An integrated approach based on cross-layer optimization

Designing a Mobile Broadband Wireless Access Network

Rajiv Laroia, Sathyadev Uppala, and Junyi Li

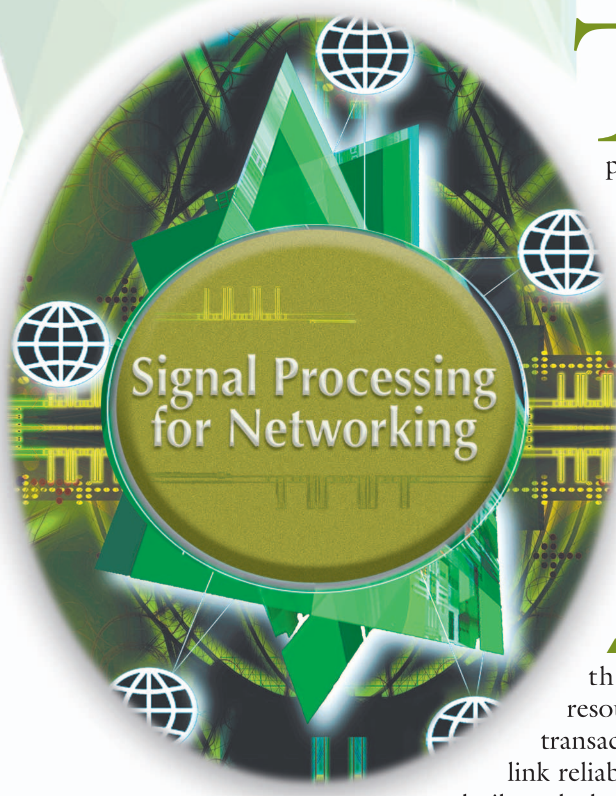

The fast growth of the Internet and the broadband technology has fundamentally changed the way to distribute and access information. Meanwhile, portable digital devices have become increasingly popular among businesses and consumers. A broadband mobile wireless access network is therefore widely envisioned to have tremendous market potential. When designing a broadband mobile wireless network, it is important to keep in mind its special requirements. The network should facilitate the seamless operation of the various Internet applications and protocols, so that a user perceives it to be of equivalent quality to a broadband wire line service. In particular, the key requirements of a desirable system include the following:

- *High bandwidth:* Such a network should provide a high bandwidth reliable link that can be shared by multiple users. A control policy that allows the resources to be efficiently switched between mobiles, on a packet-by-packet basis, is needed.
- *Low latency:* In addition to high bandwidth, latencies of the link should be low. When users get access to the resources, they should be able to sustain high burst rates and transact packets with minimal delay. Mechanisms for achieving link reliability, such as automatic repeat requests (ARQs), should be built such that delays are minimized.
- *Quality of service (QoS):* Wireless communication is typically resource constrained. QoS provides a means for effectively partitioning the limited resources. Users can be provided different levels of services, and service providers can use this to derive revenues.
- *Power savings:* The mobile nodes should be capable of achieving power savings. Mechanisms should exist for seamlessly switching to power save modes in between

1053-5888/04/$20.00©2004IEEE

data bursts while keeping latencies and control overhead in check.

In this article, we propose a new wireless system to enable a broadband mobile wireless access network. This system leverages the standard Internet protocol (IP) network elements to build the system and employs a new air interface technology based on orthogonal frequency division multiplexing (OFDM). Cross-layer issues have been the cornerstone of the system design. The design philosophy here is to jointly optimize the physical, media access control (MAC), and link layer, while leveraging the standard IP network architecture. The authors in [6] provide detailed discussions on cross layer designs for wireless systems. This article demonstrates how ideas of cross-layer design can be coherently integrated to build a complete system.

Wireless spectrum is scarce. Mobile wireless communications make achieving high spectral efficiency a challenge. Thus, the wireless link will likely be the bottleneck in any end-to-end mobile wireless broadband system. Therefore, the improvements of the wireless link performance, in terms of spectral efficiency, channel utilization and QoS capability, directly translate to the overall improvement of the end-to-end system.

System Overview

Circuit-switched technology has been the mainstay of wireless networks over the past 30 years and continues to pervade the network architecture of the recently deployed third-generation (3G) systems for legacy reasons. The system we propose (illustrated in Figure 1) is an embodiment of an end-to-end, all-IP solution. The radio access router (RAR) is connected to an IP network via a standard backhaul. It essentially routes the packets received from the network to the appropriate terminal device via the OFDM air interface. The wireless terminal (WT) can take many forms. For example, the WT can be a personal computer (PC) card that enables any generic IP host to connect to the network. Any voice or data packets transacted with the WT are treated with appropriate QoS and routed through the network. As a comparison, Figure 1 also shows a high-level architecture of a 3G system [2]. For the sake of illustration, we show a third-generation partnership project (3GPP) system. The network consists of many nonstandard IP elements that complicate the network architecture. In particular, a Node B is responsible for communication with the user terminal over the air. A radio network controller (RNC) controls several Node Bs. An RNC directs the voice traffic to the mobile switching center (MSC) and the data traffic to the serving general packet radio services (GPRS) support node (SGSN). The MSC is connected to public switched telephone network (PSTN) through the gateway mobile switching center (GMSC). The SGSN is connected to the Internet through the gateway GPRS support node (GGSN). In the OFDM system we present in this article, voice is treated as a special type of data and a single network handles both voice and data in an integrated fashion.

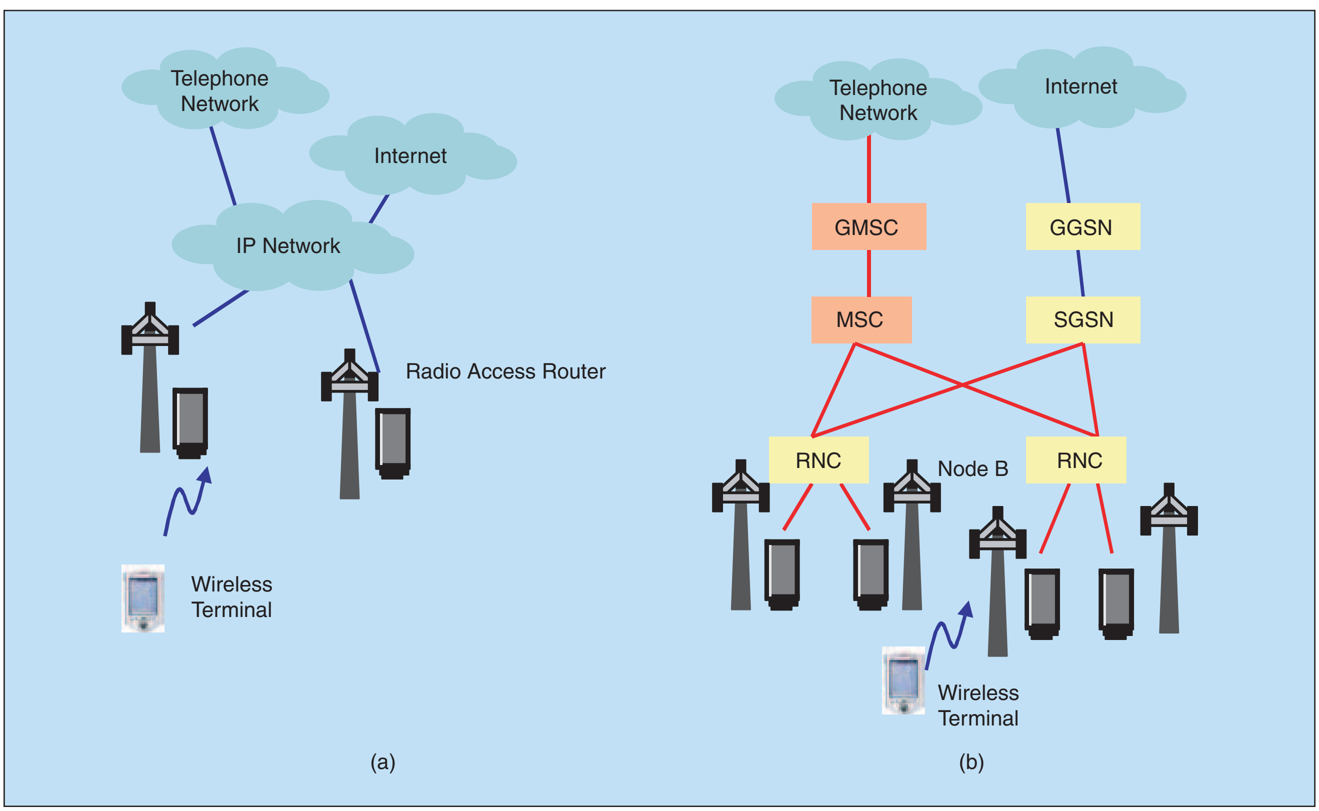

▲ *1. System architecture: (a) system architecture of an OFDM-based mobile broadband wireless access network and (b) 3G system architecture.*

Figure 2 compares the elements and protocols associated with the proposed OFDM system and a 3G system [3]. Here we show 3GPP2 architecture. In the proposed OFDM system, IP traffic is transported natively over the air interface between a WT and the RAR. Consequently, the RAR has the centralized knowledge of both IP packets and the wireless channel environment and can thus support IP QoS transparently and efficiently over the air link by facilitating cross-layer optimization. In the 3G systems, the different layer protocols terminate at different network elements between the WT and the core IP network. None of these network elements have sufficient real-time information such as instantaneous channel conditions and traffic queue status. Such real-time information is essential for making decisions, which lead to efficient utilization of resources and support IP QoS. Sharing information among distributed network elements cannot be effective due to the excessive communication delays. Moreover packets from the network get encapsulated in point-to-point protocol (PPP) by the network access server (NAS) and sent to the wireless terminal. Similarly, the terminal sends PPP packets that the NAS decapsulates and sends to the network. The addition of the PPP layer and isolation between layers limits the ability of the transceiver and the controller to be IP QoS aware. Consequently, cross-layer optimization is fundamentally hard in 3G systems.

The system we propose is based on an all-IP network. This allows us to reuse many of the protocols developed for the Internet. Mobility management is carried out using mobile IP protocols. Handoffs are done through layer 3 signaling, and this eliminates the need for layer 2 aware network elements in the radio access network. This approach also facilitates handoff to other IP networks, such as IEEE 802.11 local area networks.

Physical Layer Design

Basic Idea of OFDM

OFDM is a well-established technology adopted in the digital audio and video broadcasting systems in Europe and chosen by IEEE 802.11 standards group for wireless LANs (see [1]). The basic idea of OFDM is that the total bandwidth is divided into a number of orthogonal tones, over which multiple data symbols are transmitted in parallel. Consider the construction of an OFDM signal using N tones. Each of the N tone signals is a modulated sinusoid at a certain frequency. The frequencies of those tone signals are equally spaced.

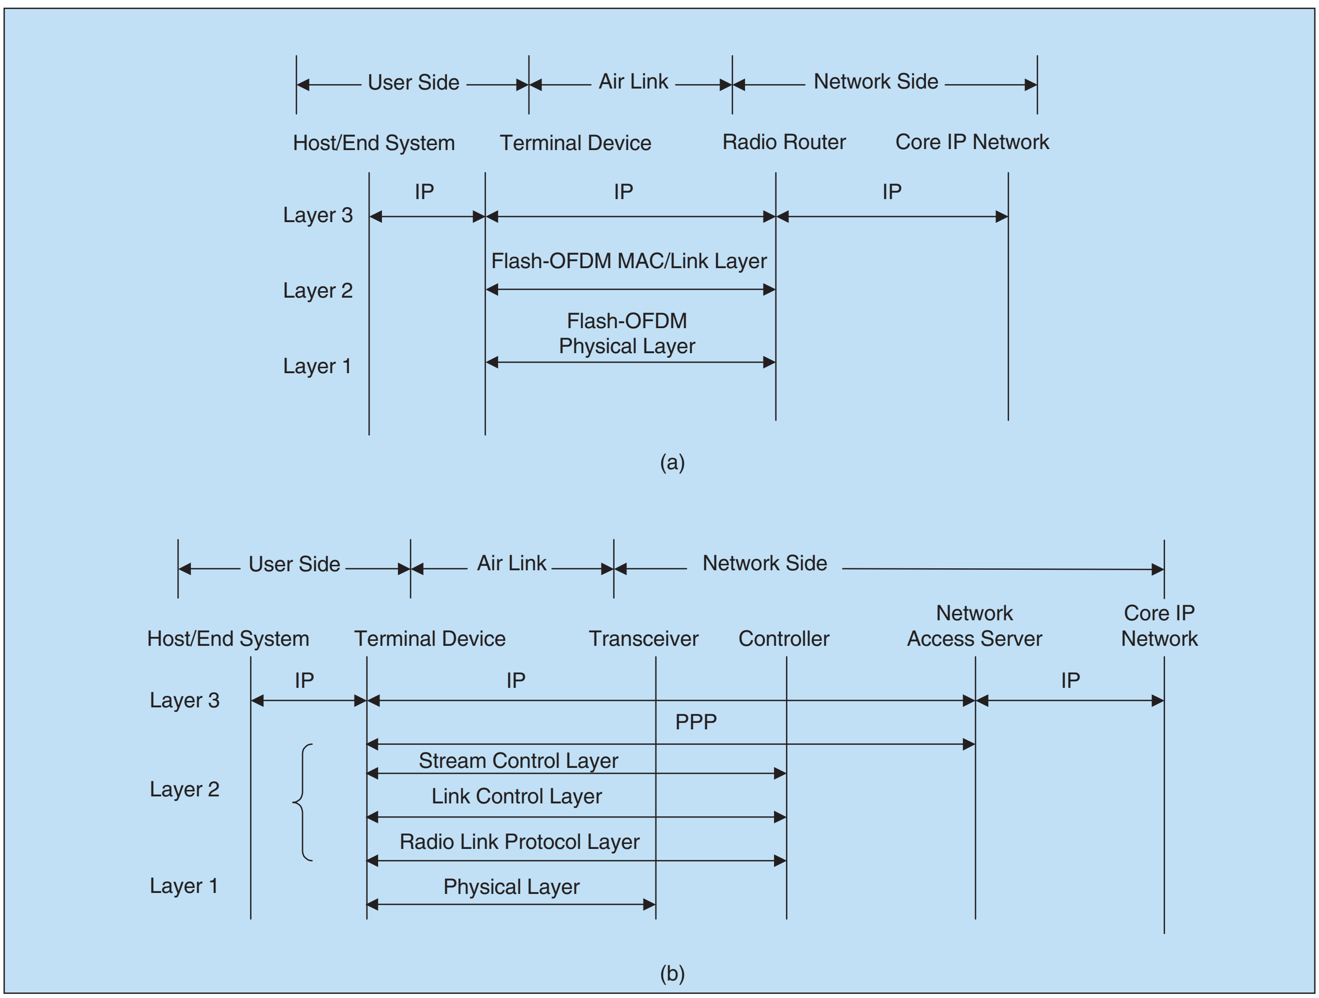

▲ *2. Protocol overview: (a) protocol stack of an OFDM-based mobile broadband wireless access network and (b) 3G system protocol stack.*

Therefore, all the tone signals are orthogonal to each other in a time duration T, which is the reciprocal of the tone-frequency spacing. An OFDM symbol is the sum of the tone signals for a time duration T, prepended with a cyclic prefix, which is a cyclic extension of the tone signals. The introduction of the cyclic prefix ensures that in a multipath channel, as long as the delay spread does not exceed that of the cyclic prefix, the multipath replicas of the OFDM symbol at the receiver always have an integer number of sinusoid cycles within a time duration T, thereby maintaining the orthogonality at the receiver.

Tone-Hopping OFDMA

OFDM can also been used as a multiple access scheme, referred to as OFDMA, where different tones may be assigned to multiple WTs. The OFDMA idea can be further combined with tone hopping to become a new form of spread spectrum technology, namely tone hopping TH-OFDMA (see [5], [7], and [8]). Tone hopping helps to improve frequency diversity and, more importantly, to average out-of-cell interference in a cellular environment. The air link resource in TH-OFDMA can be viewed in a time-frequency plane, as illustrated in Figure 5, where a small box represents a particular tone in a given OFDM symbol. Thus, each column represents all the tones in a given OFDM symbol, and each row represents a given tone over successive OFDM symbols. Tone hopping is specified with tone hopping sequences, each of which indicates the corresponding tone frequency of the hopping sequence at a given OFDM symbol. Figure 3 also illustrates a tone-hopping sequence.

In the proposed system, the total bandwidth is divided into a number of orthogonal tones and reused in every cell (or every sector in a sectorized cell). Adjacent cells employ distinct patterns of data tone hopping to construct various control and traffic channels. Data tone-hopping sequences are designed such that any two tone-hopping sequences in the same cell never collide (no intracell interference), and the maximum collision between any two tone-hopping sequences in different cells is minimized (averaged intercell interference).

Symbol-Based Hopping in Downlink and Dwell-Based Hopping in Uplink

In the downlink, a pilot signal is always transmitted to facilitate many operations at the WTs, such as system determination, synchronization, and channel estimation (see [9]). In the proposed system, the pilot consists of known symbols being transmitted with a few pilot-hopping sequences. In addition to satisfying the above requirement of data-hopping sequences, the pilot sequences are designed to uniformly cover the time-frequency plane. In the downlink, a pilot- or data-hopping sequence occupies different tones from one OFDM-symbol to another, as illustrated in Figure 4.

Uplink design possesses a different constraint. Sharing pilots among multiple mobiles is not feasible, and therefore coherent modulation would be hard with the same symbol-by-symbol tone hopping as in the downlink. Instead, uplink tone hopping occurs once every few symbols in this system. The interval during which tone hopping does not occur is called a dwell. Figure 5 shows the dwell-based uplink tone hopping. During a dwell, one symbol is used to transmit a known symbol, which can be used as a training symbol for channel estimation or as a reference symbol for differential modulation.

Channel Segments

In this system, the transmission unit of a traffic or control channel is called a segment, which consists of one or a few tone-hopping sequences over one or a few OFDM-symbols. Figure 8 illustrates traffic-channel

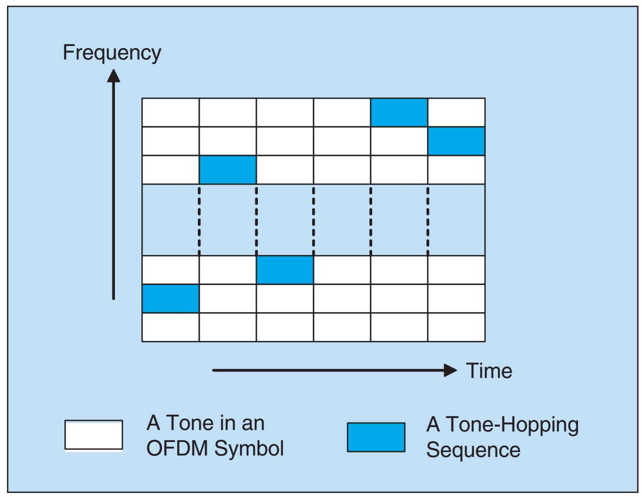

▲ *3. Air-link resource in tone hopping OFDMA in the time-frequency plane and symbol-by-symbol tone hopping sequence in the downlink.*

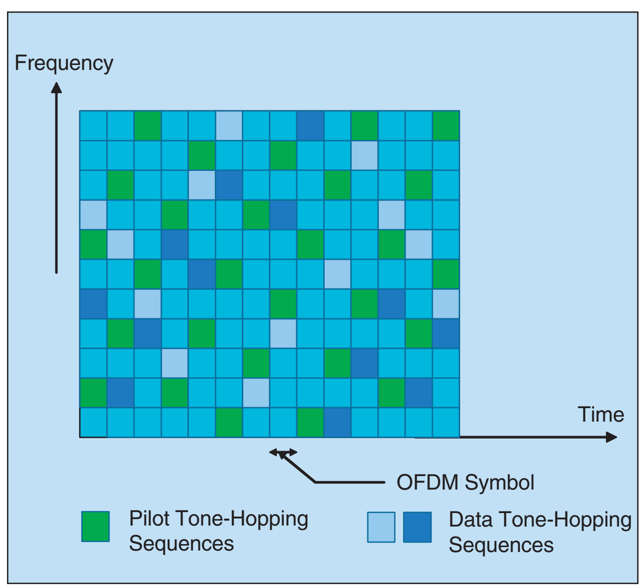

▲ *4. Symbol-by-symbol hopping in the downlink.*

segments and control-channel segments. In the figure, to simplify the illustration, the effect of tone hopping is omitted and each segment is thus represented as a rectangle in the time-frequency plane. As shown in Figure 6, the segments can be of different sizes to suit the specific requirements of the traffic or control channels. The channel coding and modulation are carried out on a per-segment basis. Apparently, the coding and modulation schemes can vary with respect to different segments and indeed have been individually optimized for each of the channels. For example, low-density parity check (LDPC) codes [4] can be used in a traffic-channel segment, which is relatively large, but cannot be used in a small control-channel segment.

Power Control and Timing Control

Power control and power allocation are used to reduce interference, improve link reliability, and save power. A WT periodically reports the measurement of downlink channel quality so that the RAR can determine the required transmission power at any coding and modulation rate when the WT is scheduled. The RAR also controls the uplink transmission power by measuring the quality of the uplink channel in which the channel measurement reports are sent and periodically sends power control instructions to adjust its transmission power.

To maintain the orthogonality in the uplink, the OFDM symbol boundaries of the signals from all WTs have to be aligned with the RAR receiver window, which is fixed in time. Closed-loop timing control is used for timing synchronization. Specifically, each WT periodically sends a timing measurement signal. The RAR measures the arrival time of the signal and sends back a command to adjust the WT's transmitter symbol timing. As closed-loop timing control is intended to compensate for the variation of the propagation delay due to human mobility, which is at much slower rate than the speed of the signal propagation, the required rate of closed-loop timing control is very slow. For example, a rate of once every few seconds is sufficient for the mobile velocity of 60 miles per hour. Therefore, the overhead is very little.

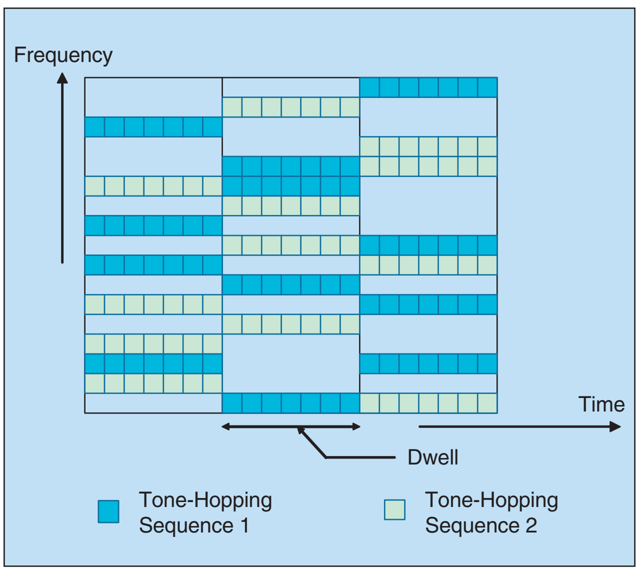

▲ 5. Dwell-by-dwell hopping in the uplink.

Salient Features

Improved Spectral Efficiency

TH-OFDMA has important advantages over CDMA. Specifically, in a CDMA system, different WTs are separated with different direct sequences. For example, Walsh codes are often used as a set of orthogonal basis functions to ensure that the signal of one WT does not interfere with that of another WT. However, the wireless multipath channel, as well as the pulse-shaping filters employed at both transmitter and receiver, destroys the orthogonality at the receiver. TH-OFDMA instead uses sinusoids, which have a fundamental advantage over any other orthogonal basis functions in that they are the *only* eigen functions of any time-invariant linear system. The wireless channel and the pulse-shaping filters can be treated as a time-invariant linear system for a short time period of one OFDM symbol. Therefore, as long as the cyclic prefix covers the total delay spread, the received signal is still a sum of sinusoids at the same frequencies as those of the transmitted signal, thereby preserving the orthogonality at the receiver. Hence, a TH-OFDMA-based system can achieve higher spectral efficiency than a CDMA-based system.

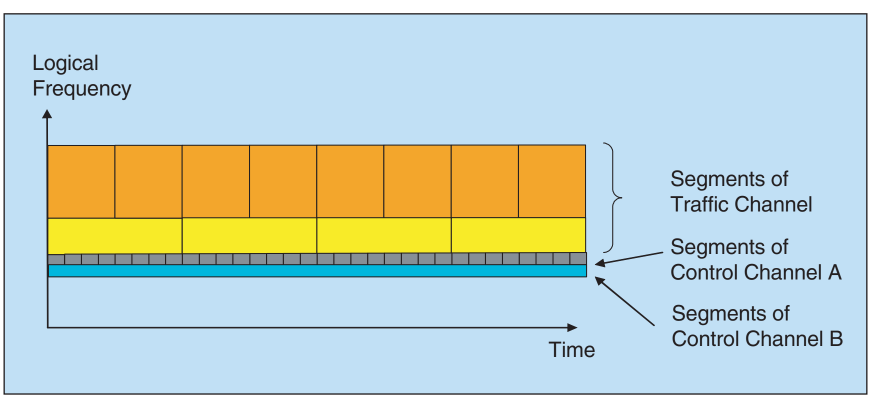

▲ 6. Illustration of traffic and control channel segments. Each rectangle represents a segment.

Fine Resource Granularity

We described the physical layer of the new OFDM-based air interface in the previous section. The MAC layer takes advantage of the fine resource granularity that can be achieved with the physical layer. As noted before, a segment is a unit of air link resource. The size of a segment can be made arbitrarily small or large depending on the number of information bits to be transported with the segment. For example, the MAC layer can use segments of very small size to transport a small number of control information bits. In fact, even a single information bit can be transported with little overhead. The segment need not contain any extensive training sequences because of the nature of the OFDM waveforms and the specific choice of symbol times and spacing between the tones. On the downlink, the demodulation information is obtained from the pilot. On the uplink, the segments have the necessary reference symbols to provide the phase reference.

This ability to send segments of arbitrary size is very useful to improve the MAC layer efficiency. Resources utilization can be tuned to meet the specific needs of the task at hand. For example, an acknowledgement is essentially a single bit so we can dedicate a very small acknowledgment segment, thereby providing a quick feedback of whether a traffic-channel segment is successfully received. In the MAC layer, a given segment can be dedicated for use by a specific WT. This allows the WT to convey control information to the RAR in a contention-free manner, thereby guaranteeing timely delivery of control information. Moreover, a given segment can be dedicated for use with predefined functionality so that there is no need of sending overheads, such as message headers, in the segment, thereby improving the efficiency and reducing the latency.

Our design philosophy is to jointly optimize the physical, media access control, and link layer, while leveraging the standard IP network architecture.

MAC Layer Design and Interactions with the Physical Layer

Above we have outlined the features of the proposed OFDM physical layer and motivated how they play a role in enhancing the spectral efficiency of the system. We now describe how media access control is done and how link reliability is achieved in the system. Considering the design philosophy of this system, treatment of the media access control layer in a vacuum is meaningless, and therefore there will be frequent references to other layers of the system.

Dynamically Broadcast Assignments and Link Reliability

This system is a packet-switched network. When WTs enter the system, they do not establish a circuit, i.e., there is no notion of certain traffic channels on downlink or uplink belonging to a specific WT. The traffic channel is a series of traffic segments and the base station (BS) controls their usage through the assignment channel. All WTs that are eligible to receive traffic segments decode the assignment channel. The assignments indicate a user identifier and the type of forward error correction and modulation for the traffic segment. WTs are preassigned a user identifier. There is a slaving between an assignment segment and a traffic segment.

Figure 7 shows the downlink traffic flow. The WT decodes the assignment segment and on seeing an identifier that matches the one allocated to it, proceeds to decode the corresponding traffic segment using the forward error correction and modulation indicated. The traffic segment can be checked whether it has been successfully decoded by using, among other things, a cyclic redundancy check (CRC). The WT then indicates with a single bit, the success or failure to decode the segment as a positive or negative acknowledgement on the uplink. The BS, because of the slave structure, knows which traffic segment is being acknowledged. A similar slave structure exists between the assignments on the downlink, the traffic segments on the uplink, and the acknowledgements on the downlink to control the uplink traffic flow. The acknowledgments are tied to the traffic segments and, hence, are shared by all the WTs. This amortizes the control resource nicely among the WTs.

The slave structure as described above also facilitates sending a traffic segment to multiple WTs (multicast of traffic). The assignment indicates a multicast group identifier. When the WTs see their group identifier in the assignment, they receive the traffic segment. This allows for a native support of multicast and an efficient usage of the air link bandwidth.

Channel Quality and Data Reporting

In the last section, we have described the slave structure between the assignments, traffic, and acknowledgements. There are two things we shall amplify in this section. One is how the BS makes the assignments for the uplink traffic. The second is how the BS makes a determination of the rate and modulation for the up and downlink.

On the downlink, the BS has the packets for all the WTs and, hence, can make a scheduling decision. On

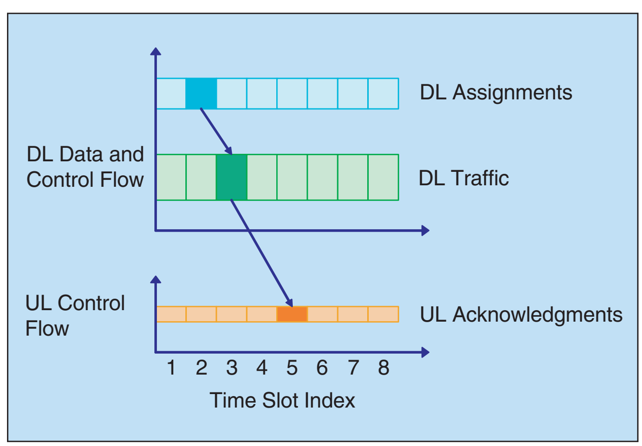

▲ 7. *Traffic control slave structure.*

the uplink, the WTs need to inform the BS of their data requirements. This allows the BS to know the status of the data queues of each of the WTs and facilitates a scheduling decision. In order to enable this exchange of information, each WT is provided with a dedicated control channel on which a WT sends its data requests periodically. This is illustrated in Figure 8. The BS factors all the data requests and the previous assignments sent in arriving at a scheduling decision. As a result of the fine resource granularity, the WT can send data requests on the dedicated control channel without interrupting other signaling such as channel quality reports. These frequent and rapid updates take place once every few milliseconds and minimize latencies in getting packets from the WT to the network.

Also facilitated on this dedicated channel are channel quality reports of various kinds. These reports allow the WT to indicate the downlink signal-to-noise ratio and interfere levels from other BSs. These reports are factored in when making a determination of which users to schedule and what modulation and forward error correction to use. Again, a periodic update of this information, using small, dedicated, control-channel segments that span a few milliseconds, allow the BS to receive updates and make scheduling decisions that are current.

State Transitions

A WT in this system can reside in various states, as shown in Figure 9. The states reflect various actions

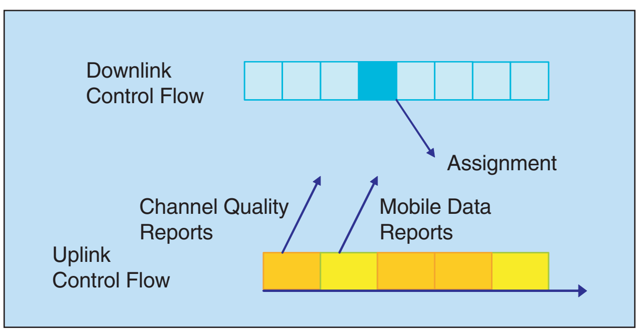

▲ *8. Channel quality and data reports from the mobile to the BS.*

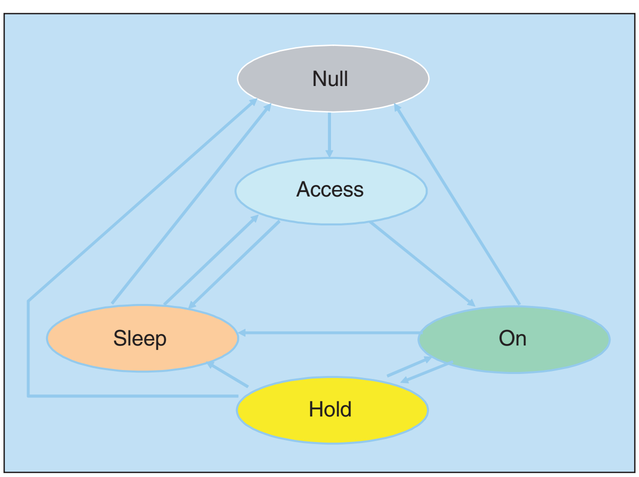

▲ *9. State transitions.*

taken at the physical, MAC, and network layers. There are several objectives for having these states. One is to facilitate various power-save modes for the WT and thus extend battery life. Another is to support a large population of WTs. On power up, a WT goes through the access state. After finishing access, it proceeds into the on state, which is the main data transaction state. The hold state augments the on state and provides some power saving. The sleep state provides the most power savings. The system can accommodate more WTs in the hold state than in the on state and in the sleep state than in the hold state. Transitions from hold to on take tens of milliseconds, whereas transitions from sleep to on take hundreds of milliseconds.

The physical and MAC layers are designed to allow fast access into the system. A WT, on powering on, identifies the strongest BS and synchronizes to it. The WT then achieves timing synchronization on the uplink by transmitting a special access signal over an access channel. The access channel is by design orthogonal to the rest of the uplink channels, and thus, the WT's access signal does not cause interference to other WTs connected to the BSs. This allows the WT to choose sufficiently high power for the access signals and do a quick access even in a faded channel. The BS, on detecting the access signal, measures the timing correction to be provided and feeds it back to the WT. The WT and BS then proceed to do further signaling, at the end of which the WT moves to the on state.

The on state is the main data transaction state. The WT in the on state is identified by a unique identifier and receives and transmits various control information. The WT is power and timing controlled. It sends data and channel quality reports on the uplink. All WTs in the on state decode the assignment channel. The scheduler (described later) at the BS determines based on several criteria, the usage of the uplink, and downlink traffic channels. The scheduling decision is signaled on the assignment channel.

The WT in the hold state is timing controlled. To support timing-control operations, the WT transmits a low periodicity signal on the uplink. The WT also listens to a state transition channel broadcast by the BS. The BS can move the WT back to the on state by paging it, or the WT can request a transition to the on state by sending a state transition request on a dedicated channel (i.e., contention free). WTs that need to transact data move to the on state. As very minimal transmissions are made in the hold state, the WT saves power.

The sleep state is the main power-saving mode. The WT periodically wakes up and listens to pages. It is neither power nor timing controlled. To transact data, the WT goes through access to the on state. When a WT crosses a cell boundary, i.e., detects a new BS, it goes through access and during the access process indicates its presence to the new BS and then goes back to sleep. Thus, the network is aware of the WT's location and can page it.

The MAC transitions allow one to capitalize on typical traffic patterns in a data network. For example, a Web-browsing scenario consists of clicking on a link, downloading a page, reading the page, and then clicking on another link. During the periods of inactivity, WTs can go into the hold state. This allows the system to cycle WTs between on and hold state and benefit from the statistical multiplexing. Long periods of inactivity would result in WTs going to the sleep state.

Scheduler

So far we have described several elements of the OFDM system. All of these come together under the control of a scheduler. In this section, rather than describe a particular scheduling algorithm in detail, we describe the architecture of the system and available information that enables diverse scheduling algorithms. One can find a survey of scheduling algorithms in [10]. The general function of the scheduler is to intelligently allocate air link resource to achieve high system performance. The scheduler, in a sense, becomes the focal point for achieving any cross-layer optimization, given that the system design allows for this. The scheduler uses information from the physical layer to the network layer to make decisions. This is a fundamental advantage over a system where the intelligence is distributed in many boxes in the network. We shall next take the segment scheduler as an example to elaborate on how the joint optimization comes about.

For every traffic segment, the segment scheduler determines which WT to assign and which coding/modulation scheme to use. The traffic channel assignments are made on a segment-by-segment basis so that the scheduler can rapidly respond to time-varying factors and generate optimal scheduling decisions. The scheduler is provided a rich set of information, such as the traffic packet queue status, the QoS profiles, and the channel conditions of all WTs. The scheduler is a central place in which the performance objectives of both the network layer and the physical layer are jointly optimized. We now elaborate on some of the performance objectives.

Spectral Efficiency

The goal is to maximize the system-wide throughput by careful selection of a WT and code rate assigned to each traffic segment. The following notions turn out to be very useful.

Multiuser Diversity

Fading is traditionally viewed as a source of unreliability that needs to be mitigated. Random channel fluctuations, however, independent across many WTs, can instead be exploited by opportunisticly scheduling a WT when its channel condition is favorable. Multiuser diversity improves the system throughput as the number of WTs in the system increases.

Joint Allocation of User, Power, Rate, and Bandwidth

As there are multiple simultaneous traffic segments, the scheduler can schedule multiple WTs at a time. For each WT, the scheduler assigns a traffic segment with proper number of tones, coding rate, modulation scheme, and power based on the channel characteristics of that WT. From an information theoretic perspective, this joint allocation can achieve higher spectral efficiency than a system where one WT is scheduled at a time and only rate allocation is possible [e.g., the downlink in the high data rate (HDR) system].

Link Adaptation

The fast acknowledgment and fast channel quality reporting mechanisms in the system enable fast link adaptation, thus improving spectral efficiency. For example, given the channel uncertainty, a margin in transmission power or coding/modulation is generally introduced to maintain a required level of the probability of successful reception. Those margins represent the tradeoff between spectral efficiency and link reliability. In the system, link reliability can be more readily achieved by fast link adaptation, thereby requiring much smaller margins.

IP QoS

Simply maximizing spectral efficiency, i.e., bits per second per hertz, may not be the ultimate goal of a wireless service provider. A service provider wants to maximize revenue. QoS awareness is crucial to achieve that goal. Instead of treating all WT traffic as undifferentiated byte streams, the idea of QoS is to deliver to end-WTs a rich set of services, each with distinct requirements of performance measures, such as priority, rate, delay, delay jitter, and dropping probability. Thus, all packets should not be treated equally in the scheduler. As the wireless link is one hop in an end-to-end IP connection, the scheduler essentially provides IP QoS over the air. Differentiated services and integrated services are two methods standardized by Internet Engineering Task Force (IETF) to implement QoS for IP networks. A scheduler needs to be aware of those standards such that interoperable QoS can be maintained consistently in an end-to-end fashion.

As many services are peer to peer in nature, delivering QoS in the uplink is as important as in the downlink. The feature of noncontention-based, segment-by-segment traffic channel assignment in the uplink enables that traffic from all WTs to flow in a controlled manner. In addition, each WT frequently reports its traffic queue status to the RAR, with a noncontention uplink channel, such that the scheduler decision can be based on instantaneous traffic requirements.

Fairness

From the WT's perspective, it is important that services be consistent and predictable. However, the actual throughput experienced by a given WT may vary over time as a function of the overall traffic loading in the cell and the wireless channel conditions. Maintaining fairness across the WTs by the scheduler is an effective way to reduce the throughput variation. In general, there is a tradeoff between the fairness and the overall spectral efficiency. For example, with a rate fair scheduler attempting to ensure identical throughput throughout the cell, WTs experience the minimum throughput variation. However, WTs with bad channel conditions are assigned more air link resource, thereby resulting in reduced overall throughput. Alternatively, in a resource fairness system, each WT gets an equal share of air link resource and the actual throughput varies for different WTs depending on their channel capacity. The benefit is that WTs with good channel condition are not sacrificed because of those with bad channel. As the channel condition varies quite significantly in the wireless environment, however, the throughput variation in the resource fairness system could be high. Fairness between the rate and resource fairness is also possible. Hence, the scheduler can support a variety of fairness criteria to allow the flexibility of trading off the fairness and the overall spectral efficiency.

Conclusion

We have presented a new OFDM-based air interface technology for a mobile broadband wireless system. Cross-layer optimization played a major role in the design. Many of the physical layer attributes are leveraged in the design of the MAC and link layers. The design choices made in the physical, MAC, and link layers are also driven by the goal of extending the Internet to the wireless space. A major physical layer benefit of this air interface comes from the orthogonality property that results in the elimination of in-cell interference. Tone hopping furthermore ensures that out-of-cell interference is averaged and a worst-case interferer does not limit the system performance. The physical layer features not only result in high capacity but also provide very fine granularity of allocating air link resource, which improves the MAC and link-layer efficiency. The MAC and link layer provide contention-free, fast control channels between the RAR and the WTs. These channels are used to ferry a variety of signaling such as assignments of traffic channel, acknowledgments, channel quality, and traffic request reports. This holistic approach allows for a scheduler that could not only achieve high spectral efficiency but also allow for a fine control over QoS attributes such as latency, reliability, and service differentiation.

Rajiv Laroia received the B.Tech. degree from the Indian Institute of Technology, Delhi, in 1985 and the M.S. and Ph.D. degrees from the University of Maryland, College Park, in 1989 and 1992, all in electrical engineering. In 1992 he joined Lucent Technologies Bell Laboratories. He founded Flarion Technologies in 2000 and is currently its CTO. His research interests are information theory and wireless communications. He has numerous publications and over 35 patents (granted and applied). He is a Fellow of the IEEE and received the 1992 Best Graduate Student of the Year Award at the University of Maryland. He was an associate editor for *IEEE Transactions on Information Theory*.

Sathyadev Uppala received the B.Tech. (hons.) degree in electronics and electrical communications engineering from Indian Institute of Technology, Kharagpur, in 1991. He received the M.S. and Ph.D. degrees in electrical engineering from the University of Washington, Seattle, in 1993 and 1998, respectively. From 1998 to 2000, he was a member of the technical staff at the wireless research center at Bell Laboratories (Lucent Technologies), Murray Hill, New Jersey. In 2000, he helped found Flarion Technologies. He is currently a director of technology at Flarion and conducts research and development of wireless communication systems. His interests include wireless systems, networking, communications, and signal processing.

Junyi Li received the B.S. and M.S. degrees from Shanghai Jiao Tong University, China, and the Ph.D. degree from Purdue University, West Lafayette, all in electrical engineering. He is an expert in the field of wireless communications systems. He helped found Flarion Technologies, where he is a key contributor of the Flash-OFDM technology and a director of Technologies. Prior to that, he was a member of the technical staff in the Digital Communications Research Department of Bell Laboratories, Lucent Technologies.

References

[1] R.D.J. van Nee, R. Prasad, and R. van Nee, *OFDM for Wireless Multimedia Communications.* Norwood, MA: Artech House, 2000.

[2] S. Uskela, "Key concepts for evolution towards beyond 3G networks," *IEEE Wireless Commun. Mag.*, vol. 10, pp. 43–48, Feb. 2003.

[3] P. Bender et al., "CDMA/HDR: A bandwidth-efficient high-speed wireless data service for nomadic users," *IEEE Commun. Mag.*, vol. 38, pp. 70–77, July 2000.

[4] T.J. Richardson and R.L. Urbanke, "The capacity of low-density parity-check codes under message-passing decoding," *IEEE Trans. Inform. Theory*, vol. 47, pp. 599–618, Feb. 2001.

[5] M. Gudmundson, "Generalized frequency hopping in mobile radio systems," in *Proc. IEEE Vehicular Technology Conf.*, 1993, pp. 788–791.

[6] A. Maharshi, L. Tong, and A. Swami, "Cross-layer designs of multichannel reservation MAC under Rayleigh fading," *IEEE Trans. Signal Processing*, vol. 51, pp. 2054–2067, Aug. 2003.

[7] J. Chuang and N. Sollenberger, "Beyond 3G: Wideband wireless data access based on OFDM and dynamic packet assignment," *IEEE Commun. Mag.*, vol. 38, pp. 78–87, July 2000.

[8] G. Leus, S. Zhou, and G.B. Giannakis, "Orthogonal multiple access over time-and frequency-selective fading," *IEEE Trans. Inform. Theory*, vol. 49, pp. 1942–1950, Aug. 2003.

[9] S. Zhou, G.B. Giannakis, and A. Swami, "Digital multi-carrier spread-spectrum versus direct-sequence spread-spectrum for resistance to jamming and multipath," *IEEE Trans. Commun.*, vol. 50, pp. 643–655, Apr. 2002.

[10] Y. Cao and V.O.K. Li, "Scheduling algorithms in broad-band wireless networks," *Proc. IEEE*, vol. 89, no. 1, pp. 76–87, Jan. 2001.