous (and easily certified) requirements on the behavior of individual devices themselves while still preserving significant freedom to innovate at both the device and the system level. The basic idea explored here is to require the PHY/MAC layers of a cognitive radio to guarantee silence during certain timeslots where the exact sequence of required silences is given by a device/system-specific code. Thus, if a system is a source of harmful interference, the interference pattern itself contains the signature of the culprit. Nevertheless, identifying the unique interference pattern becomes challenging as both the number of cognitive radios and the number of harmful interferers increases. The key tradeoffs are explored in terms of the regulatory overhead (amount of enforced silence) needed to make guarantees. The quality of regulatory guarantees is expressed by the time required to convict the guilty, the number of potential cognitive systems that can be supported, and the number of simultaneously guilty parties that can be resolved. As it turns out, the time to conviction need only scale logarithmically in the potential number of cognitive users! The base of the logarithm is determined by the amount of overhead that we will tolerate and how many guilty parties we want to be able to resolve. For example, with time-slots of four milliseconds, we can support more than two-hundred potential cognitive users each having access to more than 80% of the time-slots and still be able to resolve harmful interference to a guilty pair of users within two seconds! This means that even if one user is malfunctioning, there is still no incentive for another user to start cheating since it will fear being identified and punished.

List of Participants

Last name is listed first below, followed by first name.

Aazhang, Behnam
Abouzeid, Alhussein
Atia, George
Banerjee, Suman
Bernhard, Jennifer
Berry, Randall

Bhattacharyya, Shuvra
Bonneau, Robert
Chandra, Ranveer
Chandramouli, Rajarathnam
Chatterjee, Mainak
Cheng, Susan
Roy Choudhury, Romit
Dandekar, Kapil
Das, Samir
Devroye, Natasha
Elliott, Chip
Eriksson, Jakob
Ghosh, Amitava
Grunwald, Dirk
Haas, Zygmunt
Haykin, Simon
Honig, Michael
Katti, Sachin
Kravets, Robin
Liang, Yingbin
Liu, Xin
Love, David
MacKenzie, Allen
Mandayam, Narayan
Melodia, Tommaso
Minden, Gary
Mitola, Joseph
Nelakuditi, Srihari
O, Kenneth
Pursley, Michael
Ramanathan, Ram
Rangarajan, Sampath
Reed, Jeffery
Roberson, Dennis
Saligrama, Venkatesh
Shen, Chien-Chung
Srikant, Rayadurgam
Song, Min
Swamy, Ananthram
Tong, Lang
Vaidya, Nitin
Veeravalli, Venugopal
Vishwanath, Sriram
Viswanath, Pramod
Xin, Chunsheng
Zhao, Qing

Acknowledgement: This workshop was funded in part by the U.S. National Science Foundation award number ECCS-1059648.

Any opinions, findings, and conclusions or recommendations expressed in this report are those of the authors and do not necessarily reflect the views of the National Science Foundation.