The Future of WiMAX: Multihop Relaying with IEEE 802.16j

Steven W. Peters and Robert W. Heath, Jr., University of Texas at Austin

ABSTRACT

Relaying and cooperation have re-emerged as important research topics in wireless communication over the past half-decade. Although multihop relaying for coverage extension in wireless networks is an old concept, it became practical only recently. Nowhere is this better illustrated than in the IEEE 802.16 working group, which has devoted a task group to incorporating relay capabilities in the foundation of mobile WiMAX-IEEE 802.16e-2005. Currently, this task group is in the process of finishing IEEE 802.16j, the Multihop Relay Specification for 802.16. This amendment will be fully compatible with 802.16e-2005 mobile and subscriber stations, but a BS specific to 802.16j will be required for relays to operate. This article presents an introduction to the upcoming IEEE 802.16j amendment and provides insight about the obstacles that practical system designers face when incorporating relaying into a wireless broadband network.

INTRODUCTION

Broadband access to the Internet has become a staple of the modern home. The wireless industry, however, does not have a substantial share of the last-mile broadband access market. To address this problem, IEEE 802.16 [1–3] was formed, initially aimed at creating a standard that could compete with cable access networks as a last-mile solution for broadband access. In the beginning, end users were expected to be immobile and to have a line-of-sight (LOS) to the base station (BS). The standard since has moved into the mobile non-line-of-sight (NLOS) domain, mainly with the 802.16e-2005 amendment [4] to the 802.16-2004 standard [5].

To streamline the process of product certification and deployment, industry leaders created the WiMAX Forum, a non-profit consortium tasked with creating profiles based on IEEE 802.16 and certifying products as being compliant with these profiles. Over 500 companies are members of the WiMAX Forum, a testament to the investment of industry in WiMAX. Because of their efforts, mobile WiMAX systems are expected to begin deploying in the United States in 2008.

Even with the advanced signal processing techniques employed in WiMAX (e.g., orthogonal frequency division multiple access [OFDMA] and multiple input multiple output [MIMO]), the projected data rates will require a signal-to-noise ratio (SNR) at the receiver that may be difficult to obtain at the cell edge. Competition with wireline broadband providers means WiMAX must be highly reliable, and competition with third generation (3G) cellular networks means WiMAX must fill coverage holes for maximum mobility. These challenges are conflicting at a basic level; all else being equal, increasing data rate reduces reliability, and increasing minimum reliability service reduces the coverage area.

The most intuitive and up to now, the most widely used strategy to address these challenges is to shrink the size of the cell, effectively increasing the number of BSs over a given area. Although this strategy likely will increase capacity because users are much closer to their serving BSs (until the increased interference outweighs the increased signal power), its benefit is limited because of the exceeding cost of BSs. For each BS, the provider must pay for antenna space and the wired backhaul to the network, not to mention the digital and radio frequency (RF) equipment itself.

Instead, an increasingly attractive strategy is to insert fixed relays into the cell whose sole purpose is to aid communication from BS to mobile station (MS) and vice versa. Such networks are called multihop cellular networks (MCNs) and are the focus of recent research [6]. Although there are many unsolved problems in relaying, IEEE 802.16 has formed a task group to extend the IEEE 802.16e-2005 standard to include multihop communication, indicating that the field has reached a significant level of maturity. This amendment is called IEEE 802.16j [7], or as we sometimes refer to it, 16j.

The history of 16j is quite recent; the task group was created in March 2006, and the first technical contributions were made in November of the same year. By August 2007, a draft was sent to ballot, requesting the approval of the 350 IEEE 802.16 members. The draft failed to gar-

This work was supported by the Semiconductor Research Corporation under contract 2007-HJ-1648.

Date	Action
Aug. 2008	Final letter ballot approval
Dec. 2008	Sponsor ballot approval
Feb. 2009	Submission to RevCom
Mar. 2009	SA Approval

Table 1. *Speculative timeline for IEEE 802.16j approval.*

ner the 75 percent approval rate required so an extensive effort was put forth to create a second draft, which passed letter ballot on January 14, 2008. To resolve as many comments from disapproving voters as possible, the task group subsequently completed Drafts 3 through 6, the last of which is, as of the date of final authorship, proceeding to sponsor ballot and then to the IEEE Standards Association (SA) Standards Board Review Committee (RevCom). This committee must make a recommendation for approval or disapproval to the Standards Board itself. An optimistic speculative timeline based on previous schedules [8, 9] is shown in Table 1. Note that because the draft is not finalized, it may change during the remainder of the standardization process. This article is current as of Draft 4.

After standardization, the WiMAX Forum may create profiles based on a subset of features from 16j to ensure interoperability between vendor equipment. This subset is what actually may be implemented. The process of creating a WiMAX profile based on IEEE 802.16j has begun, but the final outcome of this process is still unclear.

Because of its short history, literature on 802.16j is sparse, consisting mainly of detailed reports of contributions to the standard proposed by the respective authors (e.g., [10–12]). Independent researchers have begun to study the amendment (e.g., [13]), but to learn even the basics, they are faced with the tedious and daunting task of sifting through contributions, comments databases, and the document itself. Although the authors have not been involved actively in the creation of IEEE 802.16j and are not 802.16 working group members, we have been following its development closely from both an academic and industrial viewpoint and found a need for a clear tutorial for researchers and executives with an interest in 802.16j. This article introduces the IEEE 802.16j standard in a tutorial form not yet found in the literature, to give newcomers without inside access, an understanding of how and why the standard was built.

The following sections provide a technical overview of IEEE 802.16j while attempting to point out interesting new aspects of MCNs. The article begins by discussing the technological challenges the 16j task group faced and the scope of the problem they are attempting to solve. Next, the most important physical (PHY)- and medium access control (MAC)-layer aspects of 16j are reviewed with an emphasis on the variety of possible relaying structures, and the article concludes with a summary and a look to the future.

Goals and Challenges

The purpose of IEEE 802.16j is not to standardize a new cellular network that includes multihop capability, but instead to expand previous single-hop 802.16 standards to include multihop capability. To expedite the standards process while still maximizing the likelihood that 16j will be useful in practice, designers limited their scope to the point-to-multipoint (PMP) OFDMA PHY mode of 802.16e-2005. Because companies are already well into the development stages of WiMAX-compatible mobile devices, 16j must be compatible with 16e MSs. This severely limits the scope and capabilities of the physical layer, but increases the likelihood that a WiMAX profile based on 16j will be implemented. As we will see, however, the BS must be modified to allow relaying. Since WiMAX BSs have yet to be deployed in many areas, this drawback may not be too severe.

Not only must the new standard be fully compatible with 16e devices, it also must satisfy this requirement for several levels of relay functionality. In 16j, not all relays are created equal. One can imagine a relay placed by the service provider near newly developed areas to extend coverage, or a commercial product bought by a subscriber wishing to extend coverage into his home. To narrow their focus, the task group created the following relay usage scenarios for IEEE 802.16j [14].

Fixed Infrastructure — Fixed-infrastructure relays, like BSs, are to be deployed by the service provider in stationary areas to serve general traffic. They are intended to increase both throughput and coverage because they are likely to be placed above roof tops to allow an LOS with the BS, but this may not always be the case. This category also may include commercial relays purchased by a subscriber, which may leave and enter the network at any time.

In-Building Coverage — Even with the relatively small demands of voice service, current mobile phones often perform poorly inside buildings. Relays are expected to be placed both by the service provider and by the end user near the shell, or just inside, of the building to fill the "coverage hole" inside. This type of relay also can be deployed near tunnels or subways to provide coverage where there is otherwise none. These relays can be nomadic and likely will operate with NLOS channels. Intriguingly, they may operate on battery power and probably will have low complexity.

Temporary Coverage — Events where a large group of people are densely packed into a small area form a unique opportunity for relays. The multihop capability of 802.16j will enable some of the traffic generated by this dense population to be routed to BSs in adjacent cells. Near stadiums, this infrastructure can be placed by the service provider as a permanent solution. Temporary relays also can be deployed in emer-

The purpose of IEEE 802.16j is not to standardize a new cellular network that includes multihop capability, but instead to expand previous single-hop 802.16 standards to include multihop capability.

The PHY of an MS expects to communicate with only one device: the BS. In 16j, the MS may receive from the BS and a relay station in the same frame. This raises issues regarding channel estimation, synchronization, and frequency offset.

gencies where some BSs may have been damaged. For this reason, temporary coverage relays may be required to run on batteries and will range from small and simple to large and complex.

Coverage on a Mobile Vehicle — A mobile vehicle, such as a train or bus, presents unique challenges to communications engineers. Usually, there are several people located very closely together, and the vehicle is moving, sometimes very quickly, through cells. To provide reliable coverage to such users, a complex relay may be deployed on the vehicle and obviously, will be highly mobile.

From a PHY perspective, backward compatibility with 16e MSs does not pose a technical difficulty, as much as it constrains possibilities. In the downlink, advanced combining of multiple signals (e.g., from both the relay and BS) is not possible, although as detailed later, cooperation is still possible. Furthermore, the relay must be able to support all the modulation and coding schemes a 16e MS supports.

The PHY of an MS expects to communicate with only one device — the BS. In 16j, the MS may receive from the BS *and* a relay station in the same frame. This raises issues regarding channel estimation, synchronization, and frequency offset.

Furthermore, PHY researchers often disregard the MAC layer sitting on top of the PHY. In receive mode, the MAC considers what kind of data the PHY received and what should be done with it; in transmit mode, the MAC neatly packages information and gives it to the PHY for transmission. For this reason, pure amplify-and-forward relaying strategies are not practical for 16j because the MAC, which figures out exactly how the data should be relayed, must operate on bits. A slight deviation from this is discussed in the next section.

Much of the burden of 16e-compatibility is placed on the MAC. An entirely new set of messages specific to relaying must be created without overlapping with the extensive set of MAC messages in IEEE 802.16-2004 and IEEE 802.16e-2005. Further, the MAC now is responsible for ensuring a required quality of service (QoS) over multiple hops and allowing for handoffs of relays serving multiple MSs. Hybrid automatic repeat request (HARQ) also must be maintained over multiple hops, as does the advanced security offered by 802.16.

The task given to the 16j task group was appreciable. In the next section, we present a broad technical overview of its solutions to the above challenges, focusing on the most important aspects.

Technical Overview of IEEE 802.16j

Terminology and Architecture

For the remainder of this article, we employ terminology used in the IEEE 802.16j draft but attempt to keep jargon to a minimum. A 16j-compliant BS is a multihop relay-BS (MR-BS), a relay station is an RS, and the end-user devices are MSs.

The 16j standard does not place a limit on the number of hops from MR-BS to MS in most cases, and the path between the MR-BS and the MS must contain only RSs. The station that directly communicates with an MS is the *access station* for that MS; the access station may be an MR-BS or an RS. The station that an RS transmits to in the uplink is its superordinate station; a relay station that an RS or MR-BS transmits to in the downlink is its subordinate RS.

With this terminology in mind, we can begin a more thorough overview of IEEE 802.16j.

PHY Modifications

Several additions to a single-hop PHY must be made to incorporate multihop capability. In particular, consider a typical single-hop cellular network. Whenever the BS transmits, it is doing so in the downlink. Whenever the MS transmits, it is doing so in the uplink. A relay, however, must transmit and receive in both the uplink and the downlink. In IEEE 802.16e-2005, a frame consists of a downlink subframe followed by an uplink subframe.[1] Thus, to accommodate the relay, 16j designers split each subframe into an *access zone* and a *relay zone*. The word "access" in access zone betrays the fact that the MS is transmitting/receiving in the uplink/downlink access zone. Similarly the relay zone consists of communication between relay stations and their superordinate stations.

To keep 16j as general as possible, its designers have included functionality allowing a non-transparent[2] relay to transmit on a different carrier frequency than its superordinate station. This can be done with a single radio that switches carrier frequencies when it switches between uplink and downlink, or it could be done with multiple radios capable of simultaneous transmission [15]. Much of the standard, however, focuses on the case where a relay transmits on the same carrier frequency as its superordinate station, and this is the case we consider here. Further, the multi-frequency case was introduced only recently (in Draft 4) and is more prone to change during the remainder of the standardization procedure.

A dichotomy of MSs arises in MCNs — those that are communicating directly with the BS and those whose data is routed through one or more relays. The MSs communicating with relays can be grouped further into mobiles that can or cannot decode the control information from the BS. The relays serving those that can decode the control information of the BS are not required to transmit control information themselves. These mobiles are in range of the BS but can achieve higher throughput by using multiple hops, each with superior signals in both directions. Thus, the benefit of relaying in this scenario is an increase in capacity. Such relays are termed *transparent*, because the mobile is not aware the relay exists.

The relays serving MSs that cannot decode the control information from the MR-BS must transmit control information at the beginning of the frame because the performance of the 16e-compatible device depends on it. These relays are called *non-transparent* because the MS syn-

[1] This article focuses on time-division duplexing (TDD) networks, although the OFDMA PHY supports FDD as well.

[2] This term will be defined shortly.

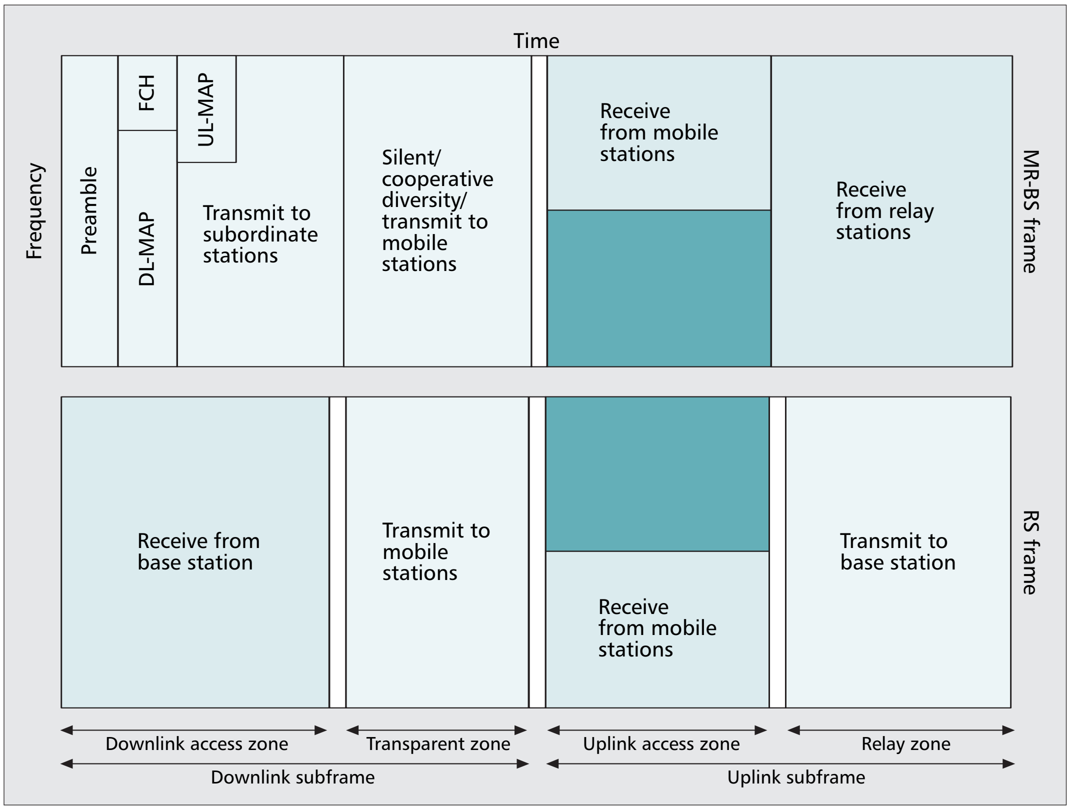

Figure 1. *An example transparent frame for both the MR-BS and the transparent RS, which does not transmit control information such as preamble or MAP. Transmissions by the MR-BS and RS are orthogonal in time or frequency. In the transparent zone, the MR-BS may cooperate with the transparent RS in sending information to the MS. A transparent relay may also be the subordinate of a non-transparent relay.*

After the downlink subframe, the uplink subframe begins. Like its downlink counterpart, the uplink subframe is split into two zones: the uplink access zone, where MSs transmit to their serving units, and the the uplink relay zone, where relays transmit to the BS.

chronizes and collects control information from it; however, the MS will have an implicit "understanding" that the non-transparent relay is actually a BS.

Transparent Relaying — Let us first consider the frame structure for a two-hop transparent relay system, an example of which is shown in Fig. 1. As in 16e, the frame begins with the downlink subframe. The BS has data to send to several MSs, all of which can decode its control data. Some of the MSs, however, may benefit if the data is first sent to a relay and then forwarded to the MS; thus, these MSs are served by relays. The way 16j designers attacked this scenario is to have all control information originate from the BS. The BS then proceeds in its normal downlink subframe, sending data to the MSs it serves directly and its subordinate relays. In this manner, the relay acts as an MS during the first downlink subframe. The downlink subframe, however, is split into two zones:

- The downlink access zone, where the BS transmits to the MSs it directly serves and its subordinate relays
- The transparent zone, where the relays transmit to the MSs they serve

In the transparent zone, the BS can either transmit to its subordinate MSs, remain silent, or transmit cooperatively with the relays, as described later in this section.

After the downlink subframe, the uplink subframe begins (with a small transition time to allow all participating RF front ends to switch from transmit to receive or vice versa). Like its downlink counterpart, the uplink subframe is split into two zones:

- The uplink access zone, where MSs transmit to their serving units
- The uplink relay zone, where relays transmit to the BS

As explained earlier, the relays generally operate in decode-and-forward mode, although this is not explicitly stated in the standard. This is because the MAC requires data contained in the header and subheaders to operate, and the PHY is generally unaware of the difference between these headers and payload data. The exception to this rule is in the optional *direct relay* zone, where the relay demodulates and deinterleaves its signal, then immediately interleaves and modulates it for transmission; without decoding, the relay may be able to receive and transmit the same data in the same frame. Such a mode of operation is permitted only for transparent two-hop relaying. Although direct relaying is an optional zone that was added to the standard relatively recently, it poses some interesting technical questions, specifically about the loss of soft information at the relay.

An immediately obvious drawback to transparent relaying is that part of the control information sent by the BS is the preamble, which may be used for channel estimation at the MS. Thus, the MS initially may train for the BS-to-MS channel, whereas its data is actually sent on the RS-to-MS channel. Although OFDM subchannel pilot symbols help correct the incorrect

The relays serving MSs that cannot decode the BS control information must transmit control information themselves. These are non-transparent relays. Because the MS cannot decode this information from the BS, it is out of range; thus, a non-transparent relay extends coverage.

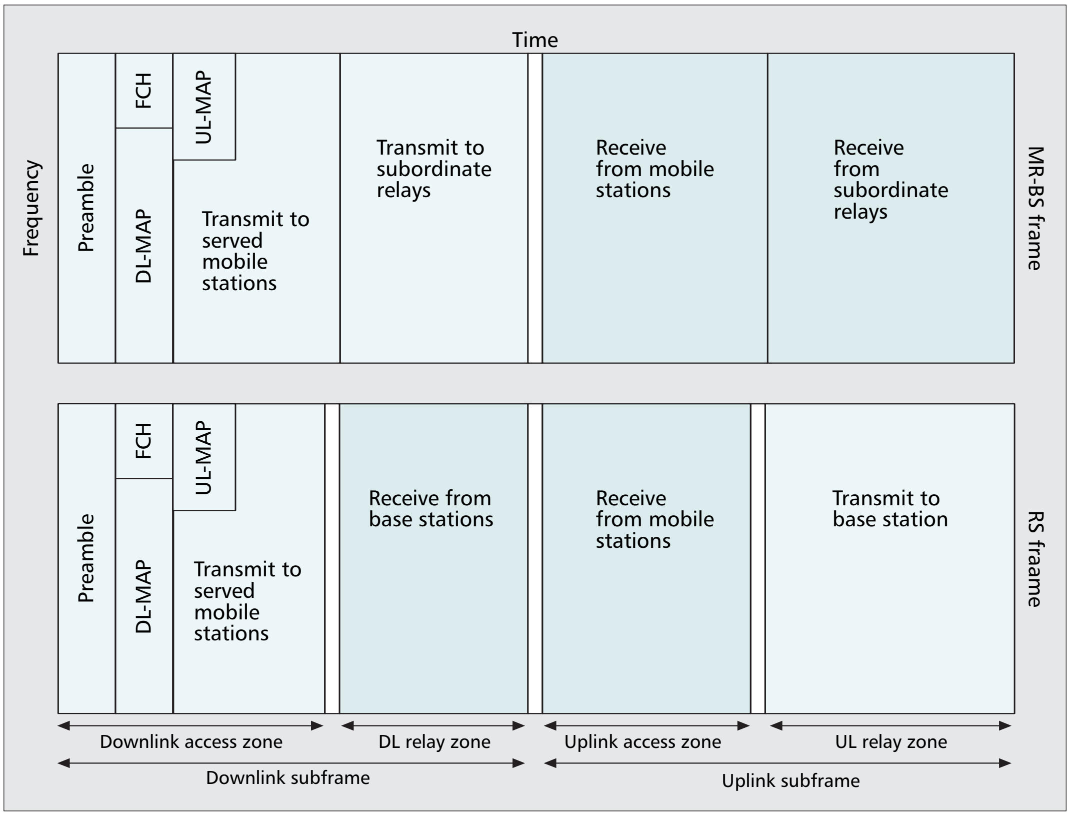

Figure 2. *An example non-transparent frame for both the MR-BS and the non-transparent RS, which transmits control information. This control information may interfere with that of the base stations and other non-transparent relays, so its main aim is coverage extension. The non-transparent relay may be capable of providing several MAC layer functions without direction from the MR-BS, and thus has potentially much higher complexity than a transparent RS.*

training, a loss in performance still can appear. The MS also expects the frequency offset from the BS; thus, the RS must alter its signal to mimic this offset.

Non-Transparent Relaying — The relays serving MSs that cannot decode the BS control information must transmit control information themselves. These are non-transparent relays. Because the MS cannot decode this information from the BS, it is out of range; thus, a non-transparent relay extends coverage.

An example frame structure for a non-transparent relay system is shown in Fig. 2. In this case, both the MR-BS and the relay transmit control data at the beginning of the frame. This way, the MS can synchronize with the relay, which is synchronized with the MR-BS. The downlink subframe starts off in the access zone so relays and BSs both are transmitting to their MSs. After the downlink access zone is the downlink relay zone, which consists of control information, data, and a midamble (called relay amble or R-amble) specifically for subordinate relays.

This raises the question of multiple hops. Suppose there are three hops between the MR-BS and the MS. In the scenario of Fig. 2, the access relay transmits to the MS in the downlink access zone and expects data from its superordinate station in the relay zone. The problem is that its superordinate station is a relay that also may be required to receive in the relay zone. There are two ways to approach this. One is to include multiple relay zones in a frame and have the relays alternate from transmitting to receiving in each one, as in Fig. 3. The other approach is to group frames together into what is called a multiframe and coordinate a repeating pattern that dictates which relays receive or transmit in each relay zone. An example given in the standard describes a two frame multiframe where relays on even hops transmit in the relay zones of even-numbered frames, and relays on odd hops transmit in the relay zones of odd-numbered frames. Note that this technique can apply to both the uplink and downlink.

The main drawback in the non-transparent case is that now the relay and BS are transmitting simultaneously in time and possibly, frequency. The immediate drawback is an increase in interference, particularly in the preamble and control channels. Obviously, power control and frequency reuse, which largely are left up to manufacturers, are crucial to non-transparent relaying. Further, non-transparent relays likely are more sophisticated (and thus, more expensive) than transparent ones because, as discussed in the next section, they may be allowed to make higher layer decisions on their own.

Cooperative Relaying — Although MSs are incapable of sophisticated receive combining or cooperation strategies, relays and multihop-

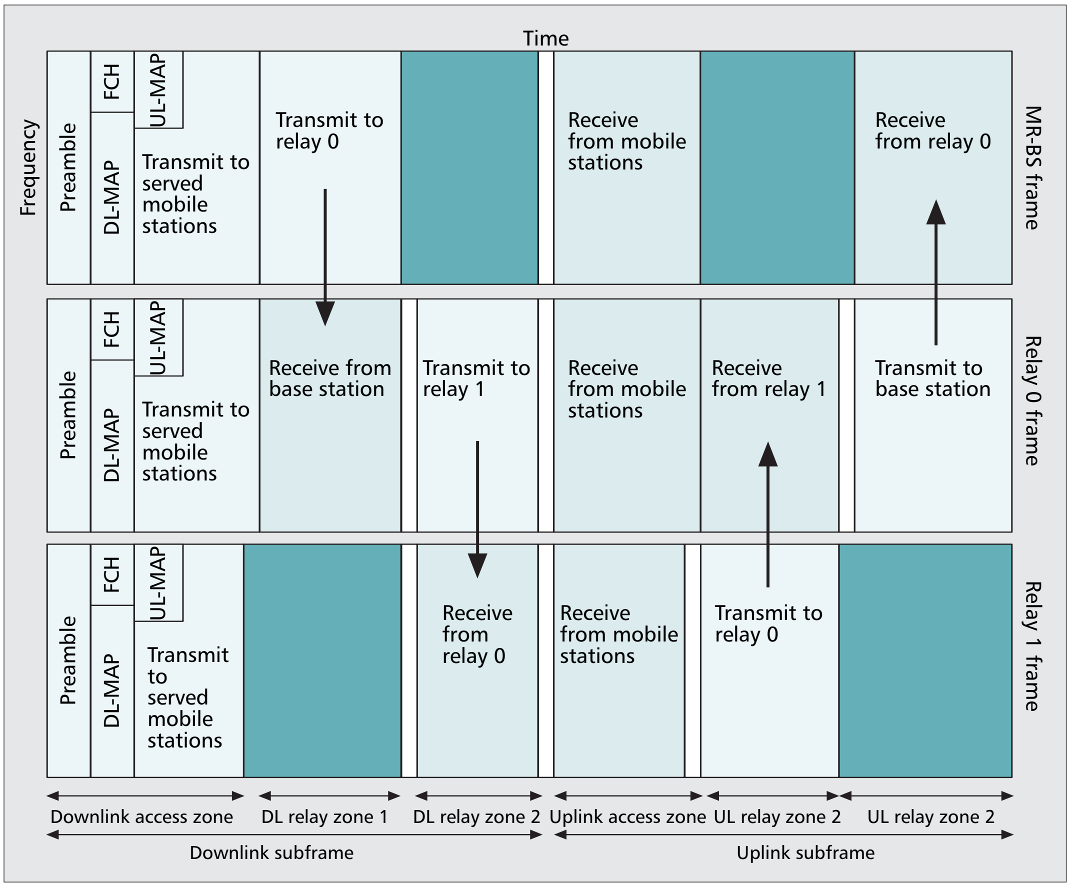

Figure 3. *An example three-hop non-transparent frame structure for the MR-BS and two RSs. To support multihop communication, the relay zone is segmented so that Relay 0 receives from the MR-BS in DL Relay Zone 1 and transmits to Relay 1 in DL Relay Zone 2. A similar concept, where the relay behavior in a DL Relay Zone depends on the frame number, is also defined in 16j.*

IEEE 802.16j provides three mechanisms for cooperative diversity. In each of them, the manufacturer is responsible for ensuring that the delay spread at the receiver is shorter than the length of the cyclic prefix of the OFDM symbol.

capable BSs can work together to provide diversity. IEEE 802.16j provides three mechanisms for cooperative diversity. In each of them, the manufacturer is responsible for ensuring that the delay spread at the receiver is shorter than the length of the cyclic prefix of the OFDM symbol [3]:

- **Cooperative source diversity**: In this mode of cooperation, antennas distributed among relays and the MR-BS transmit identical signals simultaneously in time and frequency.
- **Cooperative transmit diversity**: Cooperative transmit diversity involves utilizing previously defined (16e) space-time codes distributed across antennas at the cooperating relays or MR-BS.
- **Cooperative hybrid diversity**: Cooperative hybrid diversity is a combination of the above two diversity modes. Here, space-time codes again are spread across a subset of relay station and BS antennas, but at least two antennas are transmitting identical signals.

Cooperative relaying is an optional feature of IEEE 802.16j.

MAC Modifications

Scheduling — Much of MAC design deals with who decides which transceivers can transmit, and how those decisions are communicated to the network. In a centralized single-hop network, such as the PMP mode of IEEE 802.16e, the BS ultimately makes these decisions based on information gathered through several different techniques. In particular, 16e utilizes unsolicited bandwidth requests, polling, and contention-based procedures to determine resource allocation.

In a multihop cellular network, however, some MAC intelligence can be given to the relays. IEEE 802.16j allows non-transparent relays to operate in *distributed scheduling* mode, where they make decisions about resource allocation to their subordinate stations, possibly in coordination with the MR-BS. All transparent relays must (and non-transparent relays may) operate in centralized scheduling mode, relying on the MR-BS to allocate its resources.

Since 802.16 is connection-oriented, resource allocation is generally performed on a per-connection basis. Each connection is given a unique (in the MAC domain) connection identifier (CID). Control messages generally are sent on either *basic* or *primary* management connections, which are assigned at network entry. In 16j, these management connections can be allocated by the relay station in what is called *local CID* allocation mode. The MR-BS must grant a set of CIDs to the relay to avoid an overlap. Transport connections, which carry higher layer data, are still allocated by the MR-BS.

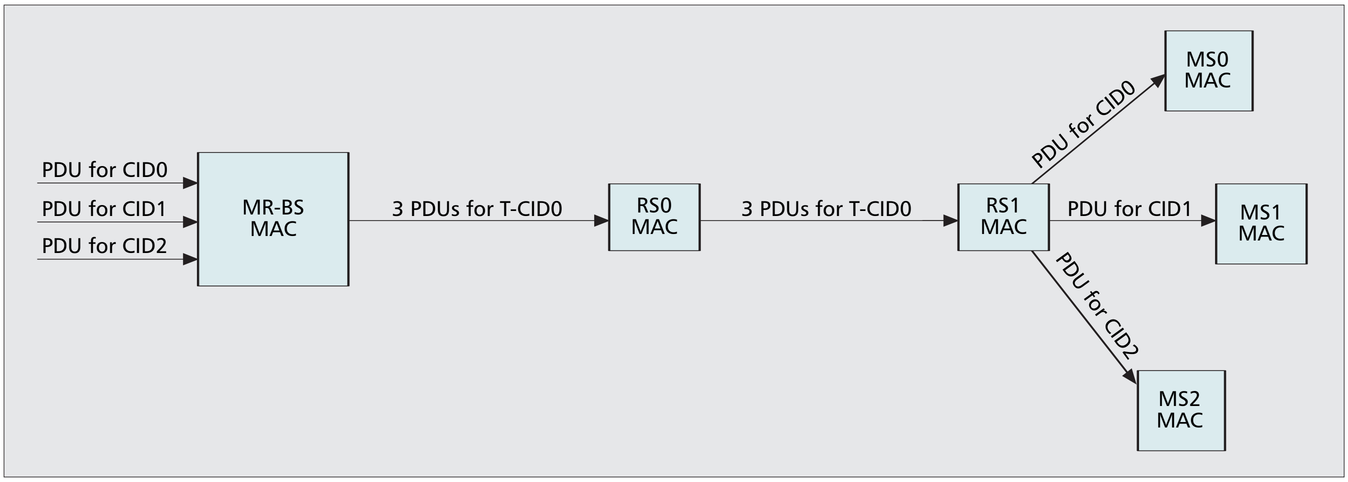

Figure 4. *An example of tunneling using a three-hop architecture serving three mobile stations. The MR-BS receives three PDUs, each destined for a unique MS. Fortunately, all of the MSs are served by the same non-transparent RS. The MR-BS can package them into tunnel packets and send them with the same destination address (T-CID). This keeps the intermediate RS, RS0, from having to keep track of which RS is serving which MS.*

Tunneling — Aside from dealing with how resources to the MS are allocated, the MAC portion of IEEE 802.16j also is devoted to defining a new MAC (called R-MAC) between the MR-BS and the access RS that will not require modifications to MSs. That is, instead of forcing each hop to act as the BS-to-MS hop of 16e, an R-MAC enables intermediate relays *to behave like relays*, instead of some mix between an MS and a BS. A fundamental part of the R-MAC is what is called a *tunnel* connection, which is identified by a tunnel CID (T-CID). A tunnel can carry connections for several different MSs, all being served by the same access RS. Consider the architecture of Fig. 4. Here, the MSs being served by RS1 each have their own connections; however, with the use of a tunnel, the intermediate RS, RS0, views them all as the same connection, thereby simplifying a number of tasks. Separate tunnels, called *management tunnel connections*, can carry management packets over multiple hops. Note that, because intermediate RSs do not distinguish between connections in a tunnel, QoS for the tunnel must meet or exceed that of each connection mapped to the tunnel.

Although connections spanning multiple hops are not required to be sent through a tunnel, tunneling does simplify the relaying process at intermediate relays. For example, suppose a relay is placed on a train, and it is serving all mobile devices on board. This train is moving around a cell, and the access relay continues to be handed off from relay to relay. In this case, only the tunnel must be reestablished during handoff, instead of reestablishing each connection on the train.

To accommodate tunneling, 16e MAC protocol data units (PDUs) are encapsulated into relay MAC PDUs over tunnel connections. The PDU begins with the relay MAC header, which distinguishes itself so that an MS will disregard the PDU. After the relay MAC header, several subheaders may be included that give the receiver information about the payload, including what the QoS demands are (to enable a distributed-scheduling-mode RS to make scheduling decisions about relaying the PDU), when to relay the packet (for a centralized-scheduling-mode RS), or how to reconstruct PDUs to pass to higher layers. The payload, which generally consists of a regular 16e MAC PDU, follows the subheaders, and a cyclic redundancy check may be appended at the end.

HARQ — IEEE 802.16e provides support for HARQ [3], which is mainly implemented in the PHY but requires extensive support from the MAC. HARQ support is extended in 16j to include all of the relay and scheduling modes detailed above. For instance, when the relays on the path to an MS are operating in non-transparent centralized-scheduling mode, a relay saves the PDU until receiving an acknowledgment from its subordinate station. The acknowledgment originates at the MS and is passed back to the MR-BS. If a relay along the path does not correctly receive the burst, it immediately notifies its superordinate station, which passes this negative acknowledgment (NAK) back up the chain to the MR-BS. The MR-BS then schedules the HARQ burst to be resent on the link that failed, and it continues down the path to the MS.

Security — Security,[3] like scheduling, also can be centralized at the MR-BS or distributed among relays. In centralized scheduling, only the MR-BS and MS hold the keys to encrypt or decrypt MAC PDUs. In distributed scheduling, the access relay station can derive keys for new MSs entering its service so that the secure link from MR-BS consists of a secure link from MR-BS to access RS, followed by another secure link from access RS to MS.

RS Grouping — Finally, relay stations can be grouped together to act as a virtual relay. An RS group has one superordinate station and is assigned a multicast connection that carries messages to all of the RSs in the group. Relays inside the group either all transmit the same

[3] *Although security is a sublayer distinct from the MAC, we include it in this section for brevity.*

control information (i.e., preamble, frame control header [FCH], downlink mobile access point [MAP]), or none of them transmit control information. Among other uses, relay grouping serves the purpose of providing transmit diversity to a serviced MS or minimizing handoff between densely packed relay stations.

Remaining Challenges

This section provides a brief summary of what the authors perceive to be the main challenges in implementing a system based on an IEEE 802.16j profile. A detailed analysis of such challenges is beyond the scope of this article.

Because implementing relaying is not the only way to extend coverage, it must be the most cost-effective approach. Relaying, as it is implemented in IEEE 802.16j, covers the broad range from very simple, inexpensive relays to complex and costly relays. A non-transparent relay acting in a distributed mode is nearly as complex as a BS. A transparent relay or a non-transparent relay in centralized mode is much simpler and less expensive but requires an order of magnitude more complexity in the BS. For service providers to adopt a new technology such as relaying, manufacturers must ensure that costs are kept low. At the time of writing, the authors are unaware of any cost analysis performed since 16j matured.

As far as technical issues, frequency reuse, relay placement, resource allocation, and scheduling are very difficult, yet extremely important, problems that IEEE 802.16j has left to manufacturers and providers to solve. In non-transparent mode, for instance, interference between the control information from the BS must somehow be minimized. Even in transparent mode, although downlink transmissions of the BS and relay are orthogonal in frequency, poor frequency reuse or scheduling still can cause unacceptable interference from surrounding relays.

Conclusion

From frame modifications to a relay MAC protocol, 16j required extensive additions to the PHY and MAC of IEEE 802.16. This article introduced the most fundamental of these additions, with emphasis on the wide variety of potential relaying capabilities.

IEEE 802.16j is nearing completion. After Sponsor Ballot approval and final IEEE-SA Standards Board approval, whether or not any part of 16j is deployed rests on the interest of the WiMAX Forum in adopting it into a profile. In the meantime, the research community can enjoy this case study in the practical design of multihop networks.

IEEE 802.16j is nearing completion. After Sponsor Ballot approval and final IEEE-SA Standards Board approval, whether or not any part of 16j is deployed rests on the interest of the WiMAX Forum in adopting it into a profile.

References

[1] C. Eklund *et al.*, "IEEE Standard 802.16: A Technical Overview of the WirelessMAN Air Interface for Broadband Wireless Access," *IEEE Commun. Mag.*, vol. 40, no. 6, June 2002, pp. 98–107.
[2] A. Ghosh *et al.*, "Broadband Wireless Access with WiMAX/802.16: Current Performance Benchmarks and Future Potential," *IEEE Commun. Mag.*, vol. 43, no. 2, Feb. 2005, pp. 129–36.
[3] J. G. Andrews, A. Ghosh, and R. Muhamed, *Fundamentals of WiMAX*, Prentice-Hall, 2007.
[4] IEEE Std. 802.16e-2005, "Amendment to Air Interface for Fixed and Mobile Broadband Wireless Access Systems — Physical and Medium Access Control Layers for Combined Fixed and Mobile Operation in Licensed Bands," 2005.
[5] IEEE Std. 802.16-2004, "Air Interface for Fixed Broadband Wireless Access Systems," 2004.
[6] O. Oyman, J. N. Laneman, and S. Sandhu, "Multihop Relaying for Broadband Wireless Mesh Networks: From Theory to Practice," *IEEE Commun. Mag.*, vol. 45, no. 11, Nov. 2007, pp. 116–22.
[7] IEEE 802.16 Relay Task Group, 2008; http://www.ieee802.org/16/relay/
[8] M. Hart, M. Nohara, and J. J. Son, "Supporting Report to EC for Request of Conditional Approval to Initiate Sponsor Ballot on P802.16j," *IEEE 802.16-08/014*, 2008.
[9] M. Nohara, "Relay TG Session #57 Summary/Closing Remarks," *IEEE 802.16j-08/014r2*, 2008.
[10] Z. Tao, K. H. Teo, and J. Zhang, "Aggregation and Concatenation in IEEE 802. 16j Mobile Multihop Relay (MMR) Networks," *IEEE Mobile WiMAX Symp.*, Orlando, FL, Mar. 2007, pp. 85–90.
[11] L. Erwu *et al.*, "Performance Evaluation of Bandwidth Allocation in 802.16j Mobile Multi-Hop Relay Networks," *IEEE 65th Vehic. Tech. Conf. (VTC2007-Spring)*, Dublin, Ireland, Apr. 2007, pp. 939–43.
[12] Z. Tao *et al.*, "Frame Structure Design for IEEE 802.16j Mobile Multihop Relay (MMR) Networks," *IEEE GLOBECOM '07)*, Washington, DC, Nov. 2007, pp. 4301–06.
[13] B. Lin *et al.*, "Optimal Relay Station Placement in IEEE 802.16j Networks," *ACM Int'l. Conf. Wireless Commun. Mobile Comp.*, Honolulu, HI, Aug. 2007, pp. 25–30.
[14] J. Sydir *et al.*, "Harmonized Contribution on 802.16j Usage Models," IEEE C802.16j-06/015, 2006.
[15] M. Hart *et al.*, "Out-of-Band Relay Clarification," IEEE C802.16j-08/079r1, 2008.

Biographies

STEVEN W. PETERS (peters@ece.utexas.edu) is a Ph.D. student at the University of Texas at Austin. He received B.S. degrees in electrical engineering and computer engineering from the Illinois Institute of Technology in 2005 and an M.S.E. degree in electrical engineering from the University of Texas at Austin in 2007. His research focuses on interference mitigation and general physical layer design in multihop cellular networks. He is also a part-time consultant on wireless systems research and development.

ROBERT W. HEATH, JR. (rheath@ece.utexas.edu) received his B.S. and M.S. degrees from the University of Virginia and his Ph.D. from Stanford University, all in electrical engineering. He is currently an associate professor in the Department of Electrical and Computer Engineering at the University of Texas at Austin and associate director of the Wireless Networking and Communications Group. He has consulted with several international companies and is the founder of MIMO Wireless Inc. He was an Editor for *IEEE Transactions on Communications*, an Associate Editor for *IEEE Transactions on Vehicular Technology*, and is currently a member of the Signal Processing for Communications Technical Committee of the IEEE Signal Processing Society. He is the recipient of the David and Doris Lybarger Endowed Faculty Fellowship in Engineering and is a registered Professional Engineer in Texas.