EXPERIMENTATION AND ANALYSIS FOR UNIFIED PACKET-BASED WIRELESS NETWORKS

Jay R. Moorman
jrmoorma@uiuc.edu
University of Illinois

John W. Lockwood
lockwood@arl.wustl.edu
Washington University, St. Louis

Abstract

This paper looks at the next generation of wireless data networks. It experiments with and analyzes the current state of wireless data communication, including the setup of an operating wireless testbed. The paper then explores the next step for wireless networks where all types of traffic including voice, multimedia, and data are sent over a common underlying packet-based network. It continues with a discussion of relevant issues needed for the transition. The paper concludes with a QoS enabled scheduling algorithm designed to address many of these issues.

1 Introduction

The current state of wireless communications can be divided into two broad categories based on the designed purpose of particular devices. These two categories are either user-centric devices which support a particular user function, or network-centric devices which add infrastructure support to the network. This paper will be focusing on the user-centric wireless devices. Network-centric devices, such as satellite communication or building-to-building LAN extensions, involve many overlapping issues with user-centric systems.

End user (user-centric) wireless communication consists of three major areas of service as shown in Figure 1. The first area most commonly associated with wireless communication is that of the wireless cell-phone. With the growth of this market exploding in recent years, this business sector has extensive leveraging power for building the next generation of protocols and devices. The second area which is beginning to emerge as a viable wireline alternative is that of the wireless network LAN. This technology continues to lag the wireline domain by an order of magnitude, but is very inviting for the mobile computing user. The final area of wireless communication that is often overlooked is that of the computationally reduced embedded wireless device. This rapidly growing area includes such items as cordless phones, infrared remotes, wireless doorbells, and even garage door openers.

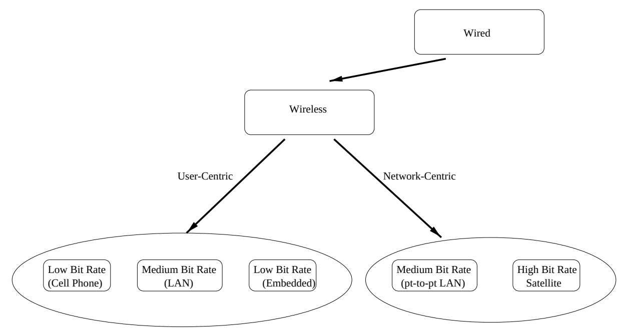

Figure 1: Wireless Communication Overview

As time progresses, we expect these three areas to converge. Wireless voice communication will continue to add data functionality with higher bandwidths for Wide Area Network (WAN) coverage. Wireless LAN's will incorporate voice traffic, and take over command and control type functionality of the embedded devices. The embedded wireless device will also become more and more intelligent, capable of executing highly advanced protocols in a less limited way. These three areas will continue to converge into a ubiquitous wireless computing environment. However, there still remains many challenges, and much work both technically and politically, before such a communication and computing environment might exist.

This paper begins the wireless communication journey in the area of wireless LAN's. It then shows how the other two areas can begin to be folded together while pushing forward with higher data rates in a scenario very similar to the wired computer network. The paper continues by presenting the current state of a wireless testbed, including the tools and tricks used to manage this network. Some measured results are given for this current network setup under a number of different scenarios. In light of these results, a few key issues are highlighted for the next stage in wireless communication development. Finally, a scheduling algorithm is briefly discussed in order to address some of the above named issues.

This paper is organized as follows. Section 2 discusses the wireless communication model and presents the path of convergence for current wireless communication. Section 3 describes an operating wireless testbed. In Section 4 some measured results of the testbed are presented along with some observations on the use of the testbed. Section 5 discusses issues for the transition to the next generation of wireless communication. Section 6 provides some information on a scheduler designed to help in this transition to the next generation of wireless networks. Finally, Section 7 concludes the paper.

2 Communication Model

There are a number of important issues that must be addressed when communicating over a wireless channel. Some of these key items are discussed below to present a better understanding of the communication model and its particular challenges.

The first and most important issue for wireless communication is the issue of wireless error. Unlike the wireline case, wireless traffic is much more prone to experience some type of channel error that disrupts the normal flow of communication. This can occur for any number of reasons including multipath fading, shadow fading, or simply straying out of range. Because these situations are unique to the wireless environment, they are the major reason that the wireless channel must be treated differently than the wired link. This is the primary area of focus for our work on wireless networks.

In the wireless environment, there are additional problems such as hidden terminals. Since the medium does not limit an interface to a single sender per time slice, multiple transmissions can be exchanged if care is taken to avoid interference with the communication of other transmissions. In other words, on the wired channel, a single sender will capture the channel to transmit and all other devices must wait for another transmission frame. Whereas on the wireless channel, multiple point-to-point messages are possible using the same medium, frequency, and time slice. The situation is particularly difficult to manage in an ad hoc scenario [1]. Typically, wireless access has been addressed through the use of RTS-CTS like mechanisms to block channel access by any devices that might interfere.

With the development of wireless devices, the issue of mobility is pushed to the forefront. Mobile communication forces the network provider to deal with additional issues such as handoff and tracking/location management. This enables mobile users to retain communication as they cross various wireless communication boundaries. Solutions to this problem involve the use of standard protocols such as Mobile IP [2], along with custom management software.

The most desirable aspect of wireless communication is the ability to perform some function without being bound to an exact location. It is still the case, however, that most of the resources are available on wired networks. What this leads us to conclude is that the most useful purpose of wireless communication is a link to wired communication, whether it is the telephone network, an ATM backbone connection, or access to an office LAN. This link needs a single hop wireless communication that is as efficient as possible. One way of approaching this is to focus on extending the wired network one additional hop for access via the wireless medium.

As this evolution of technology converges, the most important item for this unified network is Quality of Service (QoS) support. QoS mechanisms are needed to distinguish the various types of traffic, and their levels of service in a single communication network. In other words, QoS support will be the glue that holds the various network pieces together.

3 Wireless Testbed

The iWANDeR [3] wireless testbed that has been setup consists of a number of wireless LAN components coupled with a set off software tools customized through individual scripts. This has enabled us to experiment with a real wireless data network. Additionally, it has provided a means to obtain some quantifiable results from this operating communication system.

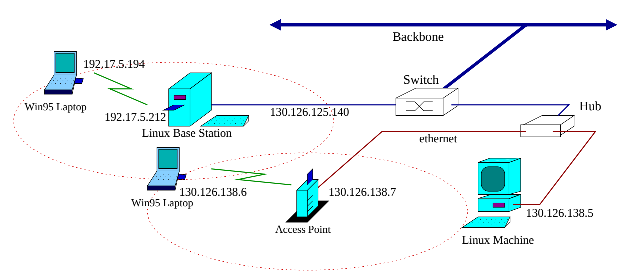

Figure 2: Wireless Testbed

The hardware of the testbed is comprised of mobile laptops that communicate to a fixed base station via wireless capable PCMCIA cards. The base station also has a fixed network connection to provide extended wired service to the wireless laptops. The PCMCIA cards are Netwave AirSurfer Plus cards [4] operating at a rate of 1 Mbit/s. Each card uses the 2.4 GHz frequency hopping spread spectrum for communication. These cards are also known as the BayStack 650 Wireless PC Card. The communication protocol employed is both the older proprietary Netwave protocol, and the newer 802.11 firmware upgrade. Mobile laptops have been setup using both the Linux and Windows operating systems. The original base station was a Netwave/BayStack Access point designed for the sole purpose of supporting a single wireless cell. A device driver was developed to allow the base station functionality to be ported over to a standard PC. This PC base station consists of an AMD K6-2 based machine with both a PCMCIA wireless network interface and a wired network connection. The machine runs the Linux operating system with the 2.2 kernel. A diagram of the testbed setup can be seen in Figure 2.

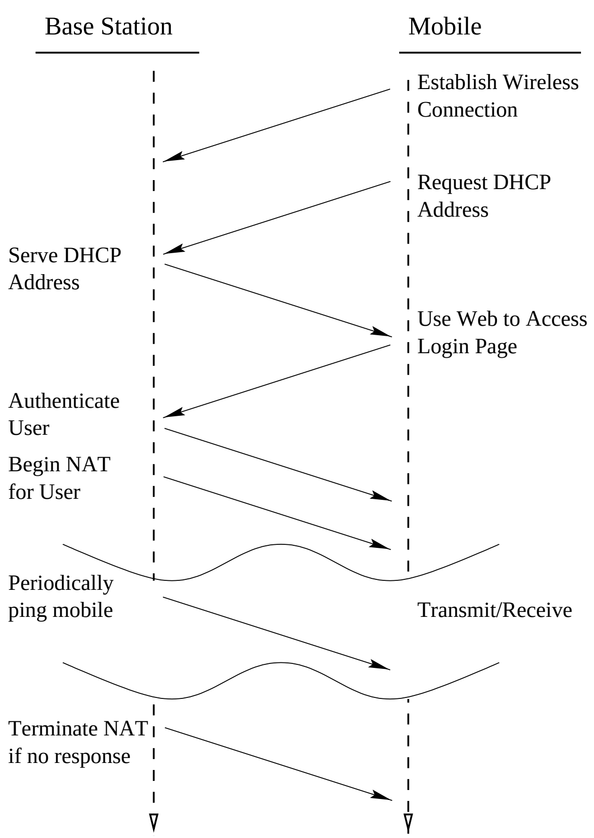

Figure 3: Testbed Communication

With this initial setup a mobile user can access the base station. To provide additional network connectivity the base station also performs IP forwarding. In order to have an environment that is more easily configurable, some additional base station services are provided. An IP address is dynamically served to new mobile stations, and then managed by the base station. These mobiles are assigned transient IP addresses, allowing the base stations to perform masquerading for access to non-local traffic. This scenario also enables the base station to serve as a firewall for traffic routing. Thus unauthorized users or mobiles can be denied access, or only provided with a limited set of resources by the base station. The base station is also responsible for terminating the Network Address Translation (NAT) capabilities when the laptop is turned off or leaves the area.

A better understanding of the system is achieved by looking at the communication transactions involved in the testbed as shown in Figure 3. When the laptop is brought up, the PCMCIA device is configured to attempt a connection with another card. This connection attempt uses the point-to-point Netwave protocol with a particular frequency hopping pattern. Once the connection has been established, a DHCP request is sent to the base station by the mobile. The base station then replies with an acknowledge and a temporary NAT address lease for the mobile to use. At this point, the interaction requires user intervention in order for the base station to authenticate the user. Alternatively, it would also be possible to configure the NAT to trust specific MAC addresses from particular PCMCIA cards. However, for less hardware dependent operation, and to alleviate problems from theft, we have chosen to perform user level authentication.

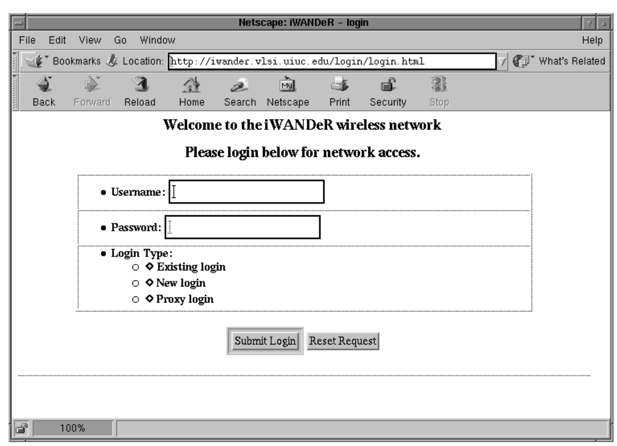

Figure 4: Login Web Page

A login web page is served by the base station as seen in Figure 4. The user then logs in to the network using this page. If the base station can authenticate the user, it will run a script to begin address translation for the assigned mobile address. This script uses the *ipchains* package in Linux to forward packets belonging to a particular IP address. The current setup will forward all packets once the user has been authenticated. It can, however, also be configured to only provide packet forwarding to a limited part of the network. For instance, some users might be limited to an office LAN, a university domain, or possibly a reduced set of services. This restricted access can be kept track of in an mSQL database which is initialized during a new login request. The login request web page is shown in Figure 5. New users requesting wireless access must be authenticated off-line by name, social security number, and group association.

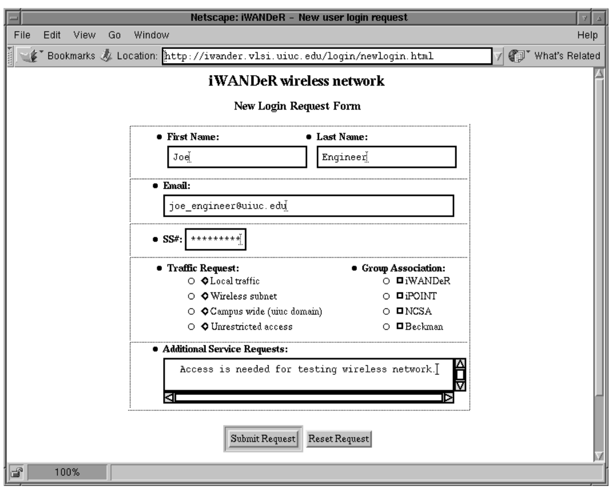

Figure 5: Account Setup

The final responsibility of the base station is to make sure the address translation is properly terminated once the mobile is no longer active. This can be triggered off the expiration of the DHCP address lease, or more simply by periodically using *ping* to verify the mobile is still alive and active. We have chosen the latter such that a separate process is spawned whose sole purpose is to track the mobile until it has finished communicating and been disconnected.

4 Wireless Measurements

The working testbed has been used to perform measurements on the wireless channel. This section presents a cross section of some of the data that was gathered during our testing.

In our first set of experiments the bandwidth of the established wireless link was measured over varying locations as shown in Figure 6. Bandwidth measurements reported are with respect to a single connection at the base station using the original Netwave protocol on the 1Mbit/s PCMCIA radio devices. The data was taken with varying distance measurements. This includes data taken with almost no separation, with slight separation, across cubicles, across the room, outside the room, and across the building. As the distance increased so did the number of objects, walls, and miscellaneous interfering items. What was most noticeable in this situation is that the relative throughput remained fairly stable until reaching a threshold distance/location that caused a significant reduction in the bandwidth. Once this interference point was reached, the achieved throughput was greatly degraded. The wireless channel errors were also seen to quickly increase at these positions.

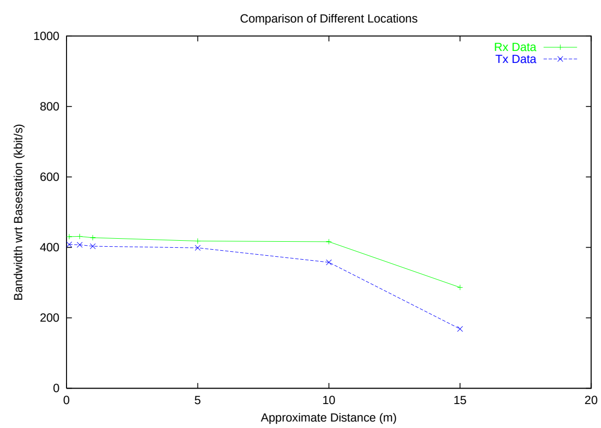

Figure 6: Bandwidth Comparison of Location

The second set of experiments shown in Figure 7 measured the throughput for a single connection using different protocols. These include the Netwave protocol under Linux with both UDP and TCP traffic, the 802.11 protocol under Windows, and normalized TCP traffic on a wired ethernet connection for comparison. It was found here that the connection nature of TCP traffic caused a slight degradation in throughput over the connectionless UDP. It was also determined that the newer 802.11 protocol runs slightly more efficiently than that of the older Netwave protocol. Finally, comparison with the wireline showed that the wireless nature of the link does cause a significant reduction in the achieved bandwidth. All protocols measured suffer from a large hit for moving traffic through the protocol stack. This has an affect of reducing the maximum achievable bandwidth by over 33%.

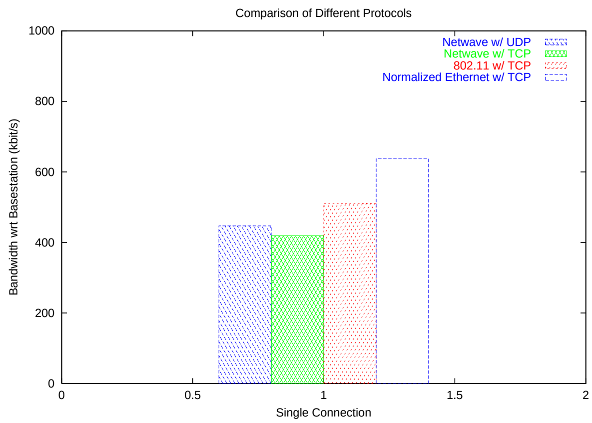

Figure 7: Bandwidth Comparison of Protocols

The final experiment shown in Figure 8 measured the achieved throughput as the number of connections to the base station increased. The aggregate throughput stayed approximately constant as the number of connections varied. However, the individual connections received a reduced, but relatively equal portion of the bandwidth compared to the other connected flows. This graph most of all demonstrates that best effort mechanisms are implemented properly in today's network, but are insufficient to support different service levels for different traffic types. It also suggests that other algorithms will be needed when some flows have requirements that cannot be reduced. These requirements call for some improved scheduling techniques in order to support the different types of traffic over the wireless network.

5 Next Generation Wireless

In the next generation of networks we will begin to see the true convergence of voice, multimedia, and data traffic. The rapid growth of the internet motivates providing access to these services for mobile users. This merging of various dedicated networks will occur both in the wired and wireless arenas. In this time of rapid fundamental changes in how we communicate, there exits a unique opportunity for the wireless world to take the lead in presenting effective solutions that can carry this information to the end user. In fact, the wireless protocols might well be able to bypass wireline technology in providing access to a unified multi-service network.

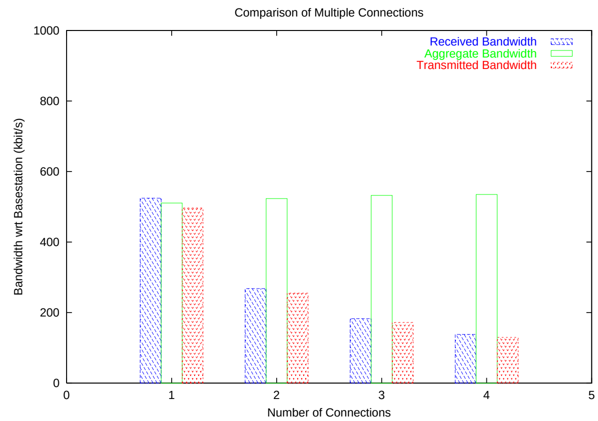

Figure 8: Bandwidth Comparison of Multiple Connections

In order to support such a wide range of traffic on a network, the infrastructure must be capable of coping with conflicting traffic requirements. It must be able to support low bit-rate traffic with minimum delay, real-time traffic with delay guarantees, and varying bit-rate traffic with low loss rates for e-mail, ftp, etc. An ideal solution will change as little existing standards as possible. QoS mechanisms are needed to distinguish the various types of traffic, and their levels of service in a single network.

The QoS support includes the call admission and the subsequent scheduling of packet transmissions. It is difficult to schedule packets in order to meet QoS guarantees, especially in the presence of wireless error and channel fading. However, the next generation of wireless networks will need to rely on a packet scheduler that can provide these traffic guarantees.

6 MPFQ Scheduler

The treatment of every flow as an equal partner in the communication schedule begs the need for additional scheduling criteria in order to support different types of traffic. In particular, to support QoS for various traffic over a wireless link, a new and smarter scheduler must be developed for fairer channel access with guaranteed bounds on bandwidth and delay.

The design of an improved scheduler is a critical step in the growth of wireless networks. To improve upon the current packet scheduling, we began work on the Multiclass Priority Fair Queuing (MPFQ) scheduler. This scheduler is designed to be part of a complete QoS mechanism in the base station. This allows it to fit into a broader framework such as in the scheduling multiplexer and QoS sublayer of [5].

The MPFQ scheduler is based on a multiclass priority algorithm that separates flows according to each general class of traffic as shown in Figure 9. These classes, as specified in [6] for backbone networks, include constant bit-rate (CBR) and real-time variable bit-rate (rt-VBR) at the higher priorities, as well as non-real-time variable bit-rate (nrt-VBR), available bit-rate (ABR), and unspecified bit-rate (UBR) at the lower priorities. The scheduler adheres to strict priorities between levels which is both desirable and necessary as shown in [7]. Within each level, a specific scheduling method is used depending on the necessary traffic requirements. By extending the backbone QoS parameters onto the wireless communication link, the end-to-end QoS requirements can be satisfied and fairness can be maintained for delay sensitive traffic.

The MPFQ algorithm defines the mapping of the traffic classes into priority levels and weights at call admission. Real-time traffic uses a fair queuing algorithm specifically adapted to the wireless medium [8]. The non-real-time traffic uses a Weighted Round Robin (WRR) scheduler with an adaptation for wireless compensation. The remaining traffic receives best effort service.

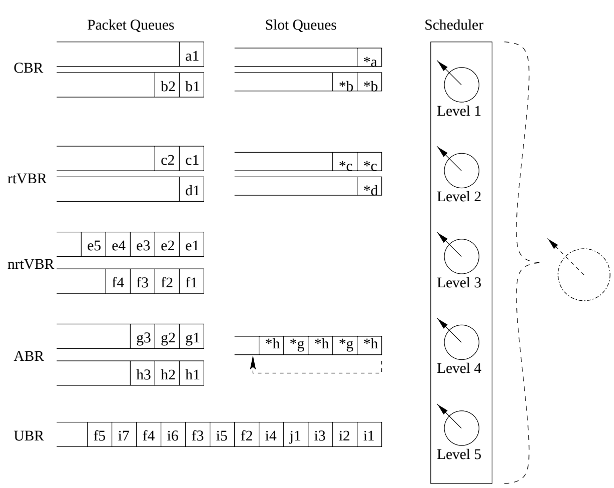

Figure 9: MPFQ Base Station Scheduler

The hierarchy of classes consists of five separate priorities. The first two priorities are used for real-time data that requires tight bounds on delay. These delays are satisfied and the channel allocated fairly through the use of a wireless fair queuing model [7]. This model includes tagging packets with start and finish times so that the schedule will be fair, will meet all deadlines, and will correctly allocate the required level of bandwidth.

The third priority in the system is used for flows that require guaranteed bandwidth, but not tight delay bounds. MPFQ uses a simple WRR scheme to ensure that traffic bounds are met over time. Traffic will be entered into per flow queues as it arrives at the base station. These queues will then be accessed when packets are needed for the outgoing queue. However, this type of traffic is sensitive to loss and requires a modification to the scheme. If the flow detects a channel error, then its packet will be skipped in that round. At a later time, the flow will have a chance to send the packet when the channel is error-free. The third priority uses the WRR scheduling to access both nrt-VBR and ABR traffic at the Minimum Cell Rate (MCR).

The fourth priority in the system also does not require tight delay bounds. It compensates errored flows by recirculated slots for a later time when the channel is good. This is done by keeping the incoming packets in separate packet queues and using a single slot queue to reserve slots for the head-of-line (HOL) packets. Incoming data is added to the per-flow packet queue and a slot is added to the recirculating queue. This slot has a pointer to the packet queue as required by the scheduler. If the HOL slot must wait, due to channel error, it is entered into the back of the recirculating queue, reserving space for later transmission. Once the channel is again error-free, the first slot in the outgoing queue is used to transmit the HOL packet in the flow's packet queue. By recirculating the slots in this way we provide wireless compensation without packet loss.

The lowest priority in the system processes flows that require no guarantees in either bandwidth or delay. Packets are entered in the order of arrival. An arriving packet is placed in the outgoing queue to be sent in FIFO order.

By scheduling flows within prioritized traffic classes, the MPFQ algorithm distinguishes itself from other wireless algorithms. This differs from class-based queueing (CBQ) [9] by supporting different scheduling policies at different levels, using ATM QoS parameters to specify flow requirements, keeping per-flow state, and compensating flows for channel error on the wireless medium.

The result is a scheduler that can provide flows with a bandwidth and delay guarantee for the last hop in the communication link. This wireless link guarantee can be combined with a wired guarantee in order to provide an individual flow an end-to-end delay and bandwidth guarantee. Thus the ability to provide QoS to a flow can be provided by the hybrid network.

7 Conclusion

A wireless testbed has been constructed to study the elements of future wireless packetized networks. A number of tests were run to gather statistics about the current state of wireless networks. Results of this work was presented. It was shown that today's wireless LAN devices work well only for best-effort packet flows. Other scheduling algorithms such as MPFQ will be needed to provide service for real-time and variable-bit rate flows.

The MPFQ algorithm was presented for supporting bandwidth and delay guarantees in a fair manner. This algorithm, operating in a base station as the wireless packet scheduler, enables mobile stations to seamlessly communicate with a backbone network using QoS-enabled traffic flows.

References

[1] C.-K. Toh, C.-H. Shih, V. Vassiliou, and M. Delway, "Emerging and future research directions for mobile wireless ATM networks," in *WCNC '99*, (New Orleans, LA), Sept. 1999.

[2] C. Perkins, "IP Mobility Support." RFC 2002, Oct. 1996.

[3] Information Available via WWW at `http://iwander.vlsi.uiuc.edu/wireless`, 1999.

[4] "CreditCard Netwave Adapter Application Program Interface (API) Specification." Netwave Technologies, Inc. Part No. 7000-1001, Version 1.48, July 1998.

[5] S. Sen, A. Arunachalam, K. Basu, and M. Wernik, "A QoS management framework for 3G wireless networks," in *WCNC '99*, (New Orleans, LA), Sept. 1999.

[6] "Traffic Management Specification v4.0." ATM Forum Document AF-TM-0056.000, Apr. 1996.

[7] J. Moorman and J. Lockwood, "Multiclass priority fair queuing for hybrid wired/wireless quality of service support," in *WOWMOM '99*, (Seattle,Washington), pp. 43–50, Aug. 1999.

[8] T. Nandagopal, S. Lu, and V. Bharghavan, "A unified architecture for the design and evaluation of wireless fair queuing algorithms," in *MOBICOM '99*, (Seattle, Washington), pp. 132–142, Aug. 1999.

[9] S. Floyd and V. Jacobson, "Link-sharing and resource management models for packet networks," *IEEE/ACM Transactions on Networking*, vol. 3, Aug. 1995.