

# CogNet - An Architectural Foundation for Experimental Cognitive Radio Networks within the Future Internet

Dipankar Raychaudhuri  
Narayan B. Mandayam  
WINLAB, Rutgers University  
New Brunswick, NJ 08902 USA  
{ray,narayan}@winlab.rutgers.edu

Joseph B. Evans  
Benjamin J. Ewy  
ITTC, University of Kansas  
Lawrence, KS 66045 USA  
{evans,bewy}@ku.edu

Srinivasan Seshan  
Peter Steenkiste  
ECE & CS, Carnegie Mellon Univ.  
Pittsburgh, PA 15213 USA  
{srini,prs}@cs.cmu.edu

## ABSTRACT

This paper describes a framework for research on architectural tradeoffs and protocol designs for cognitive radio networks at both the local network and the global internetwork levels. Several key architectural issues for cognitive radio networks are discussed, including control and management protocols, support for collaborative PHY, dynamic spectrum coordination, flexible MAC layer protocols, ad hoc group formation and cross-layer adaptation. The overall goal of this work is the design and validation of the control/management and data interfaces between cognitive radio nodes in a local network, and also between cognitive radio networks and the global Internet. Protocol design and implementation based on this framework will result in the *CogNet* architecture, a prototype open-source cognitive radio protocol stack. Experimental evaluations on emerging cognitive radio platforms are planned for future work, first in a wireless local-area radio network scenario using wireless testbeds such as ORBIT, and later as part of several end-to-end experiments using a wide-area network testbed such as PlanetLab (and GENI in the future).

## Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design - *Distributed networks, Network topology, Wireless communication.*

## General Terms

Protocols, Algorithms, Management, Design, Experimentation.

## Keywords

Cognitive Radio, Network Protocols, Internet Architecture

## 1. INTRODUCTION

Recent “Moore’s law” advances in programmable integrated circuits have created an opportunity to develop a new class of intelligent or “cognitive” radios [1][2] which can adapt to a wide variety of radio interference conditions and multiple protocol standards for collaboration between otherwise incompatible systems. Such a cognitive radio would be capable of very dynamic physical layer adaptation via scanning of available

spectrum, selection from a wide range of operating frequencies (possibly non-contiguous), rapid adjustment of modulation waveforms and adaptive power control. In addition, a suitably designed cognitive radio with a software-defined physical layer would be capable of collaborating with neighboring radios to ameliorate interference using higher-layer protocols. These higher layer coordination protocols could range from multi-node signal combining and coding methods to etiquette mechanisms all the way to fully collaborative multi-hop forwarding between radio nodes. Thus, suitably designed cognitive radios have the potential for creating a next-generation adaptive wireless network [3] in which a single universal radio device is capable of operating in a variety of spectrum allocation and interference conditions by selecting appropriate physical and network layer parameters often in collaboration with other radios operating in the same region. Such a “cognitive network” will lead to increased network capacity and user performance. Perhaps for the first time in the short history of networking, cognitive radios offer the potential for organic formation of infrastructure-less collaborative network clusters with dynamic adaptation at every layer of the protocol stack including physical, link and network layers [4].

While the development of cognitive radio hardware and software, especially at the physical layer, has received considerable attention, the question of how one organizes a set of cognitive radios into a cognitive network is not well understood. As such, adaptive networks of cognitive radios represent an important but demanding research challenge for both the wireless and networking communities. The extreme flexibility of cognitive radios has significant implications for the design of network algorithms and protocols at both local/access network and global internetworking levels. In particular, support for cross-layer algorithms which adapt to changes in physical link quality, radio interference, radio node density, network topology or traffic demand may be expected to require an advanced control and management framework with support for cross-layer information and inter-node collaboration. At the wireless local-area network level, an important technical challenge is that of distributing and managing this inter-node and cross-layer information then using this control information to design stable adaptive networking algorithms that are not overly complex. At the global internetworking level, clusters of cognitive radios represent a new category of access network that needs to be interfaced efficiently with the wired network infrastructure both in terms of control and data. End-to-end architecture issues of importance include naming and addressing consistent with the needs of self-organizing network clusters, as well as the definition of sufficiently aggregated control and management interfaces between cognitive radio networks and the global Internet [5].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

*MobiArch'06*, December 1, 2006, San Francisco, CA, USA.

Copyright 2006 ACM 1-59593-566-5/06/0012...\$5.00.

The *CogNet* architectural foundation was designed to support the following capabilities:

- Spectrum agility and fast spectrum scanning over multiple frequency bands, providing local awareness of radio interference and the ability to change frequency bands on a per-packet basis
- Fast PHY adaptation, or the ability to change physical-layer waveforms on a per-packet basis and PHY collaboration modes such as network coding
- Spectrum etiquette protocol and dynamic spectrum policy implementation on a per-session basis
- Fully programmable MAC layer, with the option of dynamic adaptation to meet service needs
- Cross-layer protocol implementation capabilities based on integrated PHY, MAC, network algorithms
- Ad hoc cluster formation, supporting multi-hop packet forwarding among peer groups of radio nodes.

## 2. ARCHITECTURAL FOUNDATION

In order to implement the capabilities detailed above, we envision several inter-module interfaces and protocols, specifically:

- A *Global Control Plane (GCP)* implemented as a cross layer network management overlay that can interface with the network layer and can provide aggregated representations of the cognitive subnetwork state to the future Internet.
- An *API for PHY layer adaptation* (e.g., agility, change of modulation waveform), and support for collaborative PHY in the form of network coding.
- *Spectrum coordination protocols* that facilitate dynamic sharing among radio nodes using mechanisms such as etiquette policies or spectrum server.
- *Autoconfiguration* (e.g., bootstrapping and topology discovery) protocols that can be used to establish network connectivity after a cognitive radio device is turned on or enters a new service area.
- *Flexible MAC* framework that permits programmable functionality capable of dynamic selection of channel sharing modes based on observed network conditions and traffic demands.
- *Network layer protocols* that support service discovery, naming, addressing and routing in ad hoc wireless constellations, including features that provide economic incentives for collaboration.

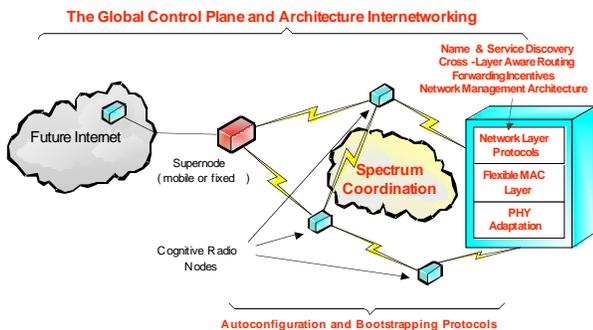


Figure 1. Research concepts explored by CogNet.

Figure 1 presents the architectural elements that we will experiment with in *CogNet*. These key elements of an architecture, and a discussion of how they provide the required capabilities and research tools, are detailed below.

## 2.1 Global Control Plane

In view of the complexity and range of control and management functions required, it is clear that the protocol functionality of a cognitive network should be partitioned into separate control and a data planes. The data plane protocol stack on each node contains the modules needed to support data communication between the wireless nodes and it exposes a set of controls for each module through an API. This API is used by a general and extensible “Global Control Plane” (GCP) to monitor, configure, and adapt the data plane modules.

In the GCP-based architecture, illustrated in Figure 2, each node has a dedicated control “interface” along with data interface. The control interface provides an initial radio bootstrapping and service discovery function that can operate either on a channel at the edge of the service band or a dedicated portion of a TDMA frame, and would have wider radio coverage than a typical service channel. The bootstrapping functions can utilize other nodes to rebroadcast control packets using a controlled flooding mechanism, thus, providing global awareness to all cognitive radios within a subnetwork.

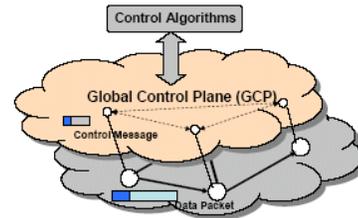


Figure 2. GCP architecture.

The GCP also provides higher level functionality including a network management and control overlay, and an interworking point where aggregated representations of the cognitive radio network state and control points are provided to the future Internet. The GCP is envisioned to be tightly integrated with the control and management of the future Internet backbone network and interact with its network management infrastructure about communication and security policies.

The complete GCP protocol stack shown in Figure 3 provides for the distribution of control messages required to optimize the various collaborative PHY, spectrum coordination, flexible MAC or ad hoc networking functions for best performance in the data plane.

## 2.2 PHY Adaptation and Network Coding

Communication networks today share the same fundamental principle of operation. Independent data streams may share network resources, but the information itself is separate. Routing, data storage, error control, and generally all network functions are based on this assumption. Network coding [6] breaks with this assumption. Instead of simply forwarding data, nodes may collaborate with each other to recombine several input packets into one or several output packets.

There are two main benefits of this approach: potential throughput improvements and increased robustness particularly for multicast and broadcast service scenarios. Robustness translates into loss resilience and facilitates the design of simple distributed algorithms that perform well, even if decisions are based only on partial information. In fact, successful reception of

information does not depend on receiving specific packet content but rather on receiving a sufficient number of independent (recombined) packets.

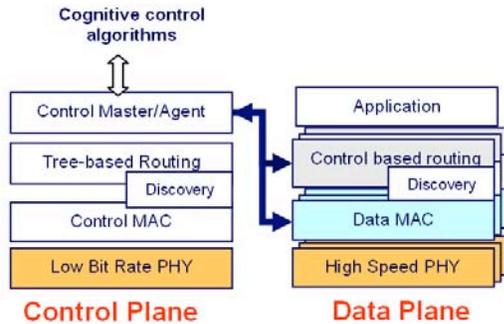


Figure 3. Protocol stack in the GCP architecture.

A number of theoretical results on network coding have been published over the past few years, of particular interest here being its robustness when introducing forwarding incentive mechanisms [7], but real-time experimental validation has lagged behind because of the need for a programmable radio with cross-layer control capabilities. In particular, the implementation of the coded multicast arrangement will require GCP support for path establishment, specification of forwarder functionality (routing and coding) and correlation of packets to be combined via a global packet ID (such as source address, destination address, port numbers and sequence number). The *CogNet* framework will allow research into collaborative PHY architectural tradeoffs by providing this implementation of network coding as a representative example.

### 2.3 Spectrum Coordination Protocols

A number of approaches have been proposed for improved spectrum sharing over the past decade. Notable methods being discussed in the technical and regulatory communities include property rights regimes [8], spectrum clearinghouse [9], unlicensed bands with simple spectrum etiquette [10], open access [11] and cognitive radio under consideration here. The distinctions between unlicensed spectrum regimes, open access and cognitive radio approaches are relatively subtle as they are all based on the concept of technology neutral bands to be used by a variety of services using radio transceivers that meet certain criteria. For example, cognitive radio may be viewed as a special case of open access or unlicensed regimes in which radio transceivers are required to meet a relatively high standard of interference avoidance via physical and/or network layer adaptation. The cognitive radio principles currently under consideration by the FCC and the research community span a fairly wide range of possible functionalities both at physical and network layers.

The “agile wideband radio” scheme [12] is the most prevalent concept for cognitive radio, in which transmitters scan the channel and autonomously choose their frequency band and modulation waveform to meet interference minimization criteria without any protocol-level coordination with neighboring radio nodes. Another simple technique is “reactive control” [13] of transmit rate/power, in which radio nodes do not have any explicit coordination with neighbors but seek equilibrium resource allocation using reactive algorithms to control rate and power, analogous to the way the TCP protocol reactively adjusts source bit-rate over the Internet.

The *CogNet* architectural foundation was designed to support research of a slightly higher level of protocol complexity by supporting research on spectrum etiquette protocols [14] to improve coordination between radio nodes, using either Internet-based spectrum services or a common spectrum coordination channel at the edge of a shared frequency band.

The *common spectrum coordination channel* (CSCC) approach has been proposed [14] as a candidate mechanism for implementing spectrum etiquette policies. In this approach, each wireless device sends periodic beacons containing spectrum usage information so that neighboring nodes can avoid using the same frequencies, or if the network is too congested, can execute specified spectrum etiquette policies. It is observed here that the CSCC approach is similar in concept to the global control plane discussed earlier, and can thus be integrated into a GCP implementation as a subset of its overall functionality. In our earlier work [14] we have shown that CSCC implementations can achieve significant spectrum efficiency improvements with relatively simple etiquette policies when compared with radios with spectrum scanning and reactive agility.

An alternative approach which can be used for spectrum coordination is the *spectrum server*. In this design, all radios use a spectrum management protocol to communicate with a centralized spectrum service within the future Internet. In our recent work in [15][16], we have examined the boundaries of system performance under the assumption that efficient access to spectrum can be resolved by an impartial “spectrum server” that can obtain information about the interference environment through measurements contributed by different terminals, and then offer suggestions for efficient coordination to interested service subscribers. There are many ways in which a spectrum server can coordinate a network of cognitive radios. Recent work in [15] has considered the role of the spectrum server in scheduling variable rate links while the work in [16] has considered the spectrum server’s role in demand responsive pricing and competitive spectrum allocation.

### 2.4 Radio Autoconfiguration Protocols

Self-organization is a key requirement for cognitive radio networks. The Global Control Plane provides a bootstrapping process that enables a radio node to be aware of itself, the surrounding nodes and current network status when it starts up. There are two phases in this process: (1) obtaining PHY parameters, reachability, and performance information by listening on a control channel or channel scanning; and (2) negotiation, during which the new node can negotiate with existing sub-networks for name/service discovery or performance optimization. The appropriate association and/or authentication process can then be initiated for the new node to join existing networks.

The bootstrapping and discovery process begins with bootstrapping beacons with information about current network status are broadcast within one hop on a specific control channel by existing active nodes in a periodic and opportunistic way. When a new node is starting up (or moving close to current sub-networks), it first listens for bootstrapping beacons on the control channel to obtain PHY and connectivity information to initialize its radio parameters by choosing proper operating frequency, transmit power, bandwidth/modulation/rate, etc. Service discovery information is obtained from the bootstrapping beacons,

which allows the node to utilize higher levels of the Global Control Plane for network layer, naming, and cross layer management and control, and other services. Current nodes can provide performance information in the beacons, such as link speed, available network capacity, congestion indication, etc. In addition to common signaling channel, nodes can scan different traffic channels and listen to beacons/send an association message to the best “cost” parent using the appropriate discovery metric.

The bootstrapping beacon is a low layer (PHY) messaging mechanism supporting neighbor discovery and determination of the logical topology. It also conveys cross-layer information, for example, PHY parameters such as available frequencies, transmit power, modulation/coding type, radio bandwidth, bit-rate, and network name as well as available services.

## 2.5 Flexible MAC Layer

Wireless medium access control (MAC) protocols [17][18][19] tend to have different operating conditions in which they operate best. For example, RTS/CTS based MAC protocols work well in hidden node scenarios, but also incur high overheads when such collisions are rare. Cognitive networks provide us the opportunity to dynamically change the MAC protocol [20] to suit both the needs of the applications running on the nodes as well as the properties of the environment around them. Our previous work has shown that this form of adaptation is useful for higher layer protocols and recent preliminary work [21] has shown some promise in adapting MAC protocols as well. The CogNet foundation will support research into changing MAC behavior on-the-fly involving several challenges as outlined below.

The first step in choosing the best MAC operating mode is to understand the propagation environment and traffic demand matrix of the involved nodes. The propagation environment must be measured/inferred using relatively sparse information and both the propagation environment and traffic matrix can change quickly over time. An added complication is that even when factors such as the propagation properties of the environment and the traffic demand matrix of the nodes are known exactly, choosing the best MAC protocol may be difficult. For example, a particular choice of MAC will rarely be best for all nodes involved. This requires carefully balancing the global utility of the system with fairness. In addition, the optimal choice may change dramatically over time. However, there may be considerable overhead in switching between different MAC behaviors and the system will need to manage such dynamic environments by balancing performance gains against mode switching overhead.

The different MAC layers that are supported may be inherently incompatible in their behavior. As a result, the system must ensure that all nodes are synchronized in any switch between MAC protocols. This ensures that the impact on reachability between nodes is minimal and short-lived. It also ensures that there is no oscillation between operating modes as different groups of nodes reach different decisions due to inconsistent information. Choosing a compatible MAC protocol ensures that a pair of nodes can communicate directly with each other. However, nodes using incompatible MAC protocols may be able to co-exist in an area much as existing incompatible wireless devices co-exist today. A key requirement of the flexible MAC is that it coordinates closely with any network topology management system. Nodes that are part of the same constellation must use

compatible MAC protocols, while independent constellations may make independent optimization decisions.

The final issue we consider here is the design of the interface between the flexible MAC layer and the Global Control Plane. The GCP is responsible for reconciling different application requirements and network policies. It is also responsible for handling much of the coordination between nodes in the system. The GCP is likely to operate on coarser time-scales. The flexible MAC protocol will adapt to more rapid changes in conditions. This partitioning of responsibilities results in an interface where the control plane may handle issues such as managing constellations and determining fairness policies but the flexible MAC chooses the current operating mode. We plan to explore alternate designs in this space by trading off GCP functionality with those implemented as part of the flexible MAC layer.

## 2.6 Network Layer Protocols

The network layer for cognitive radio networks will need to support a variety of applications including data and mobile real-time services (e.g., voice). To enable research into architectures supporting these applications given the variability of wireless communication, the *CogNet* system will provide a cognitive, multi-overlay network layer that can adapt based on sensing and learning of the cross layer environment, the communication model for the applications (e.g., one-to-one, one-to-many), and policy and security constraints. These latter issues are particularly important in a cognitive radio environment where opportunistic communications with previously unknown nodes may become common.

The cognitive network layer will utilize both non-traditional approaches such as overlay-based mechanisms for communication within a subnet, as well as support the concept of supernodes [22] which will serve as a gateway between local network layers within the cognitive subnet as well as to the future Internet and its IP-based and overlay-based networks. Overlays provide a large number of optimization points, and may be tailored to the application, inspiring the idea of multiple network layers tailored to specific applications or communication flow types. These might include routing overlays that have shown promise in ad hoc network layer routing scenarios, application-tailored overlay structures such as topologically aware overlays for group messaging, and overlays that support rich queries [23].

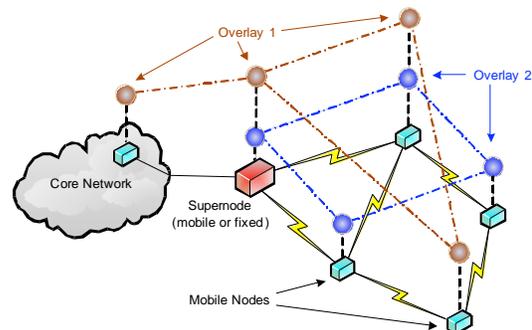


Figure 4. Multiple network overlays.

Figure 4 illustrates a cognitive radio network with mobile nodes and an interconnection to the wired core network through a supernode. Use of a particular overlay will be decided upon by cognitive techniques (e.g., case-based, expert system) based on factors including the wireless environment, application, and

policy. Overlay 2, for example, might be a multicast optimized DHT-based network layer being used to support a messaging application. Overlay 1 might be IP or some other protocol useful in interacting with nodes on the future Internet.

We will utilize our architectural framework to explore architectural issues including naming and service discovery, cross-layer aware routing, forwarding incentives, and network management.

*Naming and Service Discovery:* The cognitive radio network supports naming and addressing mechanisms that provide for self-organization, translation for global reachability, and merging / disconnection of cognitive networks themselves as well as with the wired network infrastructure. The approach is an extension of the concept of a supernode [22] to serve not only as a router between different geographically diverse cognitive radio networks, but to serve as a router between multiple network layers (e.g., DHTs, IP) within a cognitive radio subnet, and as a gateway to the future Internet. The identity of supernodes will be broadcast to new wireless nodes via the GCP's radio bootstrapping protocol. As a node joins a cognitive wireless network, layer 2 connectivity is established with neighbors and a join message for a service discovery overlay is sent to a neighbor node. This allows the node to discover how to access all of the various services such as the network management overlay, service-specific network layers, or IP configuration data. As a node joins the service discovery layer, the super node will expose the new node's identifier to the wired network via the appropriate name service (e.g., DNS). If there are multiple supernodes with connectivity to overlapping wired networks, name based overlays such as i3 [24] can be used to specify a supernode gateway.

*Cross-Layer Aware Routing:* Routing between cognitive radio nodes is influenced by control information from a number of sources. Nodes may obtain information about the application traffic, specified policies, link capabilities, MAC layer congestion status, and network reachability, and then decide on the most appropriate network layer to deliver the message to the peer. In order to support this routing decision, physical communication status and quality information from the cognitive wireless network will be exposed to participating nodes via the Global Control Plane. Note that MAC adaptation and routing decisions in the presence of cross-layer information need not be independent, and (at the expense of complexity) more general integrated routing and MAC algorithms can be investigated within the CogNet foundation.

A key architectural issue is the specification of control interfaces between routing protocols in the cognitive radio subnets and the wide-area internetwork. For scalability, these interfaces need to support suitable forms of aggregation and hierarchy in order to limit the amount of cognitive radio subnet specific information that needs to be visible to other routers in the future Internet. While routing of a packet from a corresponding host in the wired network would benefit from some visibility of network topology, MAC congestion and link quality within the cognitive radio network, this needs to be balanced against complexity and control aggregation considerations. In general, the simple abstraction of a "wireless link" does not apply to a multi-hop cognitive radio path when viewed from a host or router within the wired core. One potential solution for this is the concept of "dynamic topology control" which has been proposed [5] as a possible abstraction that permits specification of topology and link

quality in an integrated fashion to reflect the reality of the radio network. The *CogNet* architecture allows for investigation of alternative aggregated representations of cognitive radio network state and evaluate trade-offs between control granularity and end-to-end performance.

*Forwarding Incentives in Cognitive Networks:* A key design issue for ad hoc wireless networks including those formed by cognitive radios is that of creating voluntary collaborative groups with agreements to forward each others' packets. This type of collaboration between radios makes the most sense as node densities increase to a point at which PHY-layer spectrum coordination cannot sufficiently resolve congestion problems in the system. Collaboration at the network layer via formation of ad hoc networks is considered to be a powerful mechanism for better utilization of spectrum, and can be associated with low-power radio transmission to nearby neighbors as well as higher bit-rate, better quality PHY links compared with direct connection to a distant receiver. It has been shown [25] that mobility combined with forwarding can increase the capacity per node of an ad hoc network. If we assume a willingness to relay data, forwarding allows the exploitation of multi-channel diversity.

The *CogNet* system assumes a more complex context of heterogeneous users sharing spectrum, where a willingness to relay data needs to be incentivized via suitable protocol mechanisms. Using a microeconomic framework based on game theory [26], we have designed and analyzed a pricing algorithm that encourages forwarding among autonomous nodes by reimbursing forwarding costs accrued in terms of energy and lost opportunities for transmission of one's own data while forwarding for others. Our results have shown that pricing with reimbursement appears to improve the network aggregate utility (or aggregate bits per Joule) as well as utilities and revenue compared to the corresponding pricing algorithm without reimbursement. Our work has also revealed that the nodes' willingness to forward is greater when there is greater clustering of nodes in the network. Specifically, it has been observed that for large ratios of the average inter-nodal distance to the smallest distance between the access point and any source node, the tendency to forward decreases.

While the above results have been derived in the context of a simple local area network formed of an access point and a set of wireless nodes, it bears asking the question of how this affects or (or in turn is affected) by an end-to-end network architecture. In a local area network context, the incentive based resource allocation algorithms explicitly take into account wireless transmission parameters such as energy constraints, radio channel quality, and throughput performance. In the presence on an overarching end-to-end architecture, the challenge is in understanding, improving and designing the algorithms for the cognitive radio subnet after taking into account interactions with reimbursement mechanisms within the wired Internet.

*Network Management Architecture:* Network management infrastructure is often designed and deployed based on a hierarchical manager/agent architecture, with extensions for ad hoc wireless networks and techniques for passively and actively monitoring overlays. The *CogNet* framework supports evaluation of different methods for handling the variety of constraints that occur in a mobile dynamic RF spectrum wireless environment, and the unique fault localization environment associated with mapping faults that occur in a dynamic radio underlay into useful

notifications for the changing substrate of overlays providing higher level services and applications.

In order to facilitate routing and cognitive RF functionality, the Global Control Plane will include a network management overlay to provide cross layer information from the entire subnet. Each node can expose its own cross layer information into the overlay, and access other nodes' cross layer information as well, for a unique view into the hop by hop and end to end environment. The subnet wide management plane will be filled with and provide access to resource policies, node capabilities, link quality estimates, node availability, node performance metrics, fault localization, topology, and spectrum measurements that can be utilized via the GCP by protocol components from any network layer. Supernodes will be able to map topology and other network management data into evolving services in the future Internet such as topology services and information planes [27]. Applications themselves will also be able to make use of this information for fine tuning parameters such as codecs, packets sizes, timeouts, etc based on the cross layer information they see. We will explore the tradeoff of the additional overhead associated with maintaining and accessing the network management overlay to the additional performance and reliability gains from using this service.

### 3. CONCLUSIONS

The *CogNet* architectural framework provides a unique research infrastructure for the integration of cognitive networks into the global Internet. Taken together, the provided set of cognitive radio capabilities and the overall experimental architecture will be a major advance in the state-of-the-art. We continue to identify broad architecture and protocol design approaches for cognitive networks at both the local network and at the global internetwork levels. This architectural study will lead to the design of control/management and data interfaces between cognitive radio nodes in a subnetwork, and between cognitive radio subnetworks and the global Internet. The *CogNet* protocol stack software under development will allow us to apply these architectural results to a variety of cognitive networking experiments including both local area network experiments on wireless testbeds such as ORBIT, and later, for global end-to-end experiments when connected to experimental systems like PlanetLab and GENI in the future.

### 4. REFERENCES

- [1] J. Mitola III. Cognitive Radio: An Integrated Agent Architecture for Software Radio, *PhD thesis*, Royal Institute of Technology (KTH), Sweden, May 2000.
- [2] FCC 03-322 NPRM on Cognitive Radio, [http://hraunfoss.fcc.gov/edocs\\_public/attachmatch/FCC-03-322A1.pdf](http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-03-322A1.pdf)
- [3] D. Raychaudhuri, Adaptive Wireless Networks Using Cognitive Radios as a Building Block, *MobiCom 2004 Keynote Speech*, Philadelphia, PA, September, 2004.
- [4] A. Bletsas and A. Lippman, Efficient Collaborative (Viral) Communication in OFDM Based WLANs, *IEEE/ITS International Symposium on Advanced Radio Technologies (ISART 2003)*, Institute of Standards and Technology, Boulder, CO, March 2003.
- [5] D. Raychaudhuri and M. Gerla, New Architectures and Disruptive Technologies for the Future Internet: The Wireless, Mobile and Sensor Network Perspective, Report of NSF WMPG Workshop, August 2005.
- [6] R. Ahlswede, N. Cai, S.-Y. R. Li and R. W. Yeung, Network information flow, *IEEE Trans. on Info. Theory*, vol. 46, pp. 1024-1016, 2000.
- [7] S. Che and X. Wan, Network Coding in Wireless Network, *Proc. of 16th Int. Conf. on Computer Communication*, China, 2004.
- [8] G. Faulhaber and D. Farber. Spectrum management: Property rights, markets and the commons. *Proc. of the Telecommunications Policy Research Conference*, Alexandria, VA, October 2003.
- [9] Y. Benkler, Some Economics of Wireless communications. *Harvard Journal of Law and Technology*, 2002.
- [10] D. Satapathy and J. M. Peha, Etiquette modification for unlicensed spectrum: approach and impact. *Proc. of the IEEE Vehicular Technology Conference*, pages 272–276, 1998.
- [11] G. Rosston et al, Comments of 37 Concerned Economists. In Submitted to the U.S. Federal Communications Commission In the Matter of Promoting the Efficient Use of Spectrum through Elimination of Barriers to the Development of Secondary Markets, *WT Docket 00-230*. 2001.
- [12] Kiran Challapali, Stefan Mangold and Zhun Zhong, Spectrum Agile Radio: Detecting Spectrum Opportunities, *6th Annual Int. Symp. On Advanced Radio Technologies*, Boulder, Colorado, 2 - 4 March 2004.
- [13] X. Jing and D. Raychaudhuri, Spectrum Co-existence of IEEE 802.11b and 802.16a Networks using the CCCC Etiquette Protocol, *Proc. of IEEE DySPAN'05*, Baltimore, MD, Nov. 8-11, 2005.
- [14] D. Raychaudhuri and X. Jing, A spectrum etiquette protocol for efficient coordination of radio devices in unlicensed bands, *Proc. of IEEE PIMRC'03*, Beijing, September 2003.
- [15] C. Raman, R. Yates, and N. Mandayam, Scheduling variable rate links via a spectrum server, *Proc. of IEEE DySPAN'05*, Nov. 2005.
- [16] O. Ileri, D. Samardzija, T. Sizer, and N. Mandayam, Demand responsive pricing and competitive spectrum allocation via a spectrum policy server, *Proc. of IEEE DySPAN'05*, Nov. 2005.
- [17] D. J. Goodman and S. X. Wei, Efficiency of Packet Reservation Multiple Access, *IEEE Transactions on Vehicular Technology*, Vol. 40, No. 1, Feb. 1991.
- [18] V. Bharghavan, A. Demers, S. Shenker and L. Zhang, MACAW: A Media Access Protocol for Wireless LAN, *Proc. of ACM SIGCOMM'94*, 1994.
- [19] Z. Wu, D. Raychaudhuri, D-LSMA: Distributed Link Scheduling Multiple Access Protocol for QoS in Ad-hoc Networks, *Proc. of IEEE GLOBECOM '04*, November 2004.
- [20] C. Doerr, M. Neufeld, J. Fifield, T. Weingart, D. Sicker, and D. Grunwald, MultiMAC - An Adaptive MAC Framework for Dynamic Radio Networking, *Proc. of IEEE DySPAN'05*, Nov. 2005.
- [21] A. Akella, H. Bhambe, S. Seshan, The Impact of False Sharing on Shared Congestion Management, *Proc. of 11th IEEE International Conference on Network Protocols*, Atlanta, GA, November, 2003.
- [22] B. Zhao, Y. Duan, L. Huang, A. D. Joseph, and J. D. Kubiatowicz, Brocade: Landmark routing on overlay networks, *Proc. of the 1st Int. Workshop on Peer-to-Peer Systems (IPTPS)*, March 2002.
- [23] A. Bhambe, M. Agrawal, and S. Seshan, Mercury: Supporting scalable multi-attribute range queries, *Proc. of SIGCOMM, 2004*.
- [24] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana, Internet Indirection Infrastructure, *Proc. of ACM SIGCOMM, August, 2002*.
- [25] M. Grossglauser and D. Tse, Mobility increases the capacity of ad-hoc wireless networks. *Proc. of IEEE INFOCOM '01*, volume 3, pages 1360–1369, 2001.
- [26] O. Ileri, S. Mau and N. B. Mandayam, Pricing for Enabling Forwarding in Self-Configuring Ad hoc Networks, *IEEE JSAC*, vol. 23, No. 1, pp. 151-162, January 2005.
- [27] GENI: Global Environment for Network Innovations, Conceptual Design Project Execution Plan, <http://www.geni.net/GENI-10-JAN-06.pdf>