Next Century Challenges: RadioActive Networks

Vanu Bose
vanu@vanu.com
Vanu, Inc.

David Wetherall
djw@cs.washington.edu
CSE Department
University of Washington

John Guttag
guttag@lcs.mit.edu
Laboratory for Computer Science
Massachusetts Institute of Technology

Abstract

A key challenge facing wireless networking is to utilize the spectrum as efficiently as possible given current channel conditions and in the most effective way for each application. This is difficult to achieve with existing wireless devices because physical layer functionality is fixed, while channel conditions and applications can change rapidly. Instead, we argue that RadioActive networks, an adaptable wireless network architecture which draws on the strengths of software radios and active networks, are needed to meet this challenge. Active networks provide a framework for programming network services, and software radios extend this programmability into the physical layer. We believe that this approach will offer significant improvements in functionality and performance over today's wireless networks because it is no longer necessary to design *a priori* with pessimistic assumptions that describe worst case conditions. In this paper, we outline our vision, the opportunities it affords, and the challenges that must be tackled before it can become a reality.

1 Introduction

In a wireless network, channel conditions can vary significantly over time and tend to be difficult to predict, especially when nodes are mobile [ES96, DR92, NKNS96]. This is problematic because conventional communications designs are static in that they tend target a single design point that is the worst case that can be tolerated, rather than the channel conditions that are currently being encountered. The result is inefficient use of the available spectrum.

A network based on software radio technology enables a much more dynamic organization of resources. Software radios enable the characteristics of all communication layers, including the physical layer that is normally implemented in hardware, to be changed at essentially any time. While the motivation behind their development has been to solve the interoperability problems caused by different cellular standards, we believe that this technology is well-suited to wireless networking.

Recent research has shown that adaptive link layer technologies [ES98, ES, LS98, LFS97] can significantly enhance wireless network performance. A fully-programmable software radio can incorporate adaptive link layer techniques and extend the adaptability to the physical layer. This enables more effective use of the spectrum by dynamically adapting the physical layer of the network to best meet the current environmental conditions, network traffic constraints and application requirements, rather than a lowest common denominator service that must accommodate the worst case. We believe that this adaptation has the potential to significantly improve the performance of wireless networking systems, as well as enhance their functionality by taking into account different application requirements for bandwidth, latency, error rate, and security.

In order to adapt as we suggest, there must be a means of interaction between the physical layer and higher layers so that changing conditions can be observed and the most appropriate physical layer be selected. Our insight is that active networks provide a basis for achieving this adaptation. Active networks are a novel approach to network architecture in which the switches of the network perform customized computations on the messages flowing through them. By exchanging programs between the physical layer and higher layers, we gain a large degree of flexibility.

We call a network architecture based on the combination of these technologies a *RadioActive* network. Such a network presents new kinds of opportunities. For example, in a basestation to mobile system, the basestation can dynamically create channels depending upon the number of mobile units in its coverage area and their particular service requirements. When additional mobile units enter the area, the bandwidth can be appropriately apportioned to each unit. Units requiring real-time or high data rate services may be assigned a dedicated channel customized to their application, while others with bursty data requirements might be assigned to a shared channel. Furthermore, bandwidth can be apportioned to upstream and downstream channels depending on application needs. These kind of adaptations are typically not possible with conventional wireless networks.

In the remainder of this paper, we present our vision for RadioActive networks. We begin in Section 2 by describing an architecture that is based on a layered model of a wireless communication system and incremental dynamic modification of components. This is followed by a discussion of the opportunities that are presented by RadioActive networks in Section 3. We then discuss the challenges that must be tackled before we can realize these opportunities in Section 4. The software radio and active network technologies that

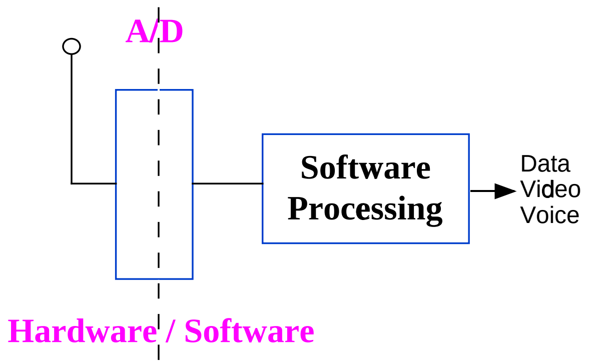

Figure 1: Ideal software radio.

can be used to realize a RadioActive network are described in Section 5 and we end the paper with concluding remarks.

2 RadioActive Networks: Freedom from the Worst Case

Our vision is to create a wireless communications network that is dynamically re-programmed as it is used to optimize for the current environmental conditions, traffic constraints and user requirements. This contrasts with conventional communication system design, which tends toward static functionality that is designed to tolerate the worst case conditions that might be reasonably expected to occur. In this manner, RadioActive networks provide freedom from worst case design.

To construct a RadioActive network, we require two key components: a model for describing the functionality of a wireless node, and a means for dynamically altering this functionality. Our key insight is that by combining software radios and active networks technologies we will be able to design, develop and explore novel wireless network protocols and applications that adapt to maximize their performance and offer new kinds of functionality. Active networks provide a framework for programming network services, and software radio technology extends this programmability into the physical layer. In the remainder of this section, we describe RadioActive networks along these two lines. The next two sections of this paper explore the opportunities and challenges posed by a system that is able to dynamically modify any of these layers. This is followed by a description of a testbed that enables the implementation and characterization of such systems.

2.1 Software Radios

A software radio is a wireless communications device in which some or all of the physical layer functions are implemented in software. The ideal software radio, illustrated in Figure 1, would directly digitize the entire band of interest and transport the stream of samples to memory where it can be directly accessed by a microprocessor. Current wideband receivers, A/D converters, I/O systems and processors cannot meet the requirements imposed by a direct implementation of this architecture. Current research in the area of software radios covers system design as well as the development of the enabling technologies of tunable wideband front-ends and the A/D converters capable of digitizing wideband signals with a signal-to-noise ratio sufficient to enable digital cellular applications.

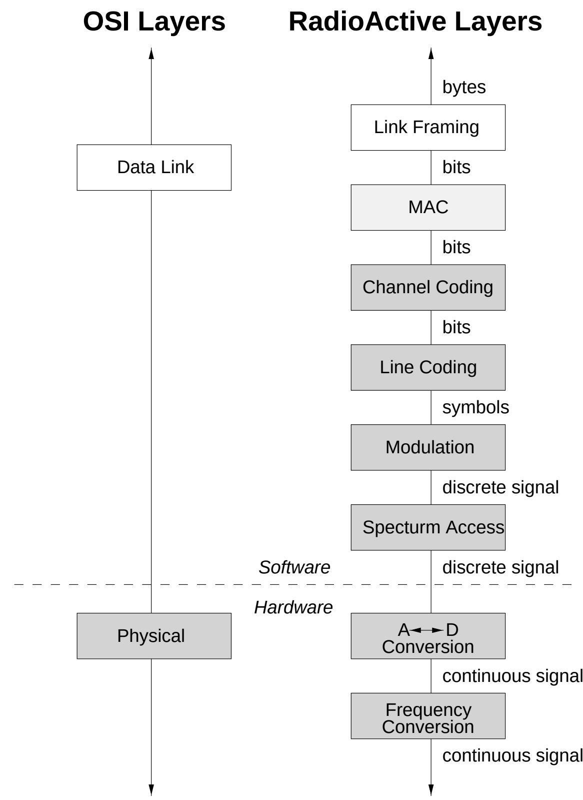

Figure 2: RadioActive Network Layering Model.

Most software radio research to date has been driven by the interoperability and deployment problems that result from the implementation of radios in dedicated hardware [Mit99]. The first significant software radio was the SpeakEasy system [LU95] which was designed to emulate more than ten different military radios. Commercial interest was spurred by an RFI from Bell South Cellular on “Software Defined Radio” [Blu95]. The principle commercial application driver in the U.S. is the multiple incompatible cellular and PCS communications, which inhibit universal roaming. In Europe, the widespread deployment of GSM has mitigated interoperability problems, but there is significant interest is using software radio to enhance services for third generation cellular systems [Tut99].

Although the primary motivation behind the development of software radios was to solve interoperability problems, we believe it provides a framework for the design of radioactive networks. The framework provided by traditional network layering models, such as the OSI model, are too coarse for our purposes. These models lump a wide array of different types of signal processing under the single heading of *physical layer*, which significantly limits the modification granularity.

The layering model shown in Figure 2 is a refinement of the OSI model that sub-divides the existing physical layer. This refinement exposes the different entities that comprise the physical layer and enables the modification of individual functional entities such as the channel coding or modulation without requiring changes to the rest of the physical layer.

The model also clarifies the role of each function performed by the physical layer. For example, the GSM system utilizes two forms of multiple access, where the spectrum is first divided into frequency channels and then sub-divided into eight time slots. Although they are both multiple access functions, they occupy different layers in the model. The frequency division is the manner in which the spectrum is shared and thus occupies the spectrum access layer. Time division is the MAC layer since it is the method for sharing a single channel.

There has been considerable research into adaptable functions and protocols for wireless networks. Most of these efforts involve adaptation in one of the sub-layers specified by the model in Figure 2. For example, adaptive FEC techniques, link layer retransmission, adaptive packet length control [ES98] and dynamic channel assignment [Ram97] are examples of functions that introduce adaptation into one of the layers.

One of the primary challenges in building a RadioActive network is understanding how changes in each of these individual layers interact with each other and affect the overall system performance. Table 1 qualitatively illustrates the principle areas in which the different layers impact network traffic parameters. Since multiple layers affect the same parameters, there are several ways in which the system could be modified in order to adapt to a given set of changes. In order to optimize the system for a given set of channel conditions, traffic loads and user requirements, it is necessary to have a quantitative framework for understanding the impact of changes in each of the layers on overall network performance. This would enable wide range of applications that provide better application level performance and better usage of the available RF spectrum.

2.2 Active Networks

Active networks are a novel approach to network architecture in which the switches of the network perform customized computations on the messages flowing through them. The concept emerged from discussions within the broad DARPA research community in 1994 and 1995, and was first described in [TW96]. Subsequently, a DARPA program was created to explore the potential of the active network approach, and today numerous research projects are underway [T^+97][1].

Active networks emerged in response to several problems with conventional networks: the difficulty of integrating new technologies and standards into the shared network infrastructure; poor performance due to redundant operations at several protocol layers; and the difficulty of accommodating new services in the existing architectural model. To address these problems, active networks exploit recently developed mobile code technologies to support dynamic network service innovation.

[1] See also http://www.darpa.mil/ito/research/anets/.

Layer	Bandwidth	BER	Latency	Jitter
MAC			X	X
Coding	X	X	X	
Modulation	X	X		
Channel MA			X	X

Table 1: Mapping between the physical layer components and their most significant effects on the network traffic parameters.

There is a synergy between software radios and active networks. Software radios provide a flexible communications infrastructure that is able to support a variety of different communications processing. Active networks provide a network architecture that takes advantage of programmable nodes in order to rapidly deploy new network services, or in the case of RadioActive networks, adapt the network services to the current operating environment.

3 Opportunities

RadioActive networks present a number of new opportunities to improve the performance of wireless networks as well as the utilization of available spectrum. These opportunities arise from the flexibility to change the physical layer processing coupled with the time-varying nature of both the mobile wireless environment and user demands on the network. In this section, we outline some of the interesting opportunities made possible by our vision. These opportunities take the form of new approaches and tools for the design of wireless network systems.

3.1 Physical Layer Support for QoS

Most quality of service work has focussed on the sharing of a network based on a static physical layer. Software radio technology can provide better quality of service in two different ways. An individual connection between two mobiles, or between a mobile and a basestation can optimize it's tradeoffs between throughput, BER and jitter based on application requirements. Alternatively, the network as a whole can support QoS by more flexibly partitioning the communications channel in order to support different quality levels.

Individual nodes in a network can tolerate different tradeoffs between latency, throughput error rate and jitter depending upon the application. Consider the case of a wireless network that utilizes frequency division to support multiple mobile nodes simultaneously. Each node is assigned a channel, but can dynamically modify the physical layer within the channel to best suit its needs. For example, if the background noise in the channel increases, a file transfer application may opt for a more robust channel coding algorithm at the expense of some latency, whereas a voice application may opt to modify the symbol constellation and drop to a lower data rate without incurring additional latency. Some existing systems, such as WaveLAN, utilize different data rates depending upon signal quality or distance from the hub. This is an example of modifying an individual connection in order to achieve better performance for the current conditions. RadioActive networks enables a far greater range of adaptability since any aspect of the system can be modified via software.

The parameters that can be modified by an individual node are limited by overall network constraints. For example, in many networks individual nodes are not free to arbitrarily modify their transmit power, since power control must be globally managed in many mutli-user networks to insure connectivity for all nodes. However, in some cases, such as TDMA networks, in may be permissible to modify power if only one node is transmitting at a given time.

In a basestation-mobile system, the ability to dynamically optimize the all aspects of the system can provide better application level performance for the mobile nodes as well as more efficient utilization of the available spectrum. As an example, consider again the conflicting requirements of voice and data communications: voice requires low latency constant bit-rate service, while data is typically bursty, with high peak throughput requirements and tends to be more latency tolerant. In a software radio network, a dedicated part of the spectrum could be assigned to support real-time traffic with guaranteed latency, and the remainder of the spectrum could be set up as a shared network that supports high transfer rates but requires arbitration to access. The network would simply view each of these channels as a distinct network link with specific characteristics, and make routing decisions based on how well these characteristics meet the needs of the application.

3.2 Self-Designing Networks

A more speculative strategy with a potentially larger payoff is to view the network as an ad hoc collection of nodes and to continually negotiate the most appropriate physical layers as the system evolves.

Ad hoc networking is the cooperative engagement of a collection of mobile computers [LP97]. The need for ad-hoc networks is often spontaneous, arising when a number of people gather in a particular location. Existing research in ad-hoc networks assumes the existence of a common physical layer over which to communicate. Software radio technology eliminates this dependence and allows the creation of a physical layer to best suit the needs of the individual participants given the spectrum availability. For example, the participants could monitor the spectrum to determine the available bands and then negotiate, through pre-defined robust control channels, the design of the physical network that utilizes the best available bands and incorporates the modulation and coding required to support the desired applications. Furthermore, the physical characteristics of the network can be re-designed as environmental conditions and user requirements change.

In addition to facilitating the creation of ad-hoc networks and optimizing performance, this approach can also improve the security of ad-hoc wireless networks. Jamming of a radio transmission is essentially denial of service at the physical layer and may be intentional or accidental. If security against jamming is a concern, then the physical layer can be designed to use spread spectrum techniques. In existing systems, if the spreading codes are known, the system can be jammed. If the spreading code in a software radio system is compromised, that particular instance of the system is vulnerable, but future instance of the system are not since the entire radio system, including the pseudo-random keys, is re-designed for each use.

3.3 Cross-Layer Optimizations

Because the physical layer has traditionally been treated as a black box, it has been difficult to optimize the performance of the network as a whole. For example, the physical layer may be ideally suited to transmit bursts of information of a particular length, or the power usage requirements may dictate certain transmission constraints. If the packet sizes are too long or too short this could lead to inefficient use of the channel. Communicating this information back to the network layer and applications that use it would allow data to be appropriately bundled and would result in better performance from the network link [LS98].

These effects may be significant, as cross-layer optimizations have been shown to be of substantial value. For example, the Snoop protocol [BPSK96] uses loss information from the wireless channel to augment TCP processing at wireless basestations. Since TCP interprets loss as congestion and this is often not the case for wireless channels, the impact on performance is considerable, and can be as great as 100% to 200%.

In a RadioActive network, there are redundancies in functionality between the physical layer and higher protocol layers. For example, the physical layer often includes some form of channel coding to reduce errors, but the network layer may also include mechanisms such as forward error correction on packets to reduce the error rate. Coordinating these efforts could lead to better error correction and/or a reduction in the overall number of bits transmitted over the channel.

An interesting approach would be to push all differences in error coding to applications, so that the channel performs "lowest common denominator" coding. In this manner, each application would be able to achieve the highest level of performance permitted by its acceptable error rate, however the application would need to have considerable knowledge of the channel in order to perform all of the coding. Joint source channel coding is a good example of a cross-layer optimization of this kind. If the channel conditions are fairly static, then the joint algorithm can be optimally designed. However, in wireless applications, where the channel varies significantly a single optimal coding scheme is not possible. The coding algorithm must be able to dynamically incorporate information about the channel and modify the coding algorithm appropriately.

4 Challenges

In order to realize the many opportunities presented in the previous section, there are several research challenges that must be addressed, including:

- The collection of information about the current state of the channel.
- Mechanisms for executing the changes to a network node.
- Policies for determining the appropriate modifications to the radio.
- Development of routing algorithms that incorporate information from the physical layer.
- An evaluation of the effect that a dynamic infrastructure might have on existing network protocols.

This section outlines some of the specific challenges that must be addressed in order to realize RadioActive networks.

4.1 Interface with the Physical Layer

Two mechanisms will be needed to allow the network node to program the physical layer. First there must be a way to provide information about the capabilities and current performance of the physical layer to the network layer. Then there must be a mechanism by which the network layer can specify and create the desired physical layer. The key challenge is the design of an interface that provides a simple programming model to the network, but captures the broad range of functionality that is implemented by the physical layer. The network parameters are derived from a combination of the node capabilities and the packet properties of the current connections.

4.2 Physical Layer Handoff

A physical layer handoff protocol will extend active network capabilities to the physical layer. This mechanism would allow for the downloading of all of the code necessary to implement the physical layer of a network connection from a trusted agent via a reliable communications protocol. A physical layer handoff protocol has been proposed in [Chi99]. The current physical layer is used to download the code for a new physical layer. Channels and protocols used for link initialization will have to be carefully addressed, as loss of connectivity in a given link could also mean the loss of the mechanism needed to upgrade the link.

4.3 Policies for Modification

Perhaps the biggest challenge that must be addressed is the development of policies that determine what modification should be made to the network and when they should occur. This requires gathering all of the pertinent information from the network and physical layers and then determining the physical layer functions that should be implemented. [ES98] demonstrates that effective adaptation policies exist for particular sub-layer functions, such as link layer FEC and packet trimming. Recent work on vertical handoffs in wireless overlay networks [WKG99] has addressed the issues of policies for choosing between complete networks to connect to when there are a number of different choices available. A RadioActive network extends the possibilities by essentially making available *any* network link, and the challenge is to then make the best choice from this space of possibilities.

An important consideration for policies is the timescale of modification. Changes in the condition of the channel occur on several timescales. For example, mobility may cause Rayleigh fading and hence changes in the channel that occur on a millisecond timescale, while changes due to human interaction, such as the beginning or end of a cell phone call, occur on the timescale of seconds. Environmental factors such as the weather may change on even longer timescales. Adaptation policies must be defined in light of these timescales if they are to be effective.

4.4 Routing in a RadioActive Network

In a RadioActive network, a single physical link can actually appear as multiple links to the network layer. For example, if a network link is physically partitioned into a dedicated channel for supporting real-time data and a shared channel for less time sensitive data, this appears as two separate network links between the two nodes. These links have different characteristics, which can be used to appropriately route different types of traffic through the network. Other information from the physical layer, such as the current bit error rate, security parameters and battery reserves can also be used by routing algorithms to handle heterogeneous traffic or optimize network wide qualities such as power dissipation or available level of security.

4.5 Interaction with Network Protocols

A necessary area of study is the interaction of existing end-to-end network protocols with a dynamically changing physical layer. Experience with vertical handoffs in overlay networks suggests that occasional changes in the physical layer do not pose a significant performance problem [SK98]. [ES98] also reports that adaptations in the link layer are compatible

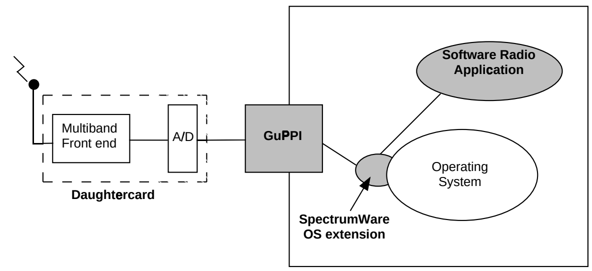

Figure 3: SpectrumWare System Architecture

with excellent TCP performance. These results are encouraging. However, care must be exercised in a RadioActive network because individual links may change more rapidly and may cause more variation in the routes that packets traverse. These effects should be studied in both simulation and testbed situations to determine the appropriate limits to place on the nature and rate of adaptability.

5 Experimental Plan

Considerable work must be done to support software radio capabilities in a mobile, hand-held device. Solutions such as reconfigurable computing [Rab97] and methods for controlling the power dissipation of processors [GC96] hold the promise of making hand-held software radio a reality. The development of low power A/D converters capable of digitizing the entire cellular band with enough resolution to implement the digital cellular standards in software is another significant research challenge.

Fortunately, research and development of RadioActive networks does not have to wait for the development of these technologies. The combination of the SpectrumWare [Bos99] software radio system and the ANTS [Wet99] Active Network toolkit provide an excellent prototyping and measurement platform which will permit the development of RadioActive network applications which can then be ported to hand-held devices as they emerge. We have developed both of these systems separately over the past two years, and have now begun the process of combining them. In the remainder of this section, we introduce these two enabling technologies and describe ongoing research that is relevant to RadioActive networks.

Our inspiration in building this experimental platform is Softnet [ZF83]. Softnet was a distributed experimental packet radio network developed in in Sweden in the early 1980s. The goal of Softnet was to allow users to define their own high level services (such as datagrams, virtual calls, file transfers and mailboxes) as well as to allow changes at the lower level of link access protocols. Packets in Softnet are considered to be programs of a network language. While much of the physical layer was fixed in hardware due to technological constraints of the day, Softnet is clearly similar in spirit to our RadioActive network, and is arguably the first RadioActive network.

5.1 SpectrumWare

The SpectrumWare software radio [BIWG99] demonstrates the feasibility of using a general purpose processor coupled with wideband digitization to implement a software radio.

The SpectrumWare system architecture shown in Figure 3 makes only one concession to the ideal architecture depicted in Figure 1, which is to first downconvert a wideband of the spectrum to an IF frequency and then digitize it. The center frequency of the RF band is selectable in software and it is important to note that what is downconverted is not just a single channel but a wideband (e.g. 10 - 20 MHz). For most systems this enables dynamic modification of the multiple access protocol, since it is implemented in software. Other than the restriction of looking at one particular band at a time, the system is functionally equivalent to the ideal software radio.

The system has two primary components. The signal acquisition and I/O sub-systems digitize a wideband of the RF spectrum and place it in memory that can be accessed by the processor. The programming environment is a modular system that allows for data-intensive real-time signal processing applications to be created and dynamically modified.

The system has been used to implement several different wireless communications systems, including a 4-channel AMPS cellular receiver, a wireless 2.5 GHz network link employing frequency hopping and a virtual patch between to incompatible radios: an analog cordless telephone and a CB radio. More details on the SpectrumWare system and applications can be found in [BS98, BIWG99].

Something about Andrew's stuff.

5.2 ANTS

ANTS, a prototype active network [Wet99], demonstrates how programmability can be used in the network at a fine granularity. It does this by adding service extensibility to Internet-style networks at the packet level. In ANTS, the packets of traditional networks are replaced with capsules that refer to a short program to be executed to forward them at active network nodes. Forwarding programs are transferred around the network to the nodes where and when they are needed by using mobile code techniques. Applications are able to construct and use novel network services at any time. Measurements of the toolkit show that the forwarding mechanism requires little processing beyond that of IP, such that even the user-level Java toolkit runs at 10 Mbps Ethernet rates.

In terms of RadioActive networks, ANTS provides a good framework for expressing adaptations in communications system processing. For example, packets sent over a RadioActive network might route themselves over different virtual channels by using different channel codings, depending on the application requirements. Further, as a concrete system that runs on commodity hardware, ANTS is well placed as a platform to explore the flexibility offered by software radios systems such as SpectrumWare.

6 Conclusions

In this paper we have argued that in order to provide better application performance and make more efficient use of the RF spectrum, wireless networks must incorporate adaptability into all aspects of a communications link. This is because the channel conditions that are encountered in a wireless network vary significantly and unpredictably over time, to a much greater extent than wired networks. The worst conditions that might be encountered can be orders of magnitude worse than a typical operating environment. Thus wireless networks and protocols statically designed for the worst case conditions do not make the best use of available resources. Further, networks carry an increasingly heterogeneous mix of traffic, including traditional bursty and elastic data communications, interactive Web traffic, and real-time sample streams.

The combination of large unpredictable variations in the channel, heterogeneous traffic and changing topology due to mobile nodes poses a significant challenge to the design of a system that can be optimized for the current conditions. To meet this challenge, we have proposed that RadioActive networks draw on the strengths of software radio and active network research. As wireless networking matures, it will become more important to make the most effective use of the available spectrum. We have argued that software radios and active networks provide a good basis for doing so because they provide the greatest level of flexibility. Further, existing research in both software radios and Active Networks forms the basis for a prototype that will allow the RadioActive network concepts to be evaluated in parallel with work on the technologies required to make them practical.

Acknowledgments

The ideas presented in this paper have benefited from discussions with a number of our colleagues. We would like to thank the members of the SpectrumWare team, in particular Andrew Chiu, as well as Srini Seshan, Hari Balakrishnan and David Tennenhouse.

References

[BIWG99] Vanu G. Bose, Michael Ismert, Matthew Welborn, and John Guttag. Virtual Radios. *JSAC issue on Software Radios*, April 1999.

[Blu95] Stephen M. Blust. Software Defined Radio - Industry Request for Information. Technical report, Bell South Cellular, December 1995.

[Bos99] Vanu G. Bose. *Virtual Radios.* PhD thesis, Massachusetts Institute of Technology, February 1999.

[BPSK96] H. Balakrishnan, V. N. Padmanabhan, S. Seshan, and R. H. Katz. A Comparison of Mechanisms for Improving TCP Performance over Wireless Links. In *SIGCOMM '96*, pages 256–269, Stanford, CA, August 1996. ACM.

[BS98] Vanu G. Bose and Alok B. Shah. Software Radios for Wireless Networking. In *Infocomm'98*, San Fransisco, 1998. IEEE.

[Chi99] Andrew G. Chiu. Adaptive Channels for Wireless Networks. Master's thesis, M.I.T., May 1999.

[DR92] D. Duchamp and N. Reynolds. Measured Performance of a Wireless LAN. In *IEEE 17th Conf. on Local Computer Networks*, September 1992.

[ES] D. Eckhardt and P. Steenkiste. A Trace-based Evaluation of Adaptive Error Correction for a Wireless Local Area Network. *Mobile Networks and Applications (MONET).* To Appear.

[ES96] D. Eckhardt and P. Steenkiste. Measurement and Analysis of the Error Characteristics of an In Building Wireless Network. In *SIGCOMM'96*, San Francisco, CA, August 1996.

[ES98] D. Eckhardt and P. Steenkiste. Improving Wireless LAN Performance via Adaptive Local Error Control. In *ICNP'98*, Austin, TX, October 1998.

[GC96] Vadim Gutnik and Anantha Chandrasakan. An Efficient Controller for Variable Supply-Voltage Low Power Processing. In *VLSI'96*, 1996.

[LFS97] P. Lettieri, C. Fragouli, and M. Srinastava. Low Power Error Control for Wireless Links. In *Mobicom97*, Budapest, Hungary, September 1997.

[LP97] Hui Lei and Charles E. Perkins. Ad Hoc Networking with Mobile IP. In *EPMCC*, 1997.

[LS98] Paul Lettieri and Mani B. Srivastava. Adaptive Frame Length Control for Improving Wireless Network Link Throughput, Range and Energy Efficiency. In *Infocom'98*, March 1998.

[LU95] Raymond J. Lackey and Donal W. Upmal. Speakeasy: The Military Software Radio. *IEEE Communications Magazine*, 33(5):56–61, May 1995.

[Mit99] Joe Mitola. Software Radio Architecture A Mathematical Perspective. *JSAC issue on Software Radios*, 1999.

[NKNS96] G. T. Nguyen, R. H. Katz, B. D. Noble, and M. Satyanarayanan. A Trace-based Approach for Mofdelling Wireless Cahnnel Behavior. In *Winter Simulation Conference*, December 1996.

[Rab97] Jan M. Rabaey. Reconfigurable Computing: The Solution to Low Power Programmable DSP. In *1997b ICASSP Conference*, Munich, April 1997.

[Ram97] S. Ramanathan. A Unified Framework and Algorithm for Channel Assignment in Wireless Networks. In *INFOCOM'97*, Kobe, Japan, 1997.

[SK98] Mark Stemm and Randy Katz. Vertical Handoffs in Wireless Overlay Networks. *Mobile Networking and Applications*, 3(4), 1998.

[T^+97] David Tennenhouse et al. A Survey of Active Network Research. *IEEE Communications Magazine*, pages 80–86, Jan 1997.

[Tut99] W. Tuttlebee. Software radio technology: A european perspective. In *IEEE Communications Magazine*, February 1999.

[TW96] D. L. Tennenhouse and D. Wetherall. Towards an Active Network Architecture. In *Multimedia Computing and Networking 96*, San Jose, CA, January 1996. A revised version appears in CCR Vol. 26, No. 2 (April 1996).

[Wet99] David J. Wetherall. *Service Introduction in an Active Network*. PhD thesis, Massachusetts Institute of Technology, February 1999.

[WKG99] Helen J. Wang, Randy H. Katz, and Jochen Giese. Policy-Enabled Handoffs Across Heterogeneous Wireless Networks. In *WMCSA 99*, New Orleans, LA, 1999.

[ZF83] J. Zander and R. Forchheimer. Softnet - An Approach to High-Level Packet Communication. In *ARRL 2nd Computer Networking Conference*, San Francisco, CA, March 1983.