

Cross-layer Design for Wireless Networks

Sanjay Shakkottai, Theodore S. Rappaport and Peter C. Karlsson*

June 23, 2003

Abstract

As the cellular and PCS world collides with Wireless LANs and Internet-based packet data, new networking approaches will support the integration of voice and data on the composite infrastructure of cellular base stations and Ethernet-based wireless access points. This paper highlights some of the past accomplishments and promising research avenues for an important topic in the creation of future wireless networks. In this paper, we address the issue of cross-layer networking, where the physical and MAC layer knowledge of the wireless medium is shared with higher layers, in order to provide efficient methods of allocating network resources and applications over the Internet. In essence, future networks will need to provide "impedance matching" of the instantaneous radio channel conditions and capacity needs with the traffic and congestion conditions found over the packet-based world of the Internet. Further, such matching will need to be coordinated with a wide range of particular applications and user expectations, making the topic of cross-layer networking an increasingly important one for the evolving wireless build-out.

1 Introduction

There continues to be rapid adoption of wireless technology, and when coupled with the explosive growth of the Internet, it is clear that there will be increasing demand for wireless data services [1]. Traffic on future wireless networks is expected to be a mix of real-time traffic such as voice, multimedia teleconferencing, and games, and data-traffic such as WWW browsing, messaging and file transfers. All of these applications will require widely varying and very diverse quality of service (QoS) guarantees for the different types of offered traffic, and we are now in the early days of this eventual amalgamation. Various mechanisms have been proposed and recently deployed to support data traffic over wireless media. This has ranged from Wireless Local Area Networks (WLANs), mainly based on the IEEE 802.11b or HiperLAN standards, to wireless wide area networks (WWANs) where data services are supported in the 2.5G and 3G system versions.

In this paper, we explore some key issues which will likely impact future cross-layer network design. We begin with an overview of data services over networks (Section 2). In Section 3, we describe some

*Shakkottai and Rappaport are faculty in the Wireless Networking and Communications Group (WNCG), Department of Electrical and Computer Engineering, The University of Texas at Austin, shakkott@ece.utexas.edu, wireless@mail.utexas.edu. Peter C. Karlsson is with the Wireless Solutions Group at Telia Research AB, Sweden, peter.c.karlsson@telia.se.

fundamental differentiators between wireline and wireless networks. Then, in Section 4, we consider the gains that can be achieved by means of a cross-layer approach for wireless network design, where physical layer information is passed to the higher layers, and discuss some of the standardization efforts in Section 5 that are beginning to move toward this integrated cross-layer approach. We finally conclude with a discussion of future research challenges.

2 Network Service Types

The Internet today consists of thousands of access networks, which vary in scale from large wireline access networks, such as campus wide area networks or Internet Service Providers, supporting tens of thousands of users, to smaller wireless access networks supporting tens to hundreds of users. These access networks are interconnected by core networks (such as AT&T, WorldCom, or Sprint backbone networks) which support hundreds of millions of users. All of these networks primarily offer two types of services: **guaranteed service** and **best effort service**.

In guaranteed service, the network provides some sort of service guarantee to individual users or groups of users. These guarantees are often in the form of ensuring that the throughput for a group of users is greater than some minimum value or that the delay experienced is smaller than some threshold. Often times, the guarantee is offered in a statistical sense, e.g. providing a particular data rate $x\%$ of the time.

In best effort service, the network makes no promises. This service is typically used by *elastic traffic*. Elastic traffic consists of traffic where users do not necessarily have any minimum requirements, but would like to get as much data through to their respective destinations as quickly as possible. Individual user data flows react to congestion in the network and *adapt* their transmission rate with the aim of minimizing congestion. Email service is often provided on a best effort basis.

3 Differences between Wireline and Wireless Networks

Both wireless and wireline networks need to support the different network data service types described in the previous section. The focus of this paper is on supporting such services over wireless networks, and the gains that can be accrued by cross-layer techniques which do away with the firm boundary that currently exists between the PHY and MAC layers, and the higher layers of the Network protocol stack. Before studying mechanisms by which data services can be supported over wireless networks, we briefly describe the characteristics of wireless systems. Based on limited frequency allocations and channel considerations, wireless

networks have peculiarities which distinguish them from conventional wireline networks.

- (i) The wireless channel varies over time and space and has short-term (or small-scale) memory due to multipath. These variations are caused either due to motion of the wireless device, or due to changes in the surrounding physical environment, and lead to detector errors. This causes *bursts of errors* to occur during which packets cannot be successfully transmitted on the link. **Fast channel variations** due to fading are such that states of different channels can asynchronously switch from “good” to “bad” within a few milliseconds and vice-versa. Further, very strong forward error correction codes (i.e. very low rates) cannot be used to eliminate errors because this technique leads to reduced spectral efficiency.
- (ii) In addition to small-scale channel variations, there is also spatio-temporal variations on a much greater time-scale [2]. **Large-scale channel variation** means that the average channel state condition depends on user locations and interference levels. Thus, due to small-scale and large-scale changes in the channel, some users may inherently demand more channel access time than others based on their location or mobile velocity, even if their data rate requirement is the same as or even less than other users.

In Section 4, we discuss the implications of these effects for supporting data services over wireless networks.

4 Cross-layer Design for Supporting Data Services over Wireless Networks

In this section, we explore the gains that can be achieved by means of a cross-layer approach, where physical layer information is passed to the higher layers. We first consider supporting TCP traffic over wireless links, and then consider supporting data services in a *multi-user* wireless network. We finally discuss deployment tools and practical issues facing the build-out of access points, based on network throughput requirements.

4.1 TCP over Wireless Links

The prevalent protocol for data transport over the Internet is Transmission Control Protocol (TCP). TCP is a connection-oriented, end-to-end data transfer protocol. It has two objectives, namely (i) reliable end-to-end transmission of data, achieved by error or loss detection and re-transmission, and (ii) congestion control over the Internet. Routers in the network indicate congestion by dropping packets, which in turn causes the source to adaptively decrease its sending rate. Future deployment of TCP is expected to include the Explicit Congestion Notification (ECN) mechanism [3] which is used to notify the receiver whenever congestion

occurs in the network. This mechanism works in the following manner: included in a TCP packet's header is the *ECN bit* which is set to zero by the source. If the router detects congestion, it will *set the ECN bit to one*, and the packet is said to be marked. The marked packet eventually reaches the destination, which in turn informs the source about the value of the mark (i.e., the ECN bit value). The source adapts its transmission rate depending on the value of the mark.

The current deployment of the TCP protocol interprets all losses as being congestion related. Whenever losses occur over a wireless channel, the TCP source reacts to this as though it was due to congestion and thus decreases the packet transmission rate, causing loss in network throughput. A solution that has been proposed to mitigate this problem is to “smooth” the channel by suitable coding and link layer ARQ at a faster time-scale than that of the TCP control loop [4] (additional references available in [5]). so that the *wireless link ideally is perceived as a constant channel, but with lower capacity*. However, in practice, there is still the problem that the TCP sender may not be fully shielded from wireless link losses. This can lead to the TCP congestion control mechanism reacting to packet losses thus resulting in redundant retransmissions and loss of throughput.

However, once ECN-enabled TCP is deployed, where the ECN bit can be used to mark packets to indicate congestion, there is a means of differentiating between *congestion related loss* and *wireless channel related loss*. Thus, the channel need not be smoothed because the ECN mechanism provides a means of explicitly indicating congestion. In [6], it has been analytically shown that in a single user environment, if packets are marked based solely on congestion information, there is no significant degradation of TCP performance due to the time varying nature of the wireless channel. A similar idea based on explicit loss notification has been discussed in [4], where simulations indicate improved throughput performance. Such an approach is an example where a cross-layer view of physical layer information (channel conditions) is used at the network layer to significantly improve network-layer throughput performance.

4.2 Data over Wireless in a Multi-user Context

We now discuss various issues related to supporting data services over cellular wireless networks. We illustrate how channel-state dependent techniques (cross-layer methods) can lead to improved network throughput, and survey some of the current related research.

Consider a conventional cellular system with a fixed base-station and a number of mobile users. Data flows (packets) arrive from the wired Internet to the cell base station and are destined to the mobile users, with the packets for each user being queued temporarily at the base station (a separate queue is maintained for each user). The objective of the base station is to schedule these packets to various mobiles in a timely manner. In

this section, we restrict ourselves to just considering the forward-link problem, where the direction of the data flow is from the base-station to the mobile user.

The earliest method used by the cellular industry to send data to a mobile user on a cellular network consisted of sending data over a circuit-switched voice channel by means of a modem (usually over a FDMA/TDMA network). As data traffic is typically bursty, this method had the disadvantage that when a user is not transmitting data, the bandwidth on the voice channel is wasted by the dedicated voice channel.

A more efficient way to use cellular spectrum is to dedicate some bandwidth for data transport and allocate this bandwidth in a dynamic manner among various data users. For instance, time could be divided into fixed size time-slots and users could be allocated time-slots in a dynamic manner. The simplest scheduling algorithm is the round robin mechanism, where users are periodically allocated slots irrespective of whether there is data to be sent to a particular mobile user. Note that this simple TDM scheme leads to wasted bandwidth just as in the previous case. Further, it suffers from a more subtle problem, namely, this approach is independent of the channel state (a channel state independent scheduling mechanism).

The 3G wireless systems exploit a built-in mechanism where the forward-link channel state information for each user is fed back to the base-station. Examples of such systems include High Data Rate (HDR) versions of CDMA2000 [7] and enhanced general packet radio service (EGPRS) in the EDGE extension of GSM [2]. These proposals are essentially time-division overlays over the underlying physical-layer protocol to enable data traffic, and have the flexibility of allowing smart scheduling mechanisms to improve efficiency based on instantaneous user demand and channel states for each user.

In CDMA2000 HDR, the channel state is periodically measured by the user and fed back to the base-station once every 1.667 msec. As can be shown easily by computing the Doppler frequency of a mobile user at 2GHz, this slot size is short enough so that each user's channel quality stays approximately constant within one time-slot. In each time-slot, one user is scheduled for transmission. Each user constantly reports to the base station its "instantaneous" channel capacity, i.e., the rate at which data can be transmitted if this user is scheduled for transmission. In an HDR system, a "good" scheduling algorithm should take advantage of channel variations by giving some form of priority to users with instantaneously better channels.

4.2.1 Multiuser Diversity Gain: A Simple Illustration

We illustrate the significant gains due to channel state dependent scheduling algorithms by means of a simple example. Consider a wireless system consisting of three users as in Figure 1. For this example, we consider access time to be slotted, and the channels to be constant over a time-slot. We assume that the channels are

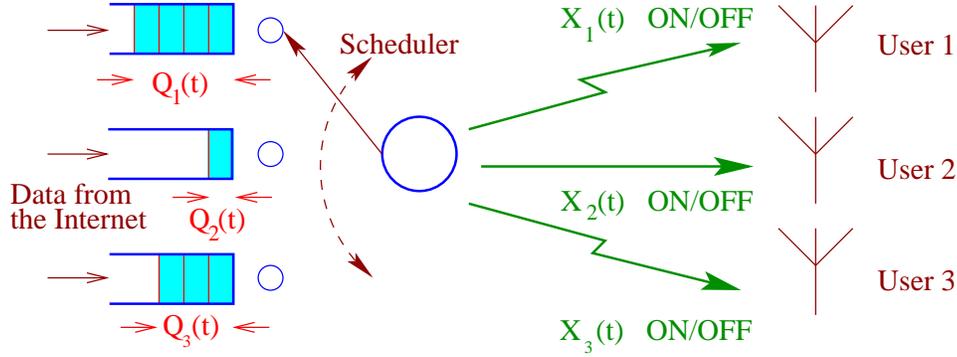


Figure 1: Packet data flows arrive from the Internet to the base-station. The data is destined to mobile users, and are temporarily buffered at the base-station. At each time t , the scheduler transmits to one of the mobile users based on the channel state $(X_1(t) X_2(t) X_3(t))$, with each component taking one of two values: ON or OFF.

either ON or OFF, equally likely, and the channels being independent of each user. Thus, in this system, there are eight possible (instantaneous) channel states for the three independent users, ranging from (ON,ON,ON) to (OFF,OFF,OFF).

When a user's channel is ON, one packet can be transmitted successfully to the mobile user during the time-slot. The system is assumed to be a TDM system; thus, the base-station can transmit to only one user on each time-slot. The associated scheduling problem is to decide which user is allowed access to the channel during each particular time-slot.

A naive scheduling rule would be to employ the well-known round-robin mechanism. In such a scheme, the users periodically are given access to the channel, with each user getting $1/3$ of the slots over time. As the channel of each user is equally likely to be ON or OFF in each time-slot, it follows that over time, on average, each user will get a data rate of $1/6$ packets/slot.

On the other hand, suppose that the base-station uses knowledge of the instantaneous channel state. Then, a simple policy would be to schedule and transmit to a user whose channel is in the ON state. If more than one user's channel is in the ON state, then, the scheduler could pick a user randomly (equally likely) among those users whose channel are ON, and send data to the selected user. This simple example assumes that all the users have identical traffic demands. For the case where some users have greater needs than others, such users would be assigned channel access with greater likelihood.

We observe that this policy ensures that *no data is sent by the base-station if and only if all users' channels are OFF* (which occurs on average, only $1/8$ of the time). Thus, the *total data rate* achieved by this state dependent rule is $1 - 1/8 = 7/8$ packets/slot. As this rule is symmetric across users, it follows that on

average, the *data rate per user* is $7/24$ packets/slot, which is almost twice the throughput as the round-robin scheme which provided $1/6$ packets/slot. In addition, the base station does not radiate power during bad channel conditions, thereby decreasing interference levels in the wireless network.

This gain achieved due to channel-state dependent scheduling is called the *multi-user diversity* gain [8, 9]. This example illustrates the significance of multi-user diversity gain in network scheduling. However, a central question is how to design on-line algorithms which achieve this gain, while also supporting diverse quality of service requirements for various users for realistic channel scenarios. While this problem is not yet completely answered, there has been extensive research on various aspects of this problem. Some recent references include [10, 9, 11, 12, 13, 14, 15, 16]. Additional references are available in a longer version of this paper [5].

4.2.2 Multiuser Gain with QoS

The previous sections illustrate gains that can be accrued based on cross-layer methods. In this section, we discuss some of the important issues related to multiuser gain and briefly survey related current literature. In particular, some important issues that will need to be addressed are:

- (i) How can multiple *real-time* data users be supported simultaneously with good quality of service (QoS) for all real-time users, namely, with packet delays not exceeding given thresholds with high probability?
- (ii) How can a *mixture* of real-time and non-real-time users be supported simultaneously with real-time users receiving their desired QoS and non-real-time users receiving the maximum possible throughput without compromising the QoS requirements of real-time users?
- (iii) How can bandwidth be fairly allocated among various users, especially when some users inherently demand more channel access time than others, even if their data rate requirement is the same, or less than others.

Though these questions have not been answered fully, there has been extensive research on various aspects of the above problems. Below, we discuss some of the available results. Additional discussion is available in a longer version of this paper [5].

In [10], optimal scheduling for a wireless system consisting of N queues and a single server is studied. The arrival processes to each of the queues are assumed to be independent and identically distributed (i.i.d.) Bernoulli processes. The channel perceived by each queue is also assumed to be an i.i.d. ON-OFF Bernoulli

process. The authors show that the policy which minimizes the total number of packets in the system, in a stochastic ordering sense, is the one which serves the user whose channel is ON and has the longest queue length.

The authors in [8] were among the first to study wireless scheduling in a multi-user context, and explicitly characterize the *multi-user diversity* gain. In [9], the author proposed a scheduling rule called the proportionally fair rule, which explicitly makes use of the channel state information and also provides fair allocation of bandwidth across users. In [12, 17], the authors show that an online algorithm (the Exponential rule) can be designed which can stabilize the user queues without explicit knowledge of the arrival rates or channel statistics, and also discuss QoS properties of such a rule.

In [18], the authors consider opportunistic beam-forming, and multi-user scheduling based on the proportional fair rule [9], and show that significant gains can be accrued due to multi-user diversity. In [14], researchers consider a utility based approach which address fairness with multi-user gains. There has been recent research on adaptively modifying the channel access time for individual users on an opportunistic basis depending on the “goodness” of user channels in wireless LANs [16].

Supporting reverse link data services imposes additional constraints due to the distributed nature of mobile users. Some recent work on supporting data services over the reverse link in a distributed manner is available in [13].

4.3 Deployment of Wireless Network Access Points

In today’s cellular/PCS networks, site deployment is conducted through computer-aided modeling of propagation characteristics, as well as the anticipated user densities and the corresponding RF interference levels [2]. Except for some large public WLAN systems, today’s WLAN deployments are made by IT professionals who are not accustomed to RF/interference issues. For efficient build-out for WLANs and 3G networks, site specific methods of determining locations of access points and in-situ throughput performance measurement of wireless modems offer promise [19, 20]. However, with so many different user classes and data rate requirements, future deployment and layout tools will need to “convert” from a wireless *system* design mindset with a *propagation* emphasis, to a wireless *network* design mindset with a user *throughput or QoS* emphasis. Early work shows that there are some first-order effects that relate well-known wireless propagation parameters, Signal to Noise Ratio (SNR) and Signal to Interference Ratio (SIR) to end-user network parameters (throughput, delay), independent of modem manufacturer or user classes.

Consider a multi user network based on 802.11b access points. The upshot of Figures 2 and 3 (see [19])

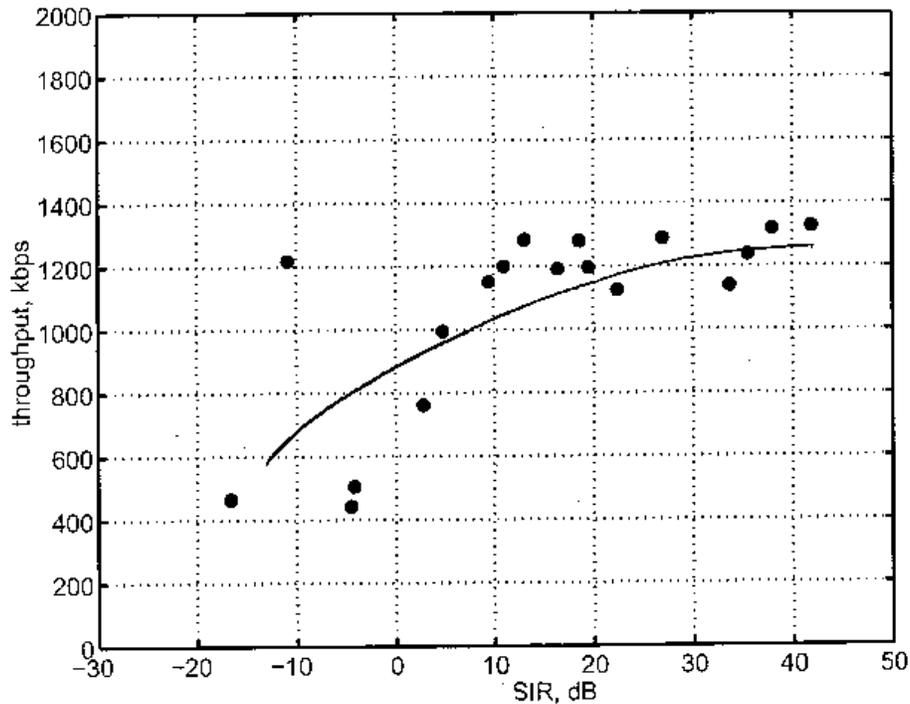


Figure 2: SIR vs. throughput for 3COM WLAN 802.11b modem in an indoor environment. Throughput was measured on a round-trip basis between two nodes, using LANFielder to generate random messages with unprotected UDP.

for more details) is that with knowledge of physical layer parameters, such as RSSI (Received Signal Strength Indication) and CIR (Carrier to Interference Ratio), it is possible to map these to network layer parameters, such as throughput.

This holds promise for network deployment CAD software tools that IT professionals can use without knowledge of specific propagation issues, while incorporating such vital data in their deployment strategy. SitePlanner and LANfielder are examples of commercial deployment and measurement tools that incorporate a network-centric approach to wireless design [20].

Another key issue facing large-scale deployment of future wireless networks in and around buildings is the need to track the location and on-going maintenance issues associated with the installed wireless equipment. Often, installers find it most desirable to hide the location of access points and antennas within buildings, primarily for aesthetic reasons, thereby making it difficult to find such equipment when providing upgrades or repairs. Unlike the early cellular telephone industry, where tall towers were easy to identify by sight, the wireless networks of the future will be hidden – “embedded” within buildings and campuses, much like building wiring, itself. There are products available today (such as those by [20]) that allow designers and

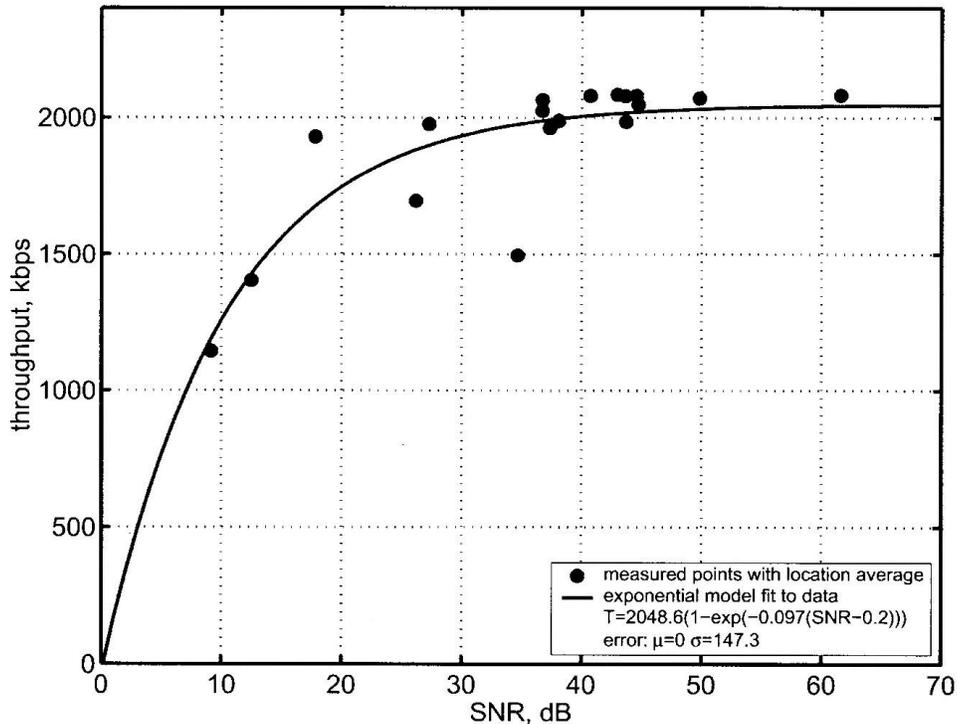


Figure 3: SNR vs. throughput for 3COM WLAN 802.11b modem in an indoor environment. Throughput was measured on a round-trip basis between two nodes, using LANFielder to generate random messages with unprotected UDP.

managers of wireless networks to handle all engineering, maintenance, and cost accounting details using simple Computer Aided Design software. This approach enables a wide range of staff, such as building owners, integrators, and vendors, to electronically share “as-built” design and measurement information that includes visual details of the location of network infrastructure within buildings. This same data may be easily ported to hand-held devices for use in the field by maintenance personnel.

5 Standardization Status for Cross-layer Design and Interoperability

As described earlier, wireless networks can be implemented as wireless LANs or cellular/WWANs. Coverage, mobility support, and throughput differ significantly between these network architectures, and there is no single system that currently stands out as a single best solution to mobile and nomadic users. The system performance of future networks will be enhanced by cross-layer design between PHY, MAC and higher layer protocols, as discussed in earlier sections. However, for a combination of services and applications demanded by users in a variety of environments, each wireless network can provide an important piece in the large mobile multimedia jigsaw puzzle. Standardization efforts are in progress to integrate various architectures.

The important co-design of physical layer, MAC and higher layers have been taken into account in some of the latest standardization issues. As discussed in Section 4.2, 3G standards such as CDMA2000 have been designed with cross-layer issues in mind. The Data Link Control (DLC) and physical layer of BRAN HiperLAN2 have been designed for high throughput, low latency and QoS support [21]. The 2 msec MAC frame with logical frame control channel, radio link control mechanisms, and scheduling provides a layered resource management. In 3GPP, the enhancement High Speed Downlink Packet Access (HSDPA) with hybrid ARQ (Automatic Repeat reQuest) and scheduling at the base-station have been introduced for reduced latency. The important topic of cross-layer design has recently been considered in the Study Group on Mobile Broadband Wireless Access Networks (MBWA) within IEEE, with the goal of improving throughput and reduced latency both in downlink and uplink.

5.1 Interoperability between WLANs and cellular WWANs

Interoperability between standards of distinct networks is crucial both for user adoption and operator management in a wide scale deployment of products and services over wireless networks. At the time of this writing, dozens of specialized applications and network protocols are being designed for dual use on IEEE 802.11a/b and cellular/PCS networks [1]. End user equipment for multi-mode terminals, such as PDAs, web phones and smart-cards, are emerging on the market.

Successful integration of existing and emerging wireless systems is a demanding task. Inherently, several standardization and regulatory bodies need to be involved, and the process is made difficult by the status of each specification process. A first WLAN/WWAN standardization initiative was taken by the ETSI BRAN project, which started the UMTS-HiperLAN2 interworking specification work. The scope has been generalized to include all WLAN standards with a loose coupling reference architecture based on the Internet Engineering Task Force (IETF) protocols. The standardization work was adopted by 3GPP, where a phased approach based on six scenarios ranging from common billing and access control toward seamless services has been defined for future releases. A new Wireless Interworking Group (WIG) was established in 2002 as a joint effort from the WLAN Community including ETSI BRAN, IEEE 802.11 and MMAC HiSWANa. The scope of WIG is to solve the interworking issues by defining an interface behind the WLAN AP toward the IP core network.

6 Conclusion and Future Research

This paper has attempted to provide an overview of the cross-layer paradigm shift that is beginning to take place as wireless communications evolves from a circuit-switched infrastructure to a packet-based infrastructure. The advantages of cross-layer networking appear to be real, and will be increasingly important as capacities offered through the wireless LAN interface approach the level of capacity that can be handled over the Internet backbone. Based on our observations, it appears that the following research areas are important for successful build-out of future broadband wireless networks. By no means, however, do we wish to suggest that these are necessarily the best problems to solve, nor is this listing meant to be exhaustive, but it does provide one viewpoint of the technical challenges and emerging research areas that will lead to new network implementations that support the amalgamation of user services and QoS requirements.

- (i) **Fundamental properties of multiuser diversity:** It has been shown (see Section 4.2) that multi-user diversity gains can substantially improve wireless network throughput. However, given different types of QoS constraints, how does the diversity gain grow with the number of users? The answer to this question is of great relevance from a practical perspective. The associated design problem would be to come up with good algorithms which make use of multiuser diversity gain. Finally, most of the studies so far are confined to a single-cell scenario. Extending these studies to the multi-cell case would be of great interest.
- (ii) **Modeling network performance:** Understanding the properties of mixed traffic and service types over wireless networks in practical propagation environments will be crucial for deploying and analyzing the performance characteristics of future scheduling approaches and protocols. Early progress has been made (see Section 4.3) but much more work remains as the industry converts from a circuit switched approach to a packet network. Creating models that translate actual channel characteristics, such as path loss and interference, into useful models at the network layer and above, while benchmarking such models with actual network field data will enable scheduling and traffic techniques to be developed accurately and expeditiously. Adaptive control of wireless networks, on a building-by-building basis, may also be needed to ensure proper scheduling and utilization of capacity as more access points and users are added in unlicensed spectrum.
- (iii) **Design for multi-carrier systems:** Most down-link scheduling algorithms today (for 3G wireless systems) are for a single-carrier (such as a shared time-division) system. Emerging WLANs and fourth generation wireless systems are expected to be based on a multi-carrier scheme (OFDM) with hundreds of carriers. In this scheduling and cross-layer networking context, an interesting area to explore is the

design of good algorithms with low complexity.

References

- [1] T. S. Rappaport, A. Annamalai, R. M. Buehrer, and W. H. Tranter, "Wireless communications: Past events and a future perspective," *IEEE Communications Magazine*, vol. 40, pp. 148–161, May 2002.
- [2] T. Rappaport, *Wireless Communications: Principles and Practice*. Upper Saddle River, NJ: Prentice Hall, 2002.
- [3] S. Floyd, "TCP and explicit congestion notification," *ACM Computer Communication Review*, vol. 24, pp. 10–23, October 1994.
- [4] H. Balakrishnan, V. Padmanabhan, S. Seshan, and R. H. Katz, "A comparison of mechanisms for improving TCP performance over wireless links," *IEEE/ACM Transactions on Networking*, December 1997.
- [5] S. Shakkottai, T. S. Rappaport, and P. C. Karlsson, "Cross-layer design for wireless networks." WNCG Technical Report TR-2003-04-00001, Wireless Networking and Communications Group, Department of Electrical and Computer Engineering, The University of Texas at Austin. Available for download at <http://www.wncg.org>.
- [6] S. Kunniyur and R. Srikant, "End-to-end congestion control: utility functions, random losses and ECN marks," in *Proceedings of IEEE Infocom*, vol. 3, (Tel Aviv, Israel), pp. 1323–1332, March 2000.
- [7] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushayana, and A. Viterbi, "CDMA/HDR: A bandwidth efficient high speed wireless data service for nomadic users," *IEEE Communications Magazine*, pp. 70–77, July 2000.
- [8] R. Knopp and P. Humblet, "Information capacity and power control in single-cell multiuser communications," in *Proceedings of the ICC*, (Seattle, WA), June 1995.
- [9] D. Tse, "Forward link multiuser diversity through proportional fair scheduling," August 1999. Presentation at Bell Labs.
- [10] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Transactions on Information Theory*, vol. 39, pp. 466–478, March 1993.
- [11] M. Andrews, K. Kumaran, K. Ramanan, A. L. Stolyar, R. Vijayakumar, and P. Whiting, "CDMA data QoS scheduling on the forward link with variable channel conditions," April 2000. Bell Laboratories Tech. Rep.
- [12] S. Shakkottai and A. Stolyar, "Scheduling for multiple flows sharing a time-varying channel: The exponential rule," *American Mathematical Society Translations, Series 2, A volume in memory of F. Karpelevich, Yu. M. Suhov, Editor*, vol. 207, 2002.
- [13] A. Eryilmaz, R. Srikant, and J. Perkins, "Stable scheduling policies for broadcast channels," in *Proceedings of the IEEE Symposium on Information Theory*, (Lausanne, Switzerland), July 2002. An earlier version appeared in *Modelling and Design of Wireless Networks, Proceedings of SPIE*, E. K. P. Chong (editor), vol. 4531, pages 70–78.
- [14] X. Liu, E. K. P. Chong, and N. B. Shroff, "A framework for opportunistic scheduling in wireless networks," *To appear in the Computer Networks Journal*, 2002.
- [15] M. J. Neely, E. Modiano, and C. E. Rohrs, "Power and server allocation in a multi-beam satellite with time varying channels," in *Proceedings of IEEE Infocom 2002*, (New York, NY), June 2002.
- [16] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic media access for multirate ad hoc networks," in *Proceedings of ACM MOBICOM 2002*, (Atlanta, GA), September 2002.
- [17] S. Shakkottai, R. Srikant, and A. L. Stolyar, "Pathwise optimality and state space collapse for the exponential rule," in *Proceedings of IEEE Symp. on Info. Theory (ISIT)*, (Lausanne, Switzerland), July 2002.
- [18] P. Viswanath, D. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Transactions on Information Theory*, vol. 48, pp. 1277–1294, June 2002.

- [19] B. Henty and T. S. Rappaport, "Throughput measurements and empirical prediction models for IEEE 802.11b wireless LAN (WLAN) installations," Tech. Rep. MPRG 01-08, Dept. of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA, August 2001. <http://scholar.lib.vt.edu/theses/available/etd-08142001-172523/>.
- [20] "Siteplanner, Wireless Valley Communications inc.." <http://www.wirelessvalley.com>. Also see US Patents 6499006, 6493679, 6442507, and 6317599.
- [21] A. Doufexi, S. Armour, P. Karlsson, M. Butler, A. Nix, D. Bull, and J. McGeehan, "A comparison of the HIPER-LAN/2 and IEEE 802.11a wireless LAN standards," *IEEE Communications Magazine*, vol. 40, pp. 172–180, May 2002.