Project	**IEEE 802.20 Working Group on Mobile Broadband Wireless Access** <http://ieee802.org/20/>
Title	**MBFDD and MBTDD Wideband Mode: Technology Overview**
Date Submitted	**2006-01-06**
Source(s)	Jim Tomcik Qualcomm, Incorporated 5775 Morehouse Drive San Diego, CA, 92121 Voice: 858-658-3231 Fax: 858-658-2113 Email: jtomcik@qualcomm.com
Re:	MBWA Call for Proposals
Abstract	This contribution (part of the MBFDD and MBTDD proposal packages for 802.20), contains the MBFDD and MBTDD Technology Overview Document. In this paper, we describe a complete and compliant technical proposal for a Mobile Broadband Wireless Access (MBWA) system that meets the requirements for the future IEEE 802.20 standard. Both TDD and FDD technologies are included in this document, since there is much in common between the two approaches. We describe the physical layer (PHY) and medium access control (MAC) protocols that are part of the proposed specification. The physical layer uses non-orthogonal and orthogonal multiple access schemes and supports deployment bandwidths from 5 MHz to 20 MHz. The technologies are optimized for high spectral efficiency and mobile operation in a wide area deployment. To improve capacity and coverage, the system employs advanced techniques such as multiple input multiple output (MIMO) transmission, fractional frequency reuse (FFR), precoding, and spatial division multiple access (SDMA). The MAC protocols are designed to provide multiple quality of service (QoS) levels and provide optimized user experience in a mobile environment. This paper provides details of how these features are incorporated in the system design.
Purpose	For consideration of 802.20 in its efforts to adopt FDD and TDD proposals for MBWA.
Notice	
Release	
Patent Policy	

This page is intentionally left blank.

Contents

1 Introduction

This paper describes a proposed system design for the future IEEE 802.20 TDD and FDD Mobile Broadband Wireless Access (MBWA) standard, designed for wide-area mobile operation in licensed frequency bands. The system is designed to provide superior performance in a wide variety of deployments including macrocellular, microcellular, and hotspots.

The features of the MAC layer are described in Section 2 through Section 5. The network architecture including the protocol layering and terminal addressing are described in Section 2. The system provides mechanisms to provide different quality of service levels. These techniques are described in Section 3. Section 4 describes techniques used by the network to initiate a connection to an idle terminal. Security schemes used for authentication and encryption are described in Section 5.

The features of the physical layer are described in Sections 6 through 26. The physical layer uses a combination of orthogonal frequency division multiple access (OFDMA) for the data channels and code division multiple access (CDMA) for some of the reverse link control channels. The system can be deployed in flexible bandwidths from 5 MHz to 20 MHz. With multi-input multi-output (MIMO) transmission, peak data rates over 260 Mbps are supported in 20 MHz bandwidth. The system supports both frequency division duplex (FDD) and time division duplex (TDD) modes of operation. The TDD mode provides flexible allocation of resources between forward and reverse links.

Physical layer framing for FDD and TDD modes are described in Sections 6 and 7, respectively. These sections also describe the various physical layer channels. The framing structure provides packet transmission latency as low as 5.5 ms, enabling efficient support of delay-sensitive applications. The system utilizes turbo coding and high order modulation to provide high spectral efficiency. The coding and modulation schemes are described in Section 8. The mechanisms for allocating resources to the data channels are described in Section 8. The procedures used by terminals to acquire the system are presented in Section 10. The system has an access channel designed to enable terminals to access the airlink with very low latency. Access channel procedures are given in Section 11. Frequency hopping is used along with OFDMA for increased diversity. Section 12 describes the two hopping modes supported in the system. The system employs control channels that have been designed to support the various transmission modes with low overhead. Forward and reverse link control channels are described in Sections 13 and 14, respectively. Power control is essential in a broadband wide area network to control interference and enhance capacity. The power control schemes are described in Section 15. Section 16 describes a robust handoff mechanism that provides seamless connectivity as terminals move from one cell to another. While the system uses orthogonal multiple access on the reverse link data channel, the capacity of such schemes does not scale linearly with the number of receive antennas at the access point. To overcome this issue, the system has a quasi-orthogonal reverse link that is described in Section 17. While the system is designed to operate with universal frequency reuse, it can also employ fractional frequency reuse (FFR) which is an interference management technique that provides better user experience at cell boundaries. System features to support FFR are described in Section 18. In a broadband fading channel, capacity gains can be obtained by scheduling each terminal on that part of the band where it sees better channel conditions. Support for subband scheduling is described in Section 19. Sections 20 through 23 describe various multiple-antenna transmission techniques that can be employed in the system for improved spectral efficiency and coverage. These include multi-input, multi-output (MIMO) transmission, precoding, beamforming (for TDD), and space division multiple access (SDMA). Section 24 describes an optional mode that can be used in wideband deployments to support terminals capable of utilizing only part of the bandwidth. Since the system is expected to co-exist with other wireless technologies and support multi-mode terminals that are capable of utilizing other technologies also, the air interface provides techniques for terminals to handoff to and from other systems. These techniques are discussed in Section 25. Section 26 describes a provision in the system for reserving bandwidth for embedding other physical layers, for future expandability of system capabilities. Finally, some concluding remarks are given in Section 27.

The system supports both TDD and FDD modes of operation. Except for Section 6 which applies only to FDD, and Sections 7 and 22 which apply only to TDD, the remaining sections of the document describe concepts that apply to both TDD and FDD operation, though there may be some differences in the details of how the concepts are applied in the two modes.

2 Network Architecture

The Figure 2-1 shows a possible network architecture for a system deployment based on the proposed MBWA air interface.

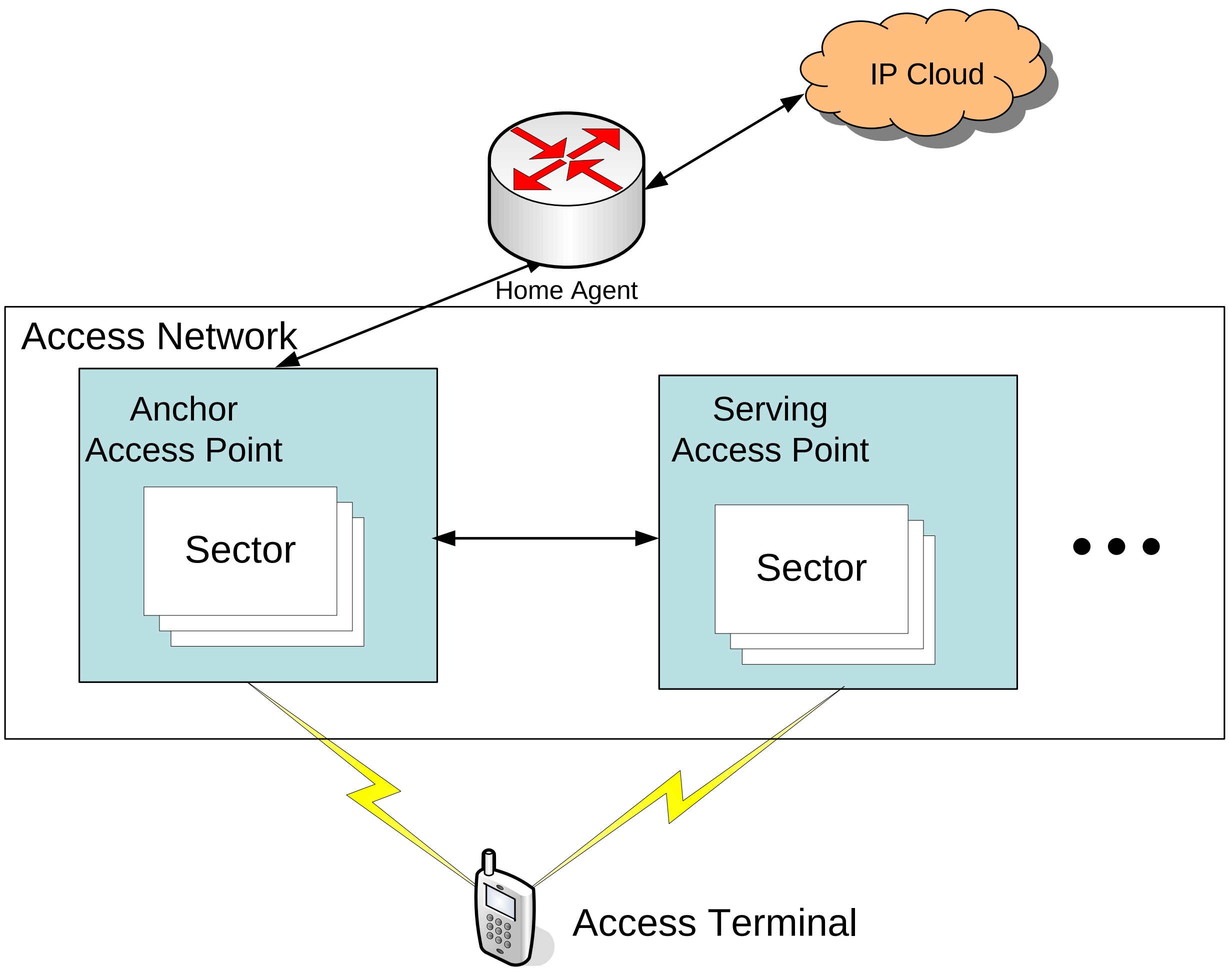

Figure 2-1 The network architecture for MBWA

The MBWA system is designed to handle high speed communication to mobile users with minimal or no degradation to user experience as the users hands off between access points. In the MBWA system, the access terminal receives service by one or more access points at a given time. Each access point may have one or more sectors to better utilize the air link resources. The access terminal keeps a list of best visible sectors in a list called active set. The active set is maintained by both the access terminal and access network, consists of sectors that the AT may choose to switch to at any time. The access network is designed to minimize the switch time between active set members. At any given time, the access terminal may be served by one sector per link, forward or reverse. The access point housing a serving sector is called Serving Access Point. At any given time only one Access Point provides connectivity to the Internet for a given access terminal. The Access Point providing Internet connectivity is called the Anchor Access Point. The Serving Access Point changes over time based on radio conditions. The change of

Serving Access Point is called Layer2 handoff. The Anchor Access Point may be changed to minimize the number of hops the packet has to travel before reaching the AT. The change of Anchor Access Point is called Layer 3 handoff. Layer 3 handoff is designed to be independent of Layer 2 handoff, enabling fast switching of Serving Access Point. Layer 3 handoff may be facilitated by a home Agent. Layer 3 handoff is also facilitated by the Route Selection Protocol in the Data Transport, which allows for make before break Layer 3 handoffs.

2.1 Access Terminal Addressing

There are two main addresses used to refer to an access terminal in a MBWA system. They are:

UATI: Universal Access Terminal Identifier. This is a 128 bit temporary identity given to the access terminal by the system. The UATI is mainly used for accessing the system. The UATI is not derived from any hardware ID given to the system such as IEEE EUI, hence it can not be used to resolve the identity of the access terminal. There is also a shortened 32 bit version of UATI called ATI, which consists of the lower 32 bits of the UATI. The ATI is used to page the access terminal.

MAC ID: The access terminal is assigned one MAC ID per sector it has in its active set. The MAC ID consists of 11 bits. The MAC ID is used by the sector to exchange unicast packets with the AT. The MAC ID is unique only within the sector that assigned the MAC ID to the access terminal.

In addition to these, the AT may have the following addresses. The MBWA system neither assumes nor relies on the existence of the following addresses to operate.

IP Address: The IP address given by the network to the access terminal. The MBWA system does not require an IP address assigned to the access terminal to operate. For management purposes, an IP address may be assigned to the access terminal.

IEEE EUI-48 or EUI-64: The EUI-48 or EUI-64 given to the access terminal at manufacturing time. To protect the identity of the access terminal, the EUI-48 or EUI-64 is never sent in the clear over MBWA channels. They may only be retrieved when full security is established with the system, over encrypted channels. Hence the MBWA system does not rely on use of this identity, instead it provides mechanisms to retrieve it over encrypted channels.

2.2 Protocol Architecture

Figure 2-2 shows the layering of protocols of the MBWA system. The air interface specification specifies the Bearer and Non-Bearer protocols only. Management protocols are not described in this document.

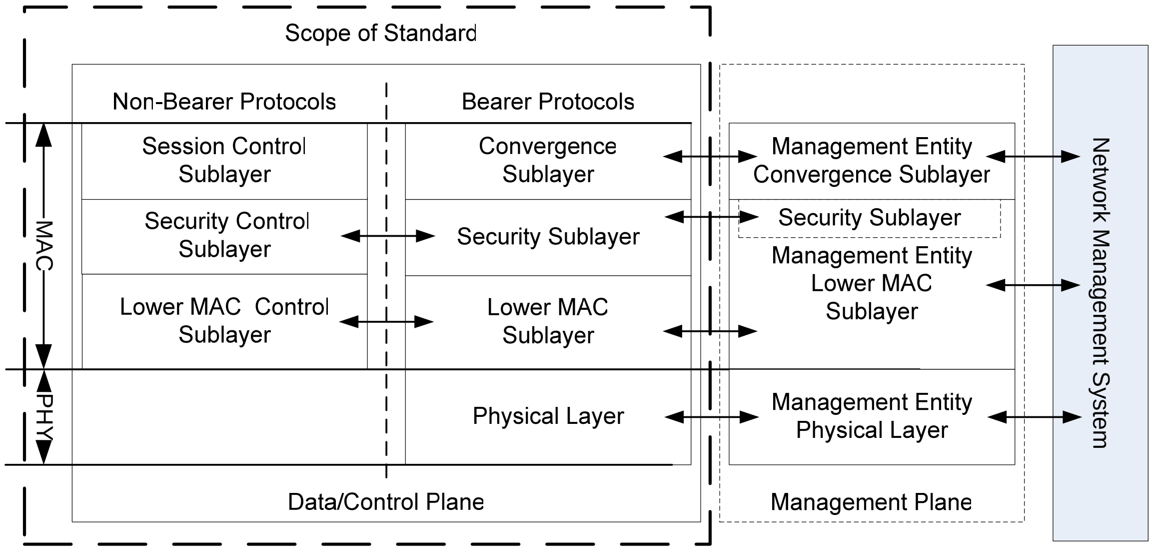

Figure 2-2 Protocol layering for MBWA air interface

Each layer and their responsibilities are:

Session Control Sublayer
: The Session Control Sublayer provides UATI management, protocol negotiation, protocol configuration, and state maintenance services. The Session Control Sublayer is a non-bearer layer and, therefore, it does not carry payload on behalf of other layers.

Convergence Sublayer
: The Convergence Sublayer provides protocols and transports used to transport messages and data, and provides multiplexing of distinct transports. For example, it provides the Default Signaling Transport for transporting air interface protocol messages and the Default Packet Transport for transporting user data. It is extensible in the sense that new transports can be defined to carry other types of packets.

Security Control Sublayer
: The Security Control Sublayer provides key exchange, and manages the Security Sublayer.

Security Sublayer
: The Security Sublayer provides authentication, and encryption services.

Lower MAC Control Sublayer
: The Lower Medium Access Control (MAC) Control Sublayer provides air-link connection establishment and maintenance services.

Lower MAC Sublayer	The Lower MAC Sublayer defines the procedures used to receive and to transmit over the Physical Layer.
Physical Layer	The Physical Layer provides the channel structure, frequency, power output, modulation, and encoding specifications for the Forward and Reverse Channels.

Each layer may contain one or more protocols or transports. Protocols use signaling messages, in-band messages, blocks, or headers to convey information to their peer protocols at the other side of the air-link. When protocols send messages they use the Signaling Network Protocol (SNP) defined in the Signaling Transport to transmit these messages. Transports also send messages using the Signaling Network Protocol. Blocks are information conveyed to a peer protocol using an encapsulation that is specific to a Physical Layer Channel. For example, the Lower MAC Control Sublayer Overhead Messages Protocol uses the SectorInfoBlock to carry information to its peer protocol at the access terminal on the primary broadcast channel (pBCH).

Figure 2-3 shows all the protocols within each sublayer.

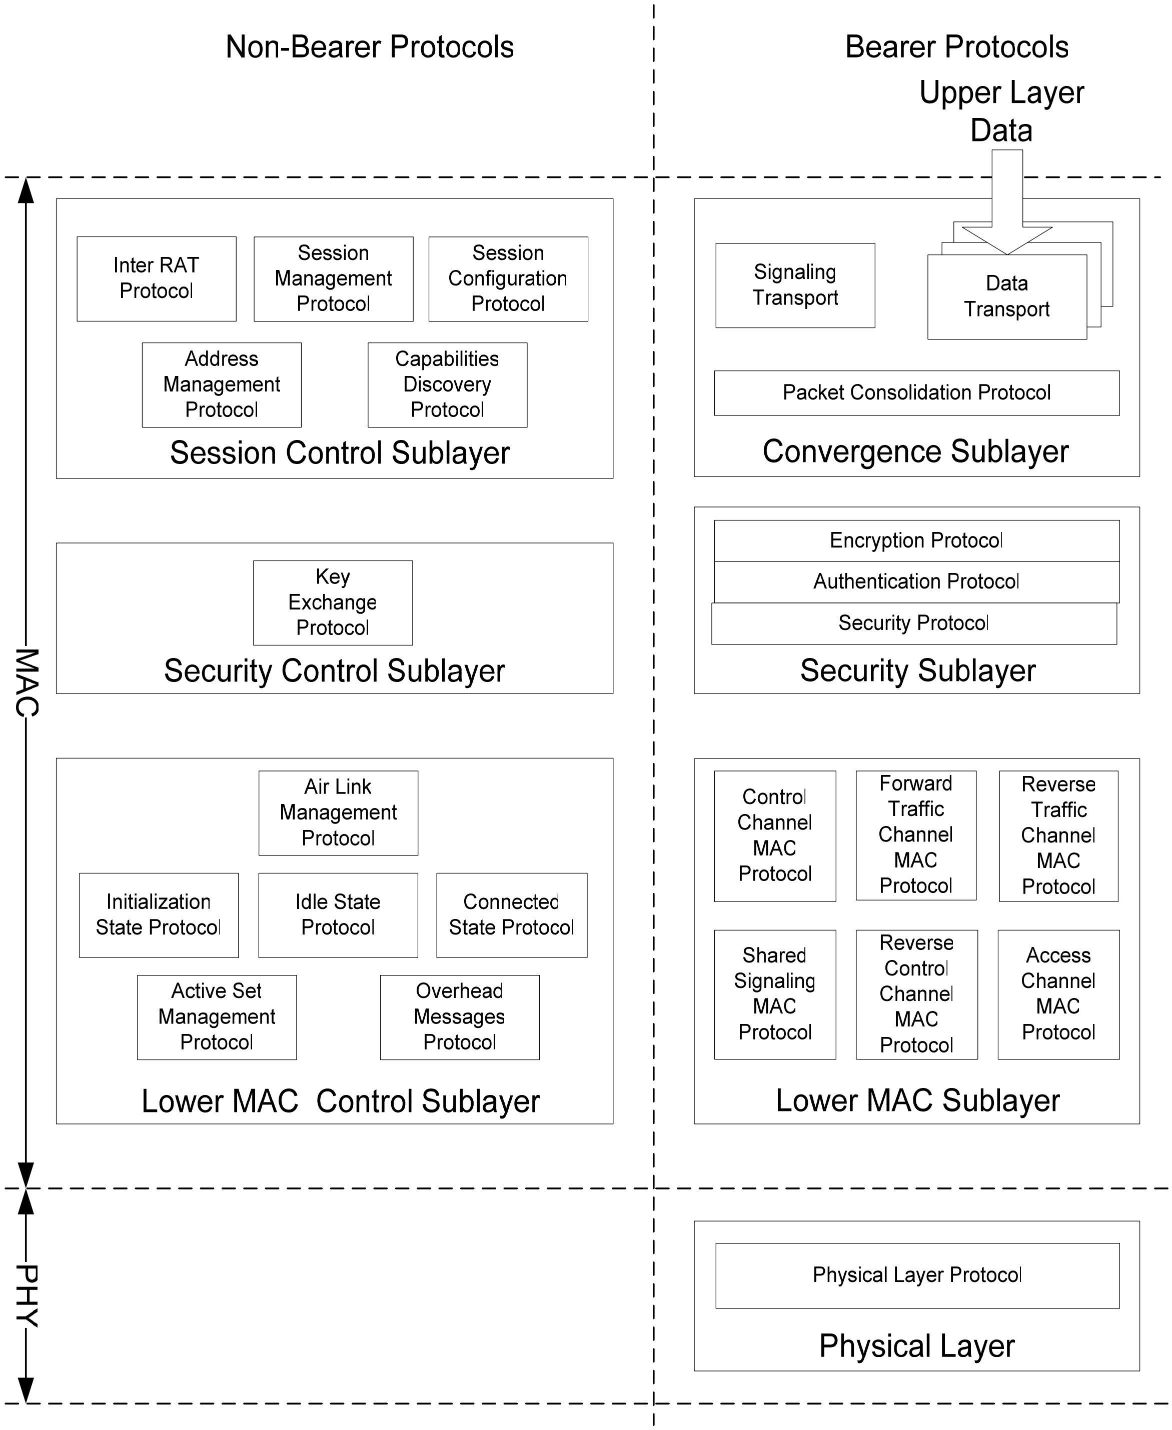

Figure 2-3 Protocols within each sublayer

- Session Control Sublayer:
  - Session Management Protocol: Provides means to control the activation and the deactivation of the Address Management Protocol, Capabilities Discovery Protocol and the Session Configuration Protocol. It also provides a session keep-alive mechanism.
  - Address Management Protocol: Provides UATI management.
  - Capabilities Discovery Protocol: Provides means for the access network to discover the capabilities of the access terminal.
  - Session Configuration Protocol: Provides means for negotiation of the SessionConfigurationToken used in the session.
  - Inter RAT Protocol: Provides means for sending and receiving other technology messages between access terminal and access network.
- Convergence Sublayer:
  - Default Signaling Transport: Provides message transmission services for signaling messages. This transport is also responsible for fragmentation, reassembly and reliable delivery of messages.
  - Default Packet Transport: Provides transmission of upper layer data. This transport is responsible for fragmentation, reassembly, reliable delivery and flow control. This transport also contains the RouteSelectionProtocol, which allows for make before break Layer3 handoff.
  - Packet Consolidation Protocol: Adds the Packet Consolidation Protocol header to transport packets prior to transmission; and, after reception, removes the Packet Consolidation Protocol header and forwards the transport packets to the correct transport. Provides transmit prioritization and packet encapsulation for the Convergence Sublayer.
- Security Control Sublayer:
  - Key Exchange Protocol: Provides the procedures followed by the access network and the access terminal to exchange security keys for authentication and encryption using industry standard 4 way key exchange.
- Security Sublayer:
  - Authentication Protocol: Provides the procedures followed by the access network and the access terminal for authenticating traffic using SHA-1.
  - Encryption Protocol: Provides the procedures followed by the access network and the access terminal for encrypting traffic using AES.
  - Security Protocol: Provides procedures for generating a cryptosync based on the information fetched from the Lower MAC Sublayer that can be used by the Authentication Protocol and the Encryption Protocol. . This protocol incurs zero over the air overhead for cryptosync.

- Lower MAC Control Sublayer:
  - Air Link Management Protocol: Provides the overall state machine management that an access terminal and an access network follow during a connection.
  - Initialization State Protocol: Provides the procedures that an access terminal follows to acquire a network and that an access network follows to support network acquisition.
  - Idle State Protocol: Provides the procedures that an access terminal and an access network follow when a connection is not open.
  - Connected State Protocol: Provides the procedures that an access terminal and an access network follow when a connection is open.
  - Active Set Management Protocol: Provides the means to maintain the active set between the access terminal and the access network.
  - Overhead Messages Protocol: Provides broadcast messages and blocks containing information that is mostly used by Lower MAC Control Sublayer protocols.
- Lower MAC Sublayer:
  - Control Channel MAC Protocol: Provides the procedures followed by the access network to transmit, and by the access terminal to receive, the Control Channels.
  - Access Channel MAC Protocol: Provides the procedures followed by the access terminal to transmit, and by the access network to receive, the Access Channel.
  - Shared Signaling MAC Protocol: Provides the procedures followed by the access network to transmit, and by the access terminal to receive, the physical layer channels controlled by this protocol.
  - Forward Traffic Channel MAC Protocol: Provides the procedures followed by the access network to transmit, and by the access terminal to receive, the Forward Traffic Channel.
  - Reverse Control Channel MAC Protocol: Provides the procedures for the access terminal to transmit, and the access network to receive, the Reverse Control Channels.
  - Reverse Traffic Channel MAC Protocol: Provides the procedures followed by the access terminal to transmit, and by the access network to receive, the Reverse Traffic Channel.
- Physical Layer:
  - Physical Layer Protocol: Provides channel structure, frequency, power output, and modulation specifications for the forward and reverse links.

2.3 Packet Framing

Figure 2-4 shows the framing of packets in the MBWA system.

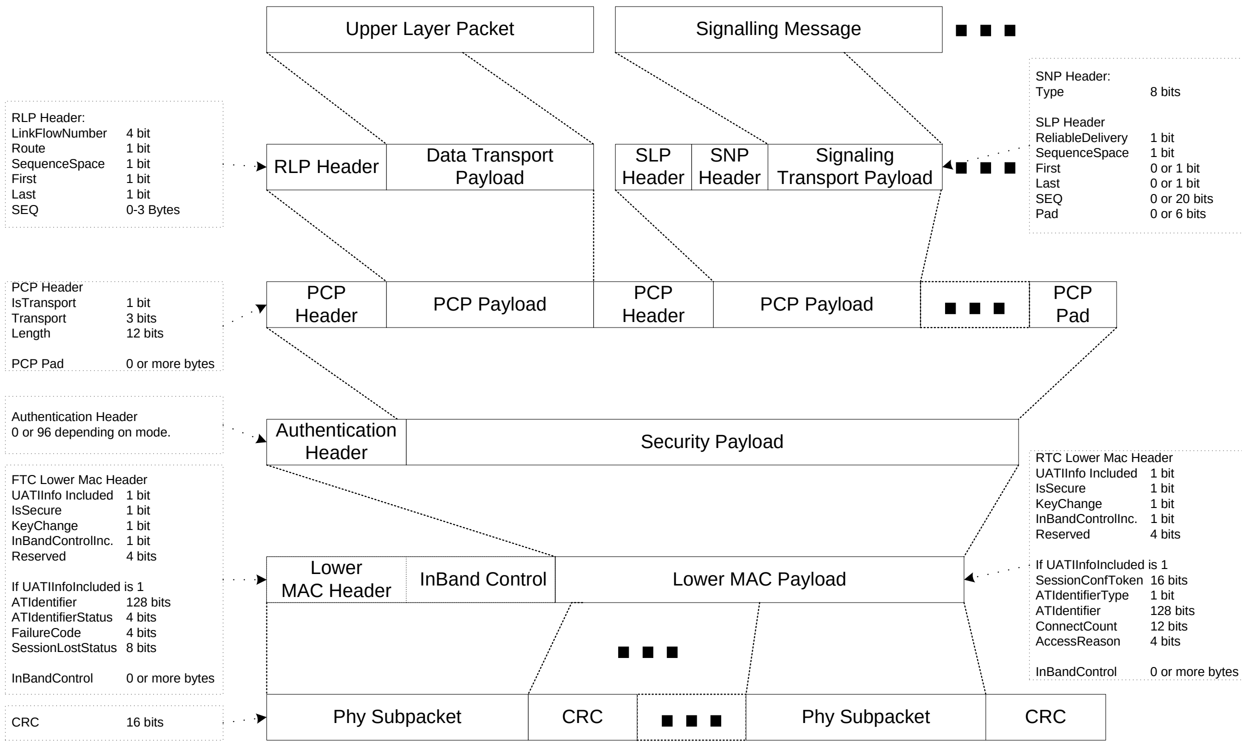

Figure 2-4 Packet Framing in MBWA

Both Data Transport and Signaling Transport provide facilities for fragmenting their payload. A PCP (PacketConsolidationProtocol) packet contains zero or more Data Transport packets and zero or more Signaling Transport packets. The PCP Pad is set to all zeros. Phy layer subdivides a MAC packet into subpackets each of which have their own CRC protection.

2.4 Session vs. Connection

A session refers to a shared state between the access terminal and the access network. This shared state stores the protocols and protocol configurations that were negotiated and are used for communications between the access terminal and the access network.

Other than to open a session, an access terminal cannot communicate with an access network without having an open session.

An access terminal with an open session may be in one of the three states.

- Connected: In this state the access terminal is assigned a MAC ID from at least one sector and it regularly monitors the overhead and assignment channels and sends CQI (Channel Quality Indicator) to the access network with certain periodicity. If the access terminal does not have traffic resources, it can be assigned traffic resources within one physical layer frame (~ 1msec). Switching to and from Connected state is referred as opening and closing the connection.
- Monitor: In this state the access terminal has no MAC ID assignment from any sector. It continuously monitors the overhead and paging channels of the access network. The access terminal may be paged within a superframe and the access terminal needs to make an access to switch to the connected state.
- Idle (Sleep): Operationally, this state is identical to Monitor state. In addition, in this state the access terminal and the access network agree on a paging cycle. The access terminal does not need to receive any transmission from the access network between the paging cycles.

During a single session, the access terminal and the access network can open and can close a connection multiple times.

2.5 Session Token for Upgradeability

MBWA air interface is designed in such a way that each protocol or transport within a sublayer can be upgraded without necessarily updating the whole sublayer or all the protocols. This allows for upgrading or enhancing portions of the air interface without a need to change the whole specification. The SessionConfigurationToken is a 16 bit value which identifies a complete set of protocol and transport instances that can be used to communicate between the access terminal and the access network.

The session of an access terminal may contain multiple SessionConfigurationTokens. A SessionConfigurationToken is InUse if the set of protocol and transport instances specified by the SessionConfigurationToken are currently being used to communicate between the access terminal and the access network. Otherwise, a SessionConfigurationToken is Suspended. Only one SessionConfigurationToken shall be InUse at a time.

The Session Configuration Protocol executes its save and commit procedures to swap the InUse SessionConfigurationToken with a Suspended SessionConfigurationToken as shown in Figure 2-5.

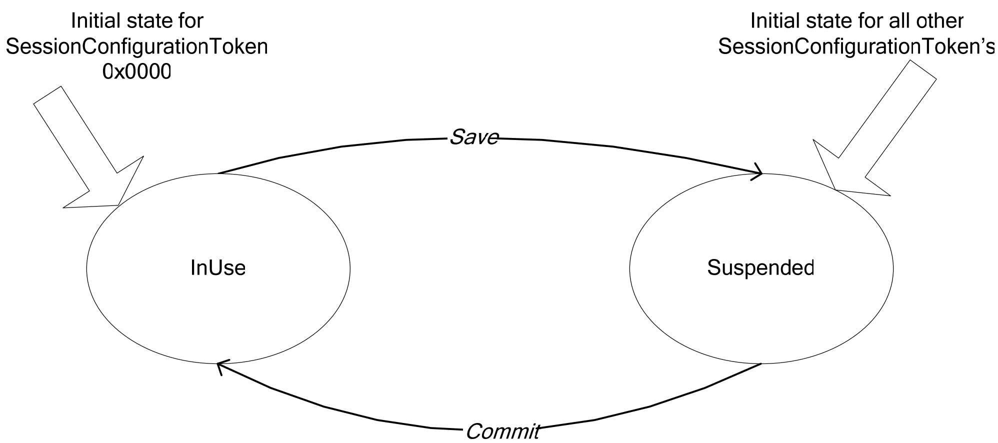

Figure 2-5 SessionConfigurationToken state diagram

3 Quality of Service

The proposed air interface supports link level QoS between the access network and the access terminal for packet data. The air interface supports IP layer per hop and end-to-end QoS services such as DiffServ (RFC 2475) and RSVP.

QoS support in the air interface protocol is provided by the:

- Radio Link Protocol (RLP): Provides configuration, negotiation, and enforcement of QoS policies.
- Reverse Request Channel (R-REQCH): Provides differentiation of user request priorities for scheduling the reverse data channel.
- Access Channel (R-ACH): Segments users into different access classes and allows for different access priority depending on user QoS.

RLP is a best effort data transmission protocol defined by the Default Data Transport. RLP provides optional retransmission of data to achieve a lower packet error rate than the physical layer can achieve alone. RLP carries one or more packet flows.

RLP supports the following functionality:

- Negotiated QoS per flow[1]
- Duplicate detection
- Optional in-order packet delivery of packets to the upper layers
- Fragmentation and reassembly
- Padding
- Packet discard, e.g., due to latency, partial packet loss, handoff
- Addressing on a per-flow basis

Each RLP instance defines 15 flows and a set of attributes to support QoS on a per flow basis. Flows are individually configured for different QoS service types. For example, a flow's attributes may be configured to support different bit error rates or packet latencies. If more than 15 QoS service types are simultaneously needed for a single access terminal, multiple data transports may be used.

A packet requiring QoS is associated with a reservation. A reservation includes an attribute defining the service requirements for the flow, as well as a filter attribute for packet classification. Another attribute defines the mapping of reservations to RLP flows. Multiple reservations may map to a single flow. All QoS reservations are initiated by the access terminal, but the access network may suggest a set of service requirements or filter attributes to the access terminal. Attributes are defined independently for the forward and reverse links

[1] Compliant to the requirements of IEEE802-20-PD-06 "System Requirements for IEEE 802.20 Mobile Broadband Wireless Access Systems – Version 14"

The main QoS operations defined by the air interface are:

- Configuration and removal of QoS parameters on a per flow and reservation basis based on authorized QoS requirements.
- Activation and deactivation of QoS upon request, including admission control and allocation of air interface resources.

3.1 Configuration and Removal of Flows and Reservations

The configuration of a flow or reservation may be done independently from its activation. Configuration should be done prior to activation, but some attributes may be modified for an active flow. For example, the flow and filter templates of a reservation or the packet error rate of a flow may be modified while a reservation is open.

Configuration may be done in advance when the session is established or when an application starts. Advance configuration allows for quick activation at the time of data communication. Configuration is a function of the QoS authorization for the access terminal and system operator QoS policies.

The access terminal uses the Reservation*KK*QoSRequest and Reservation*KK*PacketFilter attributes to signal to the access network the respective service requirements and filter attributes of a new reservation.

Reservations may be removed and reused to support a different QoS service type.

3.2 Activation and Deactivation of Flows and Reservations

A flow needs to be active and a reservation needs to be open in order to be used.

The air interface provides attributes to activate and deactivate a flow. For example, a flow may be deactivated to configure its attributes, or to limit the number of simultaneous active flows.

A reservation transitions from close to open via messaging. An access terminal may request that a reservation be opened, but only the access network may open a reservation. The access network may use admission control and QoS policies to determine if a reservation is opened and how to assign air interface resources to the reservation. Based on the QoS service type and QoS policies, a reservation may be configured to be open or closed by default when a connection is opened.

4 Paging

Paging is the process by which the access network initiates a connection with an access terminal that is in idle state, such that the access terminal wakes to listen to the forward link traffic only at certain negotiated time intervals. The proposed system provides a flexible and efficient paging architecture.

The key features of paging in the proposed system are the following

1. Page timeline
2. Page delivery using QuickPages on the Control Channel and pages on the Forward Traffic Channel.
3. Page loss recovery using Fast Repaging
4. Paging area management using RegistrationRadius field
5. Reliable page delivery upon waking up in a new sector

These techniques allow for reliable page delivery with low paging overhead and low registration overhead.

4.1 Paging Timeline

In idle state, the terminal wakes up periodically and checks if there are any pages for it. The selection of the page period involves the following tradeoff: long page periods increase the delay in initiating a connection and reduce power consumption, while short page periods reduce the delay and increase the power consumption.

In the proposed system, the paging parameters such as the paging period and offset are negotiated between the access terminal and the access network. The paging period may take values from approximately 50 ms to several minutes, allowing different terminals to operate at different points on the latency-power tradeoff.

The proposed system supports an increasing paging timeline, where the page periods are small for the first few paging cycles after entering idle state, and then increase to a larger value for a larger number of paging cycles. This design is explained in Figure 4-1, where the page period is small (Period1) for the first WakeCount paging periods, and larger (Period2) for subsequent page cycles.

Thus, the terminal is more responsive in the initial moments after entering Idle State (at the expense of slightly higher power consumption) and less responsive farther out into the Idle State (with the benefit of reduced power consumption). This terminal behavior is important for data systems where the network is more likely to send data to the terminal immediately after entering Idle State.

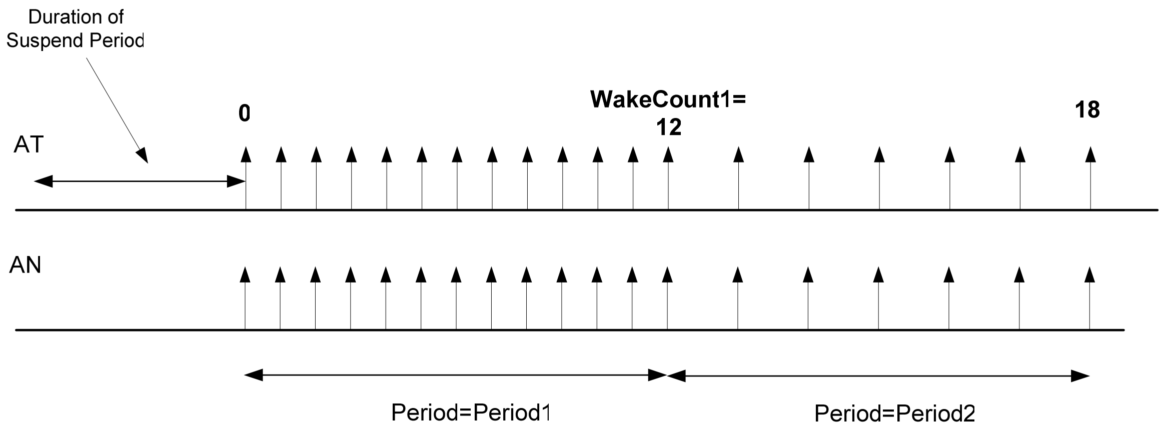

Figure 4-1 Interpage time increases after WakeCount sleep cycles

4.2 Page Delivery

The power consumption for an Idle State terminal depends on the page period and the energy spent during each page period. The page period is selected based on application responsiveness requirements .

In the proposed system, pages are delivered in a two step process using a QuickPage block, followed by a page over a traffic channel. A QuickPage block is transmitted every two superframes (approximately 50 ms).

In order to demodulate the QuickPage block, the terminal needs to demodulate only five OFDM symbols corresponding to pBCH1 in the superframe preamble. For 110 µs OFDM symbol duration, this is approximately 550µs time required to monitor a page. For a page period of 1 second, this translates to a receiver duty cycle of 0.05%. A user who receives a QuickPage monitors a subsequent paging packet that is sent on a Broadcast Forward Traffic Channel. The receipt of a paging packet then triggers the user to access the network.

A QuickPage collision is said to have occurred if a user who was not paged gets a QuickPage. The current proposal reduces the QuickPage collision probability by adopting the following design.

For the purpose of QuickPaging, each user is provided a random 128 bit number (referred to as UserHash). Assume that the QuickPage block is 36 bits long (the exact number may be found in the specification). Further, assume that two users need to be paged. In this case, the network will include the 16 LSBs of the UserHash of both users in the QuickPage block. The users whose 16 LSBs of UserHash match the 16 LSBs in the QuickPage packet will monitor the subsequent paging packet. All other users may power down their receivers. Similarly, when only one user is to be paged, the QuickPage packet contains 32 LSBs of the UserHash, and when three users are to be paged, it carries 10 bits of each user's UserHash.

The above design reduces the probability of QuickPage collision, because in a typical superframe, there may be only one or two pages, and by using 16 or 32 LSBs of the UserHash, the probability that a user who was not paged sees a collision is 1/64k or 1/4G. Such low probabilities are in contrast to QuickPaging schemes that divide the users into (say) 64 bins, and use a 64 bit bitmap to indicate what user was QuickPaged.

4.3 Page loss recovery and Fast Repaging

The fast repaging process is designed to reduce the paging delay in case the terminal misses a page. If a terminal misses a page due to a demodulation error, the terminal is aware of the miss (e.g. by keeping track of CRC on the QuickPage or page packet). If the terminal is aware of a missed page, the Fast Repaging design creates an additional paging opportunity a fixed time interval (FastRepageInterval) after the missed page. FastRepageInterval may be less than the paging period, and thus it reduces the time taken for the missed page to be delivered.

The network sends a fast repage if it does not see an access from the terminal before the FastRepageInterlval expires.

4.4 Paging area management

The paging load on a network depends on the size of the paging area. If a terminal has a large paging area, the paging overhead is high, but the terminal can be highly mobile without missing a page. If a terminal has a small paging area, the paging overhead is low, but the terminal can miss a page if it moves fast. The cost of maintaining a small paging area is to require mobile terminals to register frequently.

Though paging area management is a network side issue, the current proposal includes air interface optimizations to reduce the paging load while maintaining paging reliability. After close of a connection, the access network pages the terminal in a small area consisting of the sector where the access terminal closed the connection, and neighboring sectors. After close of connection, the access terminal registers if it moves to a new sector. By this mechanism, terminals that are stationary after connection close shall be paged in a small area.

After the first registration, the access network pages the terminal in a larger area, and the access terminal registers only after it has crossed a larger number of sectors, as measured through registration zones advertised by sectors, or latitude and longitude advertised by sectors. This allows the access terminal to detect mobility, and register in a larger radius, and also allows the access network to detect mobility, to page in a larger area.

4.5 Reliable page delivery in a new sector

In order to receive a page, the access terminal must know sector specific parameters such as hopping patterns. If these patterns are advertised in overhead messages, then the access terminal must read the overhead messages before demodulating a page. However, an Idle State terminal may cross a sector boundary during a page period, and may wake up in the coverage of a new sector. In this case, the terminal may be unable to demodulate the page because it does not know the sector specific parameters.

This proposal solves this problem by taking two steps. First, the QuickPage channel is self contained in the superframe preamble, i.e., the superfame preamble contains enough information to demodulate the QuickPage block. Second, the structure of the Forward Traffic Channel that carrier pages is described by the QuickChannelInfo message in the preamble of a superframe. Thus, a user who woke up in a new sector will be able to read the channel structure in the superframe preamble, and then demodulate the page.

5 Security

The proposed system uses industry standard procedures for encryption and authentication. To minimize the overhead on transmitted packets, the cryptosync (or nonce) is computed using implicit parameters associated with the packet transmission. The cryptosync is a shared bit-string that is known by both the sender and receiver (and also by any possible attacker). The cryptosync must change from packet to packet, and encryption algorithms forbid repetition of a cryptosync.

5.1 Cryptosync

For successful encryption, the cryptosync must never be reused for the same key. To meet this objective, the MBWA system derives the cryptosync using the following fields

1. System Time
2. ConnectionCount
3. Channel identifier (FL or RL)
4. PilotPN (on the sector that receives/transmits the packet)

The MBWA system provides a system time field that is available to both the receiver and the sender. Each frame has a unique 40 bit PHYFrameNumber, and the frame number where a packet transmission starts is used to generate the cryptosync.

The ConnectionCount field is used to guarantee that a cryptosync is not reused as the access terminal roams between different sectors with misaligned time fields. For each new connection the access terminal makes, the ConnectionCount field is incremented, preventing cryptosync reuse.

The ChannelIdentifier field guarantees that the cryptosync is different for forward and reverse link transmissions that begin in the same PHYFrame.

The PilotPN field plays a role during fast handoff between sectors that may have slight time offsets. Fast handoff occurs without incrementing the ConnectionCount field. The MBWA specification allows fast handoff only between sectors with different PilotPN values, and thus the cryptosync is not repeated.

The MBWA cryptosync requires zero over the air overhead.

5.2 Key Exchange

A four way key exchange is used to derive a session key from a pairwise master key. The pairwise master key is assumed to be negotiated by higher layer protocols that are outside the scope of the MBWA air interface specification. The higher layer key negotiation protocols may be based on the 802.1x standard.

The key exchange produces a 128 bit key for authentication and a separate 128 bit key for encryption.

5.3 Encryption

The MBWA system uses encryption based on the AES-128 standard (NIST PUB 197). Each packet is encrypted using the negotiated key and the cryptosync described earlier. All data and signaling packets except the overhead messages and the messages used for four way key exchange are required to be encrypted.

5.4 Authentication

The MBWA system uses authentication based on the HMAC algorithm specified in RFC 2104 and SHA-256 specification described in FIPS Pub 180-2. Two possible modes of authentication are supported. In one mode, all packets are authenticated using an authentication header that contains an SHA-256 authentication signature. In the other mode, only packets related to access (including ConnectionRequest and ConnectionResponse) are authenticated, in order to reduce overhead.

6 FDD Frame Structure

6.1 Channel Structure

The proposed system includes several physical layer channels on the forward and reverse links, and these channels are configured and controlled by MAC layer protocols, as illustrated in Figure 6-1 and Figure 6-2. The following is a brief description of each Physical Layer Channel.

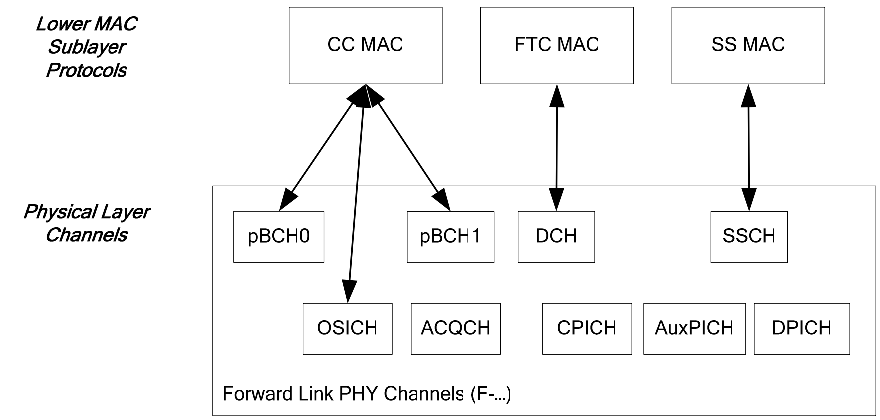

Figure 6-1 Forward channel structure

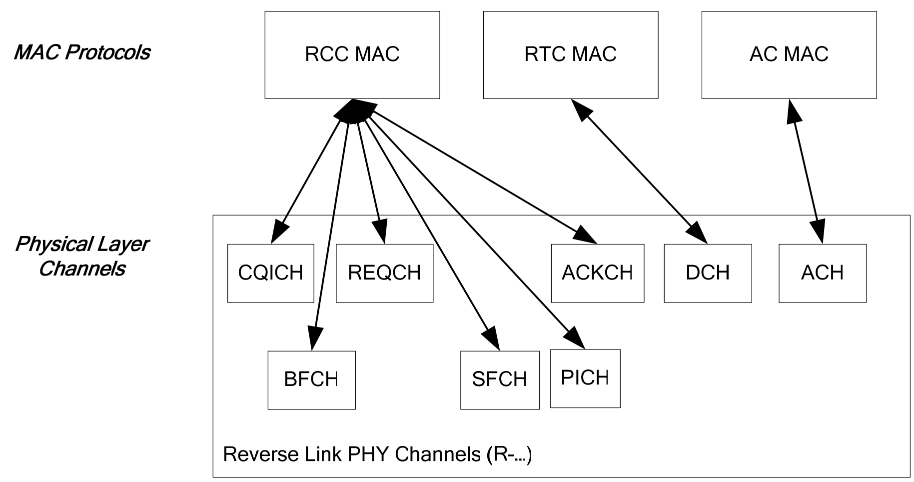

Figure 6-2 Reverse channel structure

6.1.1 Forward physical channels

Forward Acquisition Channel (F-ACQCH)
: Carries an acquisition pilot for an access terminal to use to acquire the system.

Forward Auxiliary Pilot Channel (F-AuxPICH)
: Carries auxiliary pilots for channel estimation from multiple transmit antennas. The Forward Primary Broadcast Channel 1 (F-pBCH1) indicates whether the F-AuxPICH is present.

Forward Common Pilot Channel (F-CPICH)
: Carries the common pilot.

Forward Data Channel (F-DCH)
: Carries information for a specific access terminal. A Forward Data Channel assignment is assigned to an access terminal by a Forward Shared Signaling Channel (F-SSCH) assignment. Also carries broadcast information including pages and sector specific messages.

Forward Dedicated Pilot Channel (F-DPICH)
: Carries the dedicated pilot. This channel is present in BlockHopping mode, which is indicated over the Forward Primary Broadcast Channel 0 (F-pBCH0).

Forward Other Sector Interference Pilot Channel (F-OSICH)
: Carries information about the interference from other sectors to be received by all access terminals.

Forward Primary Broadcast Channel 0 (F-pBCH0)
: Carries information about the system to be received by all access terminals.

Forward Primary Broadcast Channel 1 (F-pBCH1)
: Carries information about the sector to be received by all access terminals. Also carries quick pages.

Forward Shared Signaling Channel (F-SSCH)
: Carries forward and reverse link data channel assignments, access grants, power control commands, and acknowledgement information for Reverse Data Channel (R-DCH) transmissions.

6.1.2 Reverse physical channels

Reverse Access Channel (R-ACH)
: Used by access terminals to initiate communication with the access network. The Reverse Access Channel is also used by access terminals to obtain timing corrections.

Reverse Acknowledgement Channel (R-ACKCH)
: Carries acknowledgement information of a Forward Data Channel (F-DCH) reception.

Reverse Beam Feedback Channel (R-BFCH)
Carries information about the beam index and the quality of the forward link channel.

Reverse Channel Quality Indicator Channel (R-CQICH)
Carries information about the quality of the forward link channel of a sector as received by an access terminal. The Reverse Channel Quality Indicator Channel also carries information about the desired forward link serving sector.

Reverse Data Channel (R-DCH)
Carries information from an access terminal. The Reverse Data Channel is assigned to an access terminal by a Forward Shared Signaling Channel (F-SSCH) assignment.

Reverse Pilot Channel (R-PICH)
Carries the pilot.

Reverse Request Channel (R-REQCH)
Carries information about the buffer level at different quality of service classes for an access terminal. The Reverse Request Channel also carries information about the desired revere link serving sector.

Reverse Subband Feedback Channel (R-SFCH)
Carries information about the quality of a subband of the forward link channel.

6.2 Basic Numerology

Transmission on both links of the MBWA system is divided into units of OFDM symbols. Three basic FFT sizes are defined, namely 512, 1024 and 2048. Each FFT size corresponds to a different chip rate, as shown in Table 6-1. These chip rates are chosen so that all three FFT sizes have the same subcarrier spacing, namely 9.6KHz.

The MBWA system supports flexible bandwidth operation. Different operating bandwidths are constructed from the three basic chip rates by using different numbers of guard carriers.

The values of these parameters, along with others like the cyclic prefix duration and the OFDM symbol duration are specified in Table 6-1. Note that the OFDM symbol duration is independent of the bandwidth of operation.

Table 6-1 Basic Numerology

Parameter	512 pt FFT	1024 pt FFT	2048 pt FFT	Unit
Chip rate	4.9152	9.8304	19.6608	Mcps
Subcarrier spacing	9.6	9.6	9.6	kHz
Bandwidth of operation	≤ 5	≤ 10	≤ 20	MHz
Guard carriers	Function of bandwidth	Function of bandwidth	Function of bandwidth	
Cyclic prefix	6.51-26.04	6.51-26.04	6.51-26.04	us
Windowing duration	3.26	3.26	3.26	us
OFDM Symbol duration (for 6.51 us CP)	113.93	113.93	113.93	us

6.3 Frame Structure

Forward and reverse link transmissions are divided into units of superframes. Superframes are further divided into units of PHYFrames. The frame structure of the MBWA system defines the timing of FL and RL PHYFrames within a superframe. Additionally, it defines the relative timing of assignments, acknowledgements, and H-ARQ retransmissions associated with a data packet. This structure is designed to minimize latency of data transmissions while maintaining acceptable processing durations for encoding and decoding at the AT and the AP, as well as scheduling at the AP.

As illustrated in Figure 6-3, a forward link superframe consists of a superframe preamble followed by 24 FL PHYFrames, and a reverse link superframe consists of 24 RL PHYFrames. The superframe preamble carries acquisition sequences plus key overhead parameters that enable an AT to receive the forward link control channels and subsequently access the system. The first RL PHYFrame of each RL superframe is lengthened by the duration of the FL superframe preamble to ensure superframe timing alignment between the forward link and reverse link. Table 6-2 shows the specific durations of PHYFrames and superframes in the system. Note that these durations are a function of the cyclic prefix length which is a flexible parameter, and the number given in the table is for a cyclic prefix duration of 6.51us. This nominal value will also be used while quoting other numbers related to the frame structure, such as the retransmission interval.

Table 6-2 Superframe numerology

Parameter	Value	Unit
PHY Frame duration	8	OFDM Symbols
	911.46	us
Superframe Preamble Duration	8	OFDM Symbols
	1.07	ms
Superframe duration	24	PHY Frames (excluding Superframe Preamble)
	22.94	ms
Number of HARQ interlaces (FL & RL)	6	
Retransmission interval (FL & RL)	6	PHY Frames
	5.47	ms

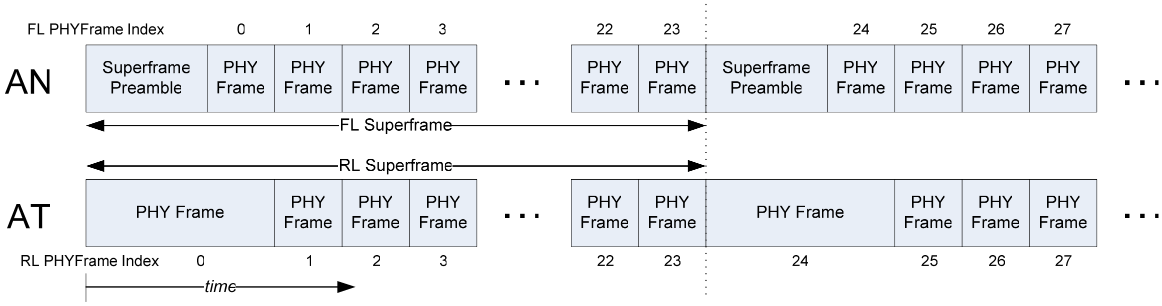

Figure 6-3 Superframe structure

6.3.1 H-ARQ Interlace Structure

Both forward and reverse link data transmissions support H-ARQ. To provide H-ARQ related processing time at the AP and AT, we use a six interlace structure for both FL and RL. Timing of transmissions associated with one of the six interlaces is shown for FL in Figure 6-4 and for RL in Figure 6-5, and the timing of the other interlaces is the same but with all transmissions shifted by the same number of PHYFrames. This interlace structure ignores the presence of the superframe preamble, i.e. PHYFrame level transmission timing occurs as if the superframe preamble were not present on FL and as if the first PHYFrame were not lengthened on RL.

For the forward link, assignments that arrive in FL PHYFrame *k* apply to the interlace containing FL PHYFrame *k*, and a FL transmission on FL PHYFrame *k* is acknowledged on RL PHYFrame *k+3*. HARQ retransmissions associated with the transmission that starts in PHYFrame *k* occur in PHYFrames *k+6n* where *n* is the retransmission index, *n*=1,2, … .

This frame structure provides an H-ARQ retransmission latency of ~5.5 ms with 1.8 ms (2 PHYFrames) of processing time at both the AT and the AP.

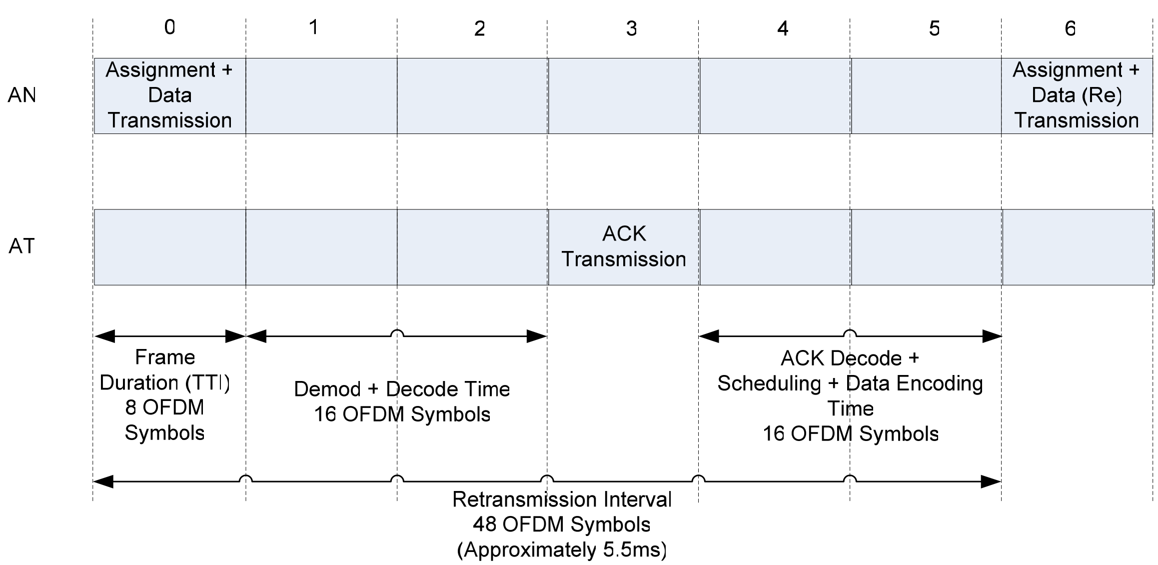

Figure 6-4 FL H-ARQ transmission timing

For the reverse link, assignments that arrive in FL PHYFrame *k* apply to the interlace containing RL PHYFrame *k+2*, and a RL transmission on RL PHYFrame *k* is acknowledged on FL PHYFrame *k+4*. HARQ retransmissions associated with a transmission that starts in PHYFrame *k* occur in PHYFrames *k+6n* where *n* is the retransmission index, *n*=1,2, … .

This frame structure provides an H-ARQ retransmission latency of 5.5 ms with 0.9 ms (1 PHYFrame) of processing time at the AT, and 2.7 ms (3 PHYFrames) of processing time at the AP. The reduced processing time at the AT is appropriate for RL since the AT only needs to perform assignment demodulation and data packet encoding/modulation – tasks that are much simpler than data packet demodulation.

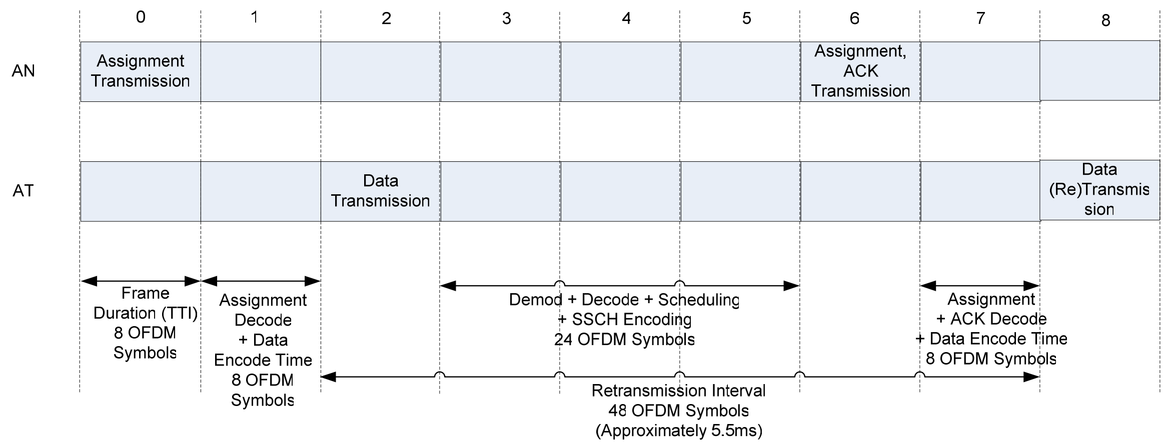

Figure 6-5 RL H-ARQ transmission timing

6.3.2 H-ARQ Interlace Structure for Extended Transmission Duration Assignments

In addition to the standard H-ARQ interlace structure, the system supports the use of "Extended Transmission Duration" assignments. Such assignments extend transmission over multiple PHYFrames and alter the timing of transmissions and corresponding ACK transmissions relative to the standard assignments. Such assignments are useful for link-budget limited users who can greatly benefit from encoding transmissions over a longer transmission duration. Extended Transmission Duration assignments create a potential for resource assignment collisions with standard assignments, and the AN should manage resource assignments to prevent such collisions.

For the forward link, Extended Transmission Duration assignments that arrive in FL PHYFrame *k* apply to the interlace containing FL PHYFrames *k* through *k+5*. A FL transmission on FL PHYFrames *k* through *k+5* is acknowledged on RL PHYFrame *k+8* through *k+9*. HARQ retransmissions associated with the transmission that starts in PHYFrame *k* start in PHYFrames *k+12n* where *n* is the retransmission index, *n*=1,2, … .

This frame structure provides an H-ARQ retransmission latency of 11 ms with 1.8 ms (2 PHYFrames) of processing time at both the AT and the AP.

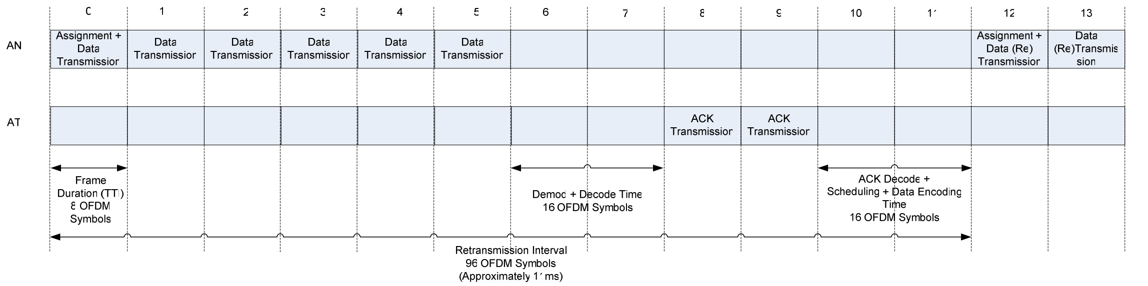

Figure 6-6 FL H-ARQ retransmission timing for Extended Transmission Duration

For the reverse link, Extended Transmission Duration assignments that arrive in FL PHYFrame *k* apply to the interlace containing RL PHYFrames *k+3* through *k+8*. A RL transmission on RL PHYFrames *k+3* through *k+8* is acknowledged on RL PHYFrame *k+12*. HARQ retransmissions associated with the transmission that starts in PHYFrame *k* start in PHYFrames *k+12n* where *n* is the retransmission index, *n*=1,2, … .

This frame structure provides an H-ARQ retransmission latency of 11 ms with 1.8 ms (2 PHYFrames) of processing time at the AT and 2.7 ms (3 PHYFrames) of processing time at the AP.

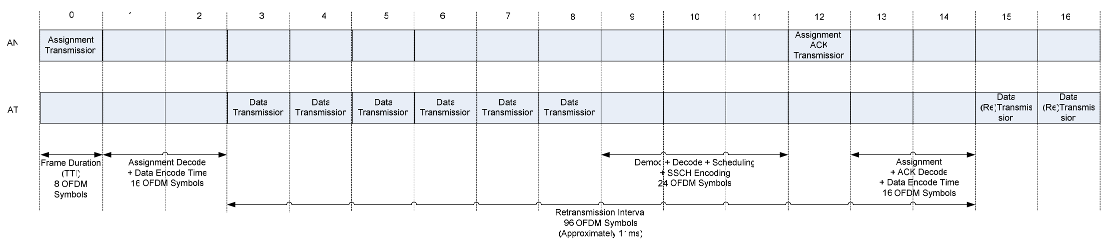

Figure 6-7 RL H-ARQ retransmission timing for Extended Transmission Duration

6.4 Periodicity of Control Channels

6.4.1 Forward Link

The forward link control channel is referred to as the Shared Signaling Channel (SSCH) and contains resource assignment messages, power control commands, and acknowledgements for RL H-ARQ. The SSCH is sent in each FL PHYFrame. The timing of resource assignments and acknowledgements follow from the frame structure described above.

The power control commands are sent using a portion of the SSCH resources and have the same periodicity as the RL Control Segment (described below) for a nominal periodicity of 5.5 ms and resulting power control rate of approximately 180 Hz..

6.4.2 Reverse Link

The reverse link control channels consist of feedback channels including CQICH for indicating channel quality and supporting handoff, SFCH for supporting sub-band scheduling and BFCH for supporting beam precoding/SDMA. Also included is REQCH to request resource allocation, a broadband pilot channel (PICH), and an acknowledgement channel (ACKCH) to support FL H-ARQ.

The ACKCH periodicity follows from the FL H-ARQ timing as described in the frame structure section above.

The remaining RL control channels (CQICH, BFCH, SFCH, REQCH and PICH) are transmitted in what is referred to as the Control Segment – a resource on the RL that is confined to a single PHYFrame. The Control Segment nominally is available once every 6 PHYFrames for a periodicity of 5.5 ms. However, the network can control the periodicity of transmissions on the Control Segment on a per-AT basis to manage the overall load on the Control Segments.

7 TDD Frame Structure

7.1 Channel Structure

The MBWA system includes several physical layer channels on the forward and reverse links, and these channels are configured and controlled by MAC layer protocols, as illustrated in Figure 7-1 and Figure 7-2. The following is a brief description of each Physical Layer Channel.

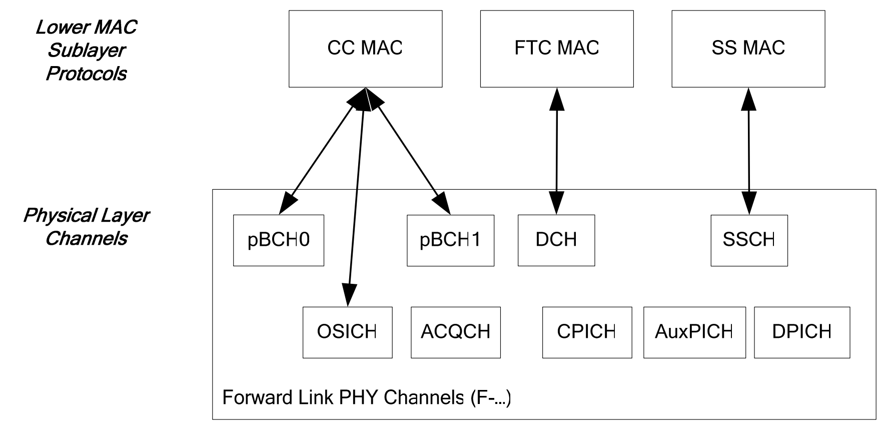

Figure 7-1 Forward channel structure

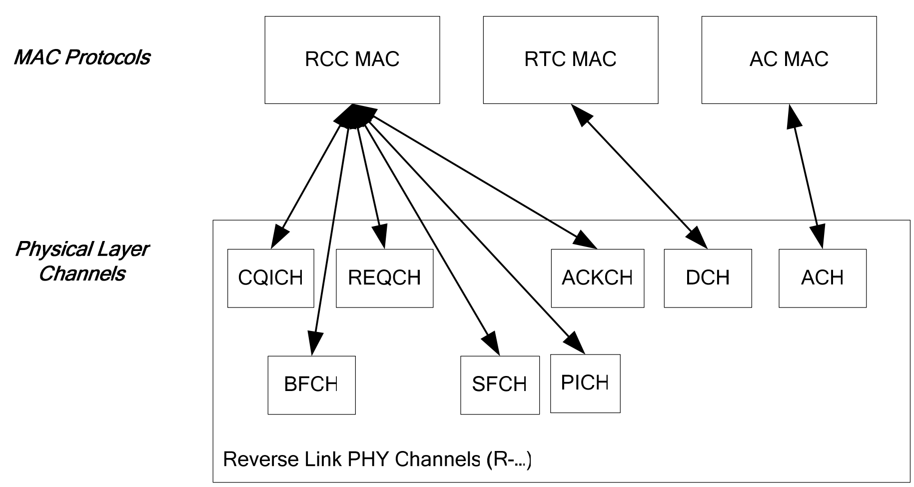

Figure 7-2 Reverse channel structure

7.1.1 Forward physical channels

Forward Acquisition Channel (F-ACQCH)
: Carries an acquisition pilot for an access terminal to use to acquire the system.

Forward Auxiliary Pilot Channel (F-AuxPICH)
: Carries auxiliary pilots for channel estimation from multiple transmit antennas. The Forward Primary Broadcast Channel 1 (F-pBCH1) indicates whether the F-AuxPICH is present.

Forward Common Pilot Channel (F-CPICH)
: Carries the common pilot.

Forward Data Channel (F-DCH)
: Carries information for a specific access terminal. A Forward Data Channel assignment is assigned to an access terminal by a Forward Shared Signaling Channel (F-SSCH) assignment. Also carries broadcast information including pages and sector specific messages.

Forward Dedicated Pilot Channel (F-DPICH)
: Carries the dedicated pilot. This channel is present in BlockHopping mode, which is indicated over the Forward Primary Broadcast Channel 0 (F-pBCH0).

Forward Other Sector Interference Pilot Channel (F-OSICH)
: Carries information about the interference from other sectors to be received by all access terminals.

Forward Primary Broadcast Channel 0 (F-pBCH0)
: Carries information about the system to be received by all access terminals.

Forward Primary Broadcast Channel 1 (F-pBCH1)
: Carries information about the sector to be received by all access terminals. Also carries quick pages.

Forward Shared Signaling Channel (F-SSCH)
: Carries forward and reverse link data channel assignments, access grants, power control commands, and acknowledgement information for Reverse Data Channel (R-DCH) transmissions.

7.1.2 Reverse physical channels

Reverse Access Channel (R-ACH)
: Used by access terminals to initiate communication with the access network. The Reverse Access Channel is also used by access terminals to obtain timing corrections.

Reverse Acknowledgement Channel (R-ACKCH)
: Carries acknowledgement information of a Forward Data Channel (F-DCH) reception.

Reverse Beam Feedback Channel (R-BFCH)
Carries information about the beam index and the quality of the forward link channel.

Reverse Channel Quality Indicator Channel (R-CQICH)
Carries information about the quality of the forward link channel of a sector as received by an access terminal. The Reverse Channel Quality Indicator Channel also carries information about the desired forward link serving sector.

Reverse Data Channel (R-DCH)
Carries information from an access terminal. The Reverse Data Channel is assigned to an access terminal by a Forward Shared Signaling Channel (F-SSCH) assignment.

Reverse Pilot Channel (R-PICH)
Carries the pilot.

Reverse Request Channel (R-REQCH)
Carries information about the buffer level at different quality of service classes for an access terminal. The Reverse Request Channel also carries information about the desired reverse link serving sector.

Reverse Subband Feedback Channel (R-SFCH)
Carries information about the quality of a subband of the forward link channel.

7.2 Basic Numerology

Transmission on both links of the MBWA system is divided into units of OFDM symbols. Three basic FFT sizes are defined, namely 512, 1024 and 2048. Each FFT size corresponds to a different chip rate, as shown in Table 7-1. These chip rates are chosen so that all three FFT sizes have the same subcarrier spacing, namely 9.6KHz.

The MBWA system supports flexible bandwidth operation. Different operating bandwidths are constructed from the three basic chip rates by using different numbers of guard carriers.

The values of these parameters, along with others like the cyclic prefix duration, the OFDM symbol duration, and the guard times for FL to RL and RL to FL transitions are specified in Table 7-1. Note that the OFDM symbol duration is independent of the bandwidth of operation.

Table 7-1 Basic Numerology

Parameter	512 pt FFT	1024 pt FFT	2048 pt FFT	Unit
Chip rate	4.9152	9.8304	19.6608	Mcps
Subcarrier spacing	9.6	9.6	9.6	kHz
Bandwidth of operation	≤ 5	≤ 10	≤ 20	MHz
Guard carriers	Function of bandwidth	Function of bandwidth	Function of bandwidth	
Cyclic prefix	6.51-26.04	6.51-26.04	6.51-26.04	us
Windowing duration	3.26	3.26	3.26	us

Parameter	512 pt FFT	1024 pt FFT	2048 pt FFT	Unit
OFDM Symbol duration (for 6.51 us CP)	113.93	113.93	113.93	us
Guard time (FL→RL)	78.12	78.12	78.12	us
Guard time (RL→FL)	16.28	16.28	16.28	us

7.3 Frame Structure

Forward and reverse link transmissions are divided into units of superframes. Superframes are further divided into units of PHYFrames. The frame structure of the MBWA system defines the timing of FL and RL PHYFrames within a superframe. Additionally, it defines the relative timing of assignments, acknowledgements, and H-ARQ retransmissions associated with a data packet. This structure is designed to minimize latency of data transmissions while maintaining acceptable processing durations for encoding and decoding at the AT and the AP, as well as scheduling at the AP.

The system supports general M:N TDD partitioning where M:N is the ratio of forward to reverse link transmission duration. For simplicity, we first describe the 1:1 and 2:1 TDD partitioning. The generalized structure for other M:N partitionings is then described in the following section.

Table 7-2 shows the specific durations of PHYFrames and superframes in the system for the 1:1 and 2:1 partitionings. Note that these durations are a function of the cyclic prefix length which is a flexible parameter, and the number given in the table is for a cyclic prefix duration of 6.5us. This nominal value will also be used while quoting other numbers related to the frame structure, such as the retransmission interval.

Table 7-2 Superframe numerology

Parameter	TDD 1:1	TDD 2:1	Unit
PHY Frame duration	8	8	OFDM Symbols
	911.46	911.46	us
Superframe Preamble Duration	8	8	OFDM Symbols
	1.07	1.07	ms
Superframe duration	24	24	PHY Frames (excluding superframe preamble)
	24.08	23.70	ms
Number of HARQ interlaces (FL/RL)	3/3	5/2	
Retransmission interval (FL/RL)	6/6	(7 or 8)/6	PHY Frames
	5.75/5.75	(6.57 or 7.57)/5.66	ms

7.3.1 1:1 TDD partitioning

As illustrated in Figure 7-3, a forward link superframe consists of a superframe preamble followed by 12 FL PHYFrames, and a reverse link superframe consists of 12 RL PHYFrames. The superframe preamble carries acquisition sequences plus key overhead parameters that enable an AT to receive the forward link control channels and subsequently access the system.

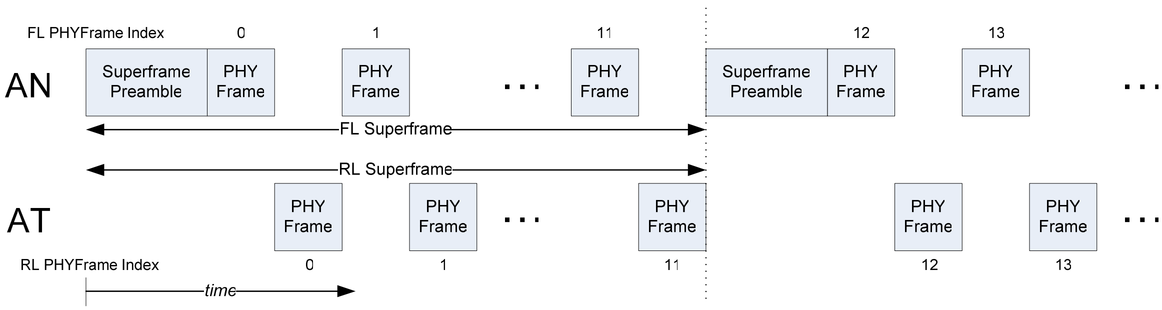

Figure 7-3 Superframe structure for 1:1 TDD partitioning

7.3.1.1 H-ARQ Interlace Structure for 1:1 TDD partitioning

Both forward and reverse link data transmissions support H-ARQ. To provide H-ARQ related processing time at the AP and AT, we use a three interlace structure for both FL and RL. Timing of transmissions associated with one of the interlaces is shown for FL in Figure 7-4 and for RL in Figure 7-5, and the timing of the other interlaces is the same but with all transmissions shifted by the same number of PHYFrames. This interlace structure ignores the presence of the superframe preamble, i.e. PHYFrame level transmission timing occurs as if the superframe preamble were not present.

For the forward link, assignments that arrive in FL PHYFrame k apply to the interlace containing FL PHYFrame k, and a FL transmission on FL PHYFrame k is acknowledged on RL PHYFrame $k+1$. HARQ retransmissions associated with the transmission that starts in PHYFrame k occur in PHYFrames $k+3n$ where n is the retransmission index, n=1,2, … .

This frame structure provides an H-ARQ retransmission latency of ~5.5 ms with 1.8 ms (2 PHYFrames) of processing time at both the AT and the AP.

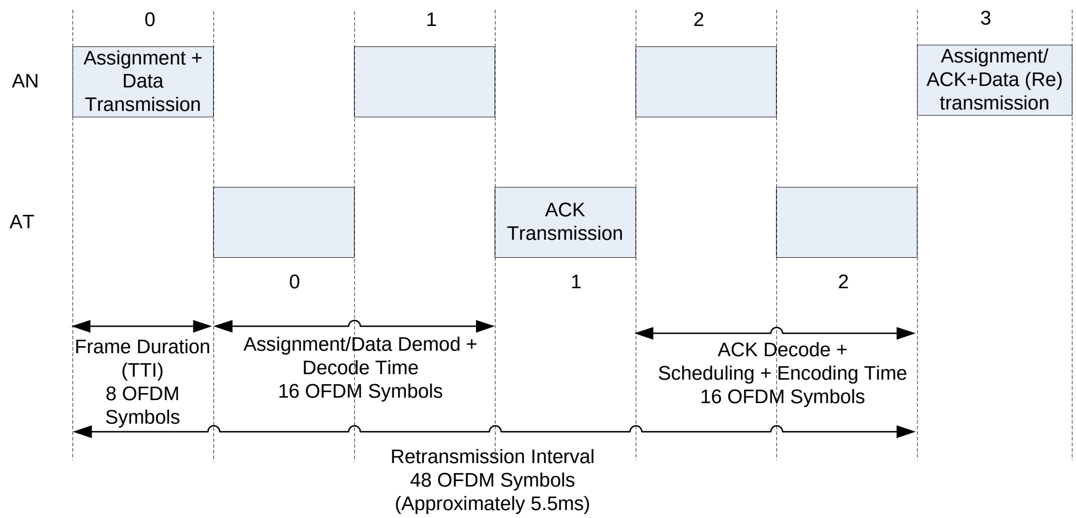

Figure 7-4 FL H-ARQ transmission timing for 1:1 TDD partitioning

For the reverse link, assignments that arrive in FL PHYFrame *k-1* apply to the interlace containing RL PHYFrame *k*, and a RL transmission on RL PHYFrame *k* is acknowledged on FL PHYFrame *k+2*. HARQ retransmissions associated with a transmission that starts in PHYFrame *k* occur in PHYFrames *k+3n* where *n* is the retransmission index, *n*=1,2, … .

This frame structure provides an H-ARQ retransmission latency of 5.5 ms with 1.8 ms (2 PHYFrames) of processing time at both the AT and the AP.

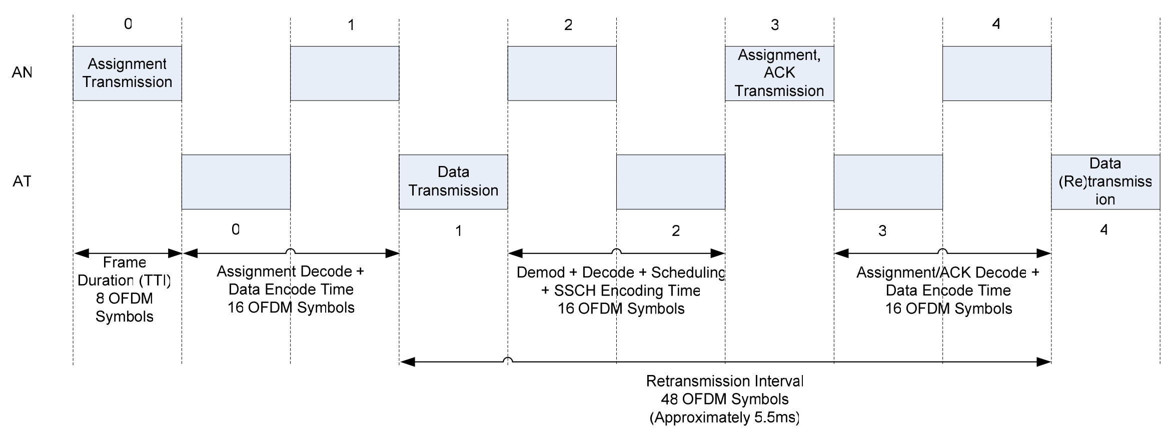

Figure 7-5 RL H-ARQ transmission timing for TDD 1:1 partitioning

7.3.2 2:1 TDD partitioning

As illustrated in Figure 7-6, a forward link superframe consists of a superframe preamble followed by 16 FL PHYFrames, and a reverse link superframe consists of 8 RL PHYFrames. The superframe preamble carries acquisition sequences plus key overhead parameters that enable an AT to receive the forward link control channels and subsequently access the system.

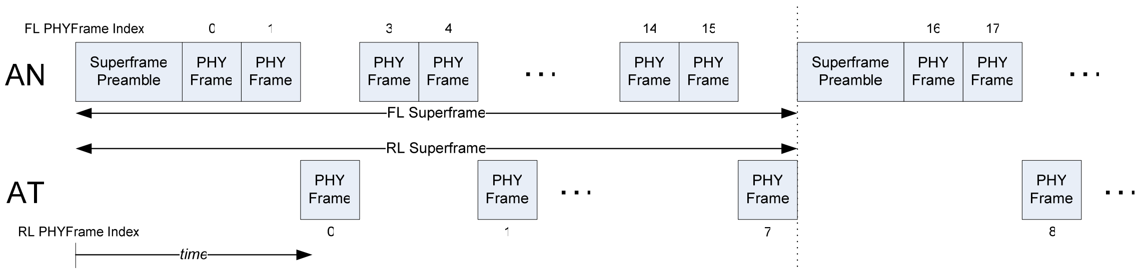

Figure 7-6 Superframe structure for TDD 2:1 partitioning

7.3.2.1 H-ARQ Interlace Structure for 2:1 TDD partitioning

Both forward and reverse link data transmissions support H-ARQ. To provide H-ARQ related processing time at the AP and AT, we use a five interlace structure for the FL and a two interlace structure for RL. Timing of transmissions associated with one of the three interlaces is shown for FL in Figure 7-7 and for RL in Figure 7-8, and the timing of the other interlaces is the same but with all transmissions shifted by the same number of PHYFrames. This interlace structure ignores the presence of the superframe preamble, i.e. PHYFrame level transmission timing occurs as if the superframe preamble were not present.

For the forward link, assignments that arrive in FL PHYFrame k apply to the interlace containing FL PHYFrame k. A FL transmission of a MAC packet FL PHYFrame k is acknowledged in RL PHYFrame $\lfloor k/2 \rfloor +1$. HARQ retransmissions associated with the transmission that starts in PHYFrame k occur in PHYFrames $k+5n$ where n is the retransmission index, n=1,2, … .

This frame structure provides an H-ARQ retransmission latency of 6.4 ms 50% of the time and 7.3 ms 50% of the time. Also, this structure provides processing time of 2.7 ms 50% of the time and 3.6 ms 50% of the time at the AT, and processing time of 0.9 ms 50% of the time and 2.7 ms 50% of the time at the AP.

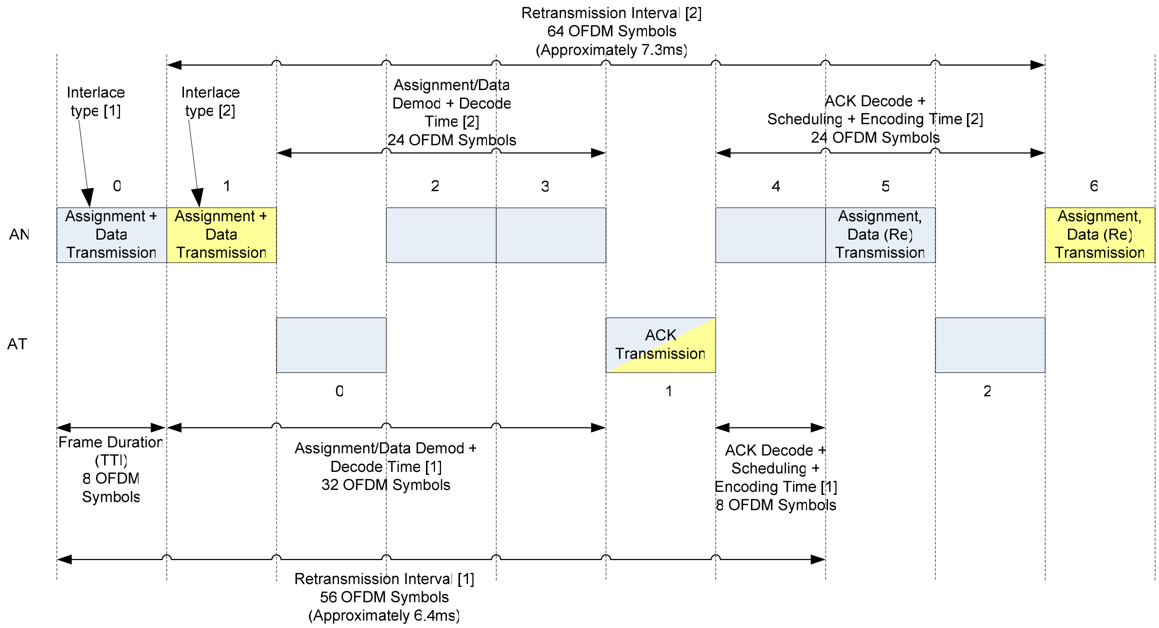

Figure 7-7 FL H-ARQ transmission timing

For the reverse link, assignments that arrive in FL PHYFrame *k* apply to the interlace containing the RL PHYFrame $\lfloor k/2 \rfloor$, and a RL transmission on RL PHYFrame *j* is acknowledged in FL PHYFrame $2j+4$. HARQ retransmissions associated with a transmission that starts in PHYFrame *k* occur in PHYFrames $k+2n$ where *n* is the retransmission index, n=1,2, … .

This frame structure provides an H-ARQ retransmission latency of ~5.5 ms with 0.9 ms (1 PHYFrame) of processing time at the AT and 2.7 ms (3 PHYFrames) of processing time at the AP.

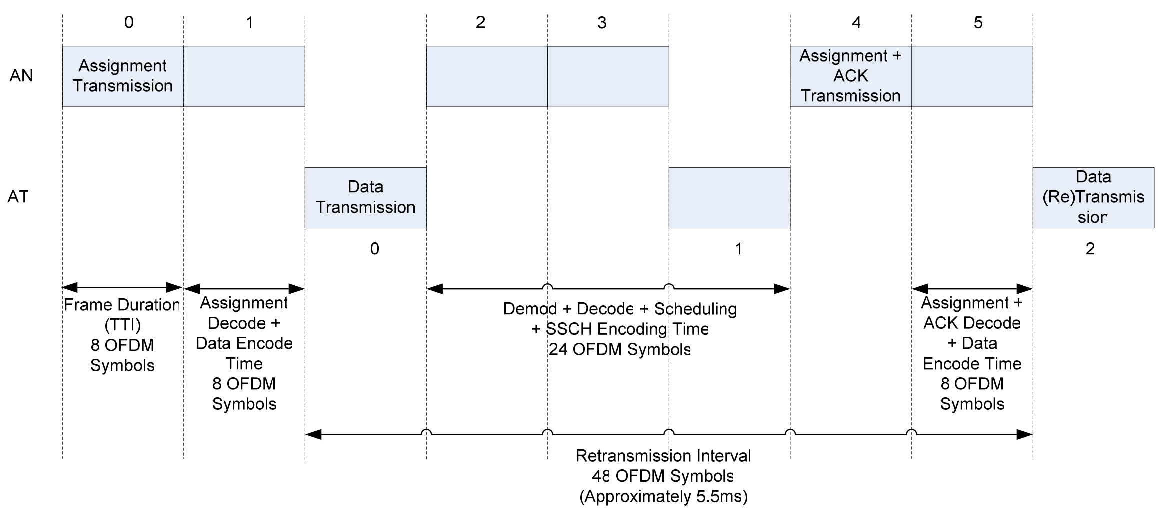

Figure 7-8 RL H-ARQ transmission timing for 2:1 TDD partitioning

7.3.3 Generalized Frame structure for M:N TDD partitioning

In addition to the 1:1 and 2:1 TDD partitioning structure described above, the system supports TDD operation with *M* FL PHYFrames followed by *N* RL PHYFrames, or an *M:N* partitioning. The grouping of consecutively transmitted PHYFrames is referred to as a "burst." This structure is illustrated in Figure 7-9. Note that this structure reduces to the TDD 1:1 structure described earlier for the case M=N=1.

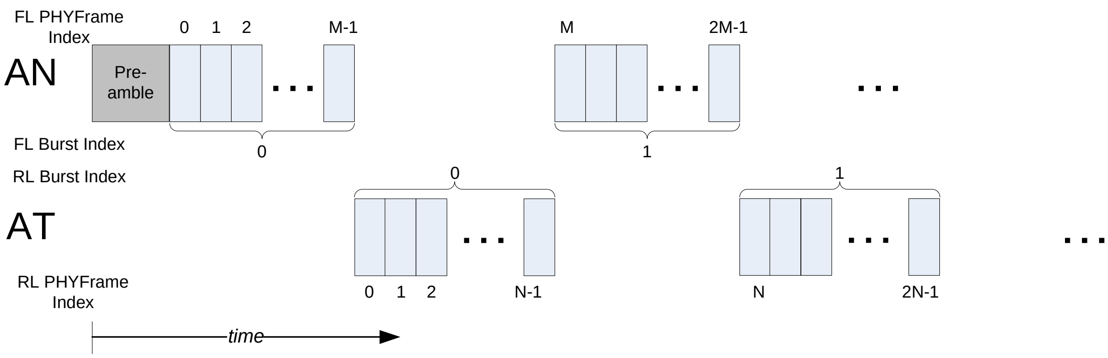

Figure 7-9 Frame structure for M:N TDD partitioning

Assignment, acknowledgement, and retransmission timing can then be specified using the following formulas, where *i* is the FL Burst Index, and *j* is the RL Burst index.

Assignments:

FL PHYFrames in *i* are assigned in FL burst *i*.
RL PHYFrames in *j* are assigned in FL burst *i=j* if M>1 or *i=j-1* if M=1.

Acknowledgement:

FL PHYFrames in *i* are acknowledged in RL burst *j=i+1*.
RL PHYFrames in *j* are acknowledged in FL burst *i=j+2*.

Retransmissions:

FL PHYFrames in *i* are retransmitted in FL burst *i+2*, if N>1 or *i+3* if N=1.
RL PHYFrames in *j* are retransmitted in RL burst *j+2*, if M >1 or *j+3* if M=1.

Resulting number of interlaces:

2*M FL interlaces if N >1 or 3*M FL interlaces if N=1.
2*N RL interlaces if M>1 or 3*N RL interlaces if M=1.

7.4 Periodicity of Control Channels

7.4.1 Forward Link

The forward link control channel is referred to as the Shared Signaling Channel (SSCH) and contains resource assignment messages, power control commands, and acknowledgements for RL H-ARQ. The SSCH is sent in each FL PHYFrame. The timing of resource assignments and acknowledgements follow from the frame structure described above.

The power control commands are sent using a portion of the SSCH resources and have the same periodicity as the RL Control Segment (described below) for a nominal periodicity of 5.5 ms and resulting power control rate of approximately 180 Hz for the 1:1 TDD partitioning.

7.4.2 Reverse Link

The reverse link control channels consist of feedback channels including CQICH for indicating channel quality and supporting handoff, SFCH for supporting sub-band scheduling and BFCH for supporting beam precoding/SDMA. Also included is REQCH to request resource allocation, a broadband pilot channel (PICH), and an acknowledgement channel (ACKCH) to support FL H-ARQ.

The ACKCH periodicity follows from the FL H-ARQ timing as described in the frame structure section above.

The remaining RL control channels (CQICH, BFCH, SFCH, REQCH and PICH) are transmitted in what is referred to as the Control Segment – a resource on the RL that is confined to a single PHYFrame. For 1:1 TDD partitioning, the Control Segment nominally is available once every 6 PHYFrames for a periodicity of 5.5 ms. However, the network can control the periodicity of transmissions on the Control Segment on a per-AT basis to manage the overall load on the Control Segments.

8 Coding and Modulation

8.1 Introduction

This document describes various aspects of the coding and modulation scheme used in the proposed system. The overall encoding and modulation structure is shown in Figure 8-1.

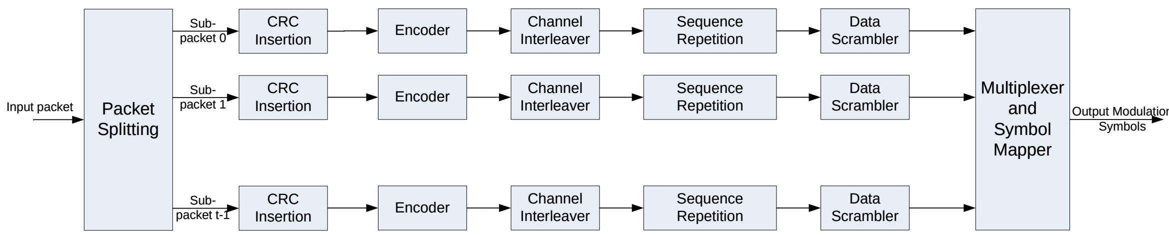

Figure 8-1 Coding and Modulation Structure

8.2 Channel Coding

Packet Splitting

Large packets are split into smaller subpackets and separately encoded. This permits different turbo decoder units to operate on the different subpackets in parallel and thus speeds up decoding. Furthermore, each subpacket has its own CRC in order to allow for early termination of decoding.

All subpackets have approximately the same size. The maximum possible size of a subpacket is 8192 bits.

Convolutional Code

A rate-1/3 convolutional code is used to encode short packets in which the number of information bits is less than or equal to 128. Such packets are primarily used in control channels. The generator polynomial of this code is:

$$G(D) = [g_0(D)\ g_1(D)\ g_2(D)\ ]$$

where $g_0(D)= 1 + D^2 + D^3 + D^5 + D^6 + D^7 + D^8$, $g_1(D)= 1 + D + D^3 + D^4 + D^7 + D^8$, and $g_2(D)= 1 + D + D^2 + D^5 + D^8$, where D represents the delay operator.

Turbo Code

A rate-1/5 turbo code is used to encode packets (or subpackets) in which the number of information bits is greater than 128. The turbo code is a parallel concatenation of two constituent systematic recursive convolutional codes with the turbo interleaver preceding the second recursive convolutional encoder. The transfer function for either constituent code is:

$$G(D) = \begin{bmatrix} 1 & \frac{n_0(D)}{d(D)} & \frac{n_1(D)}{d(D)} \end{bmatrix}$$

where $d(D) = 1 + D^2 + D^3$, $n_0(D) = 1 + D + D^3$, and $n_1(D) = 1 + D + D^2 + D^3$, where D represents the delay operator. The turbo interleaver uses a combination of bit-reversal interleaving, row-column interleaving and linear congruential interleaving. The interleaver is flexible and allows for a wide range of packet sizes.

Channel Interleaving

The proposed channel interleaver design is based on a pruned bit reversal interleaver (PBRI). An m-bit reversal interleaver operates on a sequence of length $N = 2^m$ by mapping each index to its bit reversed value. An m-bit reversal interleaver π satisfies the following distance property: If $|a-b| < 2^i$, then $|\pi(a) - \pi(b)| \geq 2^{m-i-1}$. This distance property is desirable since it provides for robustness against bursty errors.

Moreover, it is easy to see that after interleaving, the first N/2 bits are those with indices divisible by 2, the first N/4 bits have indices divisible by 4 and so on. Therefore puncturing all but the first $N/2^i$ bits would result in transmission of bits spaced 2^i apart. Thus we can achieve regular puncture patterns when the "puncturing factor" is a power of 2. Regular puncture patterns are important because the performance of turbo and convolutional codes under regular puncturing is better than the performance under irregular puncturing.

A pruned bit reversal interleaver (PBRI) operates of sequences of length N when N is not a power of 2 by appending dummy symbols, using a bit-reversal interleaver and then deleting the dummy symbols. A PBRI has the same desirable properties as the bit reversal interleaver.

The channel interleaver used with the convolutional code uses a PBRI on all the encoded bits. The channel interleaver used with the turbo encoder splits the turbo-encoded bits into three groups as shown below and then uses a PBRI to interleave each of these groups. The systematic bits are then transmitted first, followed by one of the groups of non-systematic bits and then the other non-systematic group. In addition to providing good distance and regular puncture patterns, this channel interleaver also has the property that the code reduces to a rate-1/3 turbo code when the transmitted bits are punctured down to that rate.

Incremental Redundancy by puncturing and repetition

Hybrid-ARQ in this air-link uses incremental redundancy. A certain fraction of the interleaved bits are transmitted in each HARQ transmission. The code rate seen by the receiver goes down as the number of HARQ transmissions increases. Code rates below 1/5 are achieved by repetition i.e., transmitting some of the encoded bits more than once. This is done only after all the encoded bits are transmitted at least once.

Data Scrambling

The encoded bits are XORed with a pseudorandom sequence prior to modulation. The pseudorandom sequence is generated using a PN register, which is seeded with the MAC ID of the user. MAC ID based scrambling prevents a user from accidentally decoding a packet intended for another user.

8.3 Modulation

Modulation Formats and Modulation Step-down

Four modulation formats (QPSK, 8PSK, 16QAM and 64QAM) are supported in the air-link. The signal constellation for each of these modulation formats is based on the Gray mapping. The use of multiple modulation formats and coding rates allows for a range of spectral efficiencies (also referred to as packet formats). The system supports 15 packet formats on the forward link (in SISO mode) as well as on the reverse link. The number of bits in a physical layer transmission is determined by the number of subcarriers assigned for the transmission, and the packet format chosen.

The system uses synchronous, non-adaptive H-ARQ, i.e. the channel assignments and hence spectral efficiencies are the same for all H-ARQ transmissions of a PHY packet. Some packet formats use lower order modulation formats for later H-ARQ retransmissions, a technique known as "Modulation step-down." Since the code rate is higher with a lower-order modulation format (for the same spectral efficiency), modulation step-down helps in preventing the code rate from falling below 1/5. The disadvantage of having a code rate lower than 1/5 is the need to repeat bits, which is suboptimal from an information theoretic point of view. Modulation step-down thus prevents (or at least minimizes) repetitions and thereby improves code performance.

Modulation step-down is supported on both forward and reverse links, as can be seen from the following tables of packet formats used on the FL (SISO mode) and RL respectively. Note that some of the packet formats cannot be decoded after the first transmission because the effective code rate at that point is greater than 1.

Table 8-1 FL packet formats – SISO mode

Packet Format Index	Spectral efficiency on 1st trans-mission	Max number of trans-missions	Modulation order for each transmission					
			1	2	3	4	5	6
0	0.2	6	2	2	2	2	2	2
1	0.5	6	2	2	2	2	2	2
2	1.0	6	2	2	2	2	2	2
3	1.5	6	3	2	2	2	2	2
4	2.0	6	4	3	3	3	3	3
5	2.5	6	6	4	4	4	4	4
6	3.0	6	6	4	4	4	4	4
7	4.0	6	6	6	4	4	4	4
8	5.0	6	6	6	4	4	4	4
9	6.0	6	6	6	4	4	4	4
10	7.0	6	6	6	4	4	4	4
11	8.0	6	6	6	6	4	4	4
12	9.0	6	6	6	6	4	4	4
13	10.0	6	6	6	6	6	4	4
14	11.0	6	6	6	6	6	4	4
15	NULL							

Table 8-2 RL packet formats

Packet format index	Spectral efficiency on 1st transmission	Max number of transmissions	Modulation order for each transmission					
			1	2	3	4	5	6
0	0.25	6	2	2	2	2	2	2
1	0.50	6	2	2	2	2	2	2
2	1.0	6	2	2	2	2	2	2
3	1.5	6	3	2	2	2	2	2
4	2.0	6	3	3	2	2	2	2
5	2.67	6	4	4	3	3	3	3
6	4.0	6	4	4	3	3	3	3
7	6.0	6	4	4	4	3	3	3
8	8.0	6	4	4	4	4	4	3
9	4.0	6	6	6	4	4	4	4
10	5.0	6	6	6	4	4	4	4
11	6.0	6	6	6	4	4	4	4
12	7.0	6	6	6	4	4	4	4
13	8.0	6	6	6	6	4	4	4
14	9.0	6	6	6	6	4	4	4

9 Resource Management

This section discusses how resources are allocated to the forward and reverse link data channels.

9.1 Scheduling

Scheduling refers to the allocation of subcarriers and spectral efficiency[2] to ATs over time and is centralized in the access point for both FL and RL allocations. Centralized scheduling is utilized to ensure orthogonal allocations of resources to different ATs in the system.

The goals of the scheduler are to maximize system capacity while managing QoS requirements such as latency and throughput requirements of ATs. Additionally, the scheduler manages fairness across ATs that can have widely disparate link qualities and thus can support different instantaneous spectral efficiencies.

While the details of scheduler implementation are outside the scope of the air interface specification, the design ensures that the air interface provides tools for the scheduler to utilize features such as sub-band scheduling (Section 19), fractional frequency reuse (Section 18), MIMO (Section 20), precoding (Section 21), quasi-orthogonal reverse link (Section 17), beamforming (Section 22), and SDMA (Section 23) to achieve the above goals. .

9.2 Assignment Management

The MBWA system has been designed to efficiently support a range of applications, with a corresponding range of required throughput, latency, and application packet sizes. For example, the system has been designed to support very high best-effort data capacity as well as supporting many users of voice-like applications. In addition the system has been designed to efficiently support a combination of disparate application types simultaneously over the air interface.

Resources in the system are assigned in units of hop-ports, where a hop-port is a static resource that has a mapping to a physical subcarrier. Frequency hopping is implemented by having time varying mappings from hop-ports to subcarriers.

To reduce assignment signaling overhead, the system uses “synchronous HARQ” and provides support for “sticky” assignments. With synchronous HARQ, the resources for successive retransmissions are not independently scheduled, but rather are retained for all retransmissions associated with a packet. Thus, the assignment of a set of hop-ports applies to an “interlace,” as described in Sections 6.3 and 7.3. Assignments on different interlaces are independent, and an AT may be given resources on multiple interlaces.

Assignments can be sticky or non-sticky. Sticky assignments are useful to reduce assignment overhead required when it is beneficial to schedule multiple users simultaneously, and to eliminate request latency for RL transmissions. When an assignment is non-sticky, the assignment expires on successful packet decode, or when the packet fails to decode after the maximum number of H-ARQ retransmissions allowed for the packet. When assignments are sticky, the assignment persists as long as the assigned resource is in use. An assignment is in use as long as either a packet or an erasure sequence is transmitted using the assignment. The erasure sequence is simply a one-bit “keep alive” indication used to inform the receiver that the assignment should be retained even though a data packet might not be available for transmission using the assignment. If neither a packet nor an erasure sequence is transmitted using the assignment, the

[2] Spectral efficiency is defined by coding and modulation packet formats as described in Section 8.

assignment expires and the resources are free for subsequent allocation. In addition, it is possible for the AP to send an explicit message that expires an assignment.

To reduce the overhead required to specify sets of hop ports in the system, a finite space of channel IDs are defined that map to specific sets of hop ports and that are used to communicate assignments to ATs. Because assignments can be sticky, and to combat fragmentation of resources in the system due to this finite mapping of channel IDs, the system supports supplemental assignments that add sets of hop ports to the existing set allocated to an AT for an interlace. Such supplemental assignments are sent to augment an AT's allocation between packet transmissions.

The mapping between channel IDs and hop-ports is defined using a channel tree, such as the one illustrated in Figure 9-1. Each node on the tree is given a unique channel ID. Further, each base node (nodes at the bottom of the tree) is mapped to a set of hop ports. A channel ID then maps to the set of hop ports mapped by the base nodes under the node of the channel ID.

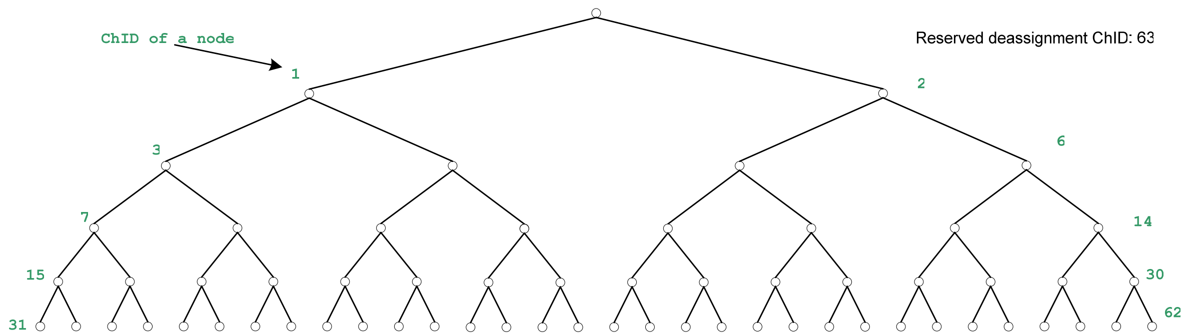

Figure 9-1 Example Channel Tree

FL and RL resource assignments are communicated to the ATs via "link assignment blocks" (LAB) transmitted over the SSCH forward link physical layer channel. For details on the SSCH structure see Section 13. A LAB contains the fields itemized in Table 9-1. Sticky and non-sticky assignments utilize different LAB types.

Table 9-1 Link Assignment Block (LAB) Fields

Parameter	Description
MACID	Unique AT identifier for a sector
Channel ID	Specifies a set of hop ports via the channel tree
Packet Format	Specifies the coding/modulation to be used for transmissions on the assigned resources.
Extended Transmission	Specifies whether the assignment is an extended transmission duration assignment (see Section 6.3.2)
Supplemental	Specifies whether the assignment should replace or supplement the existing assignment for the relevant interlace

The AT takes the following steps on reception of a LAB. If the MACID is equal to the AT's MACID, then the channel ID specified either replaces or supplements the ATs assignment for the relevant interlace depending on the "supplemental" field. If the LAB is received while a packet is in transmission on the relevant interlace, then the AT stops transmitting/receiving the current packet, updates the assignment for the interlace, and then starts with a new packet. If the channel ID specified is the reserved channel ID for deassignment, then the entire assignment for the relevant interlace is expired.

If the MACID is *not* equal to the AT's MACID, then the assignment is for another AT. In this case, if the channel ID specified conflicts with the AT's current assignment for the relevant interlace, then all conflicting hop ports are removed from the AT's assignment. This is referred to as a "decremental assignment."

10 Acquisition

10.1 Superframe preamble structure

As mentioned earlier, the forward link transmission is divided into units of superframes. Each superframe consists of a superframe preamble followed by a sequence of FL PHY Frames. The superframe preamble consists of 8 OFDM symbols, which are indexed 0 through 7. The structure of the preamble is illustrated in Figure 10-1.

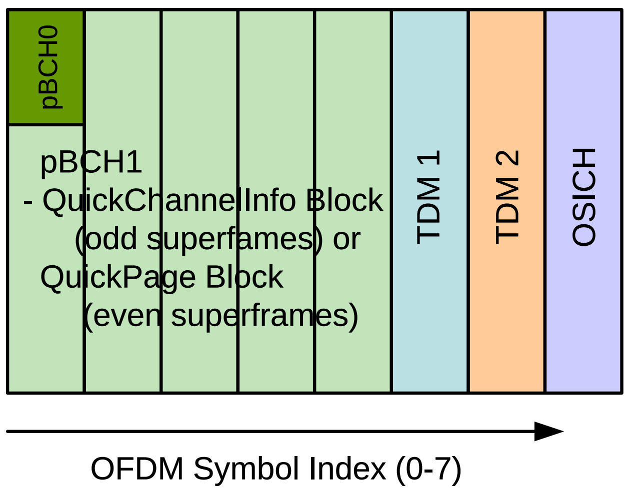

Figure 10-1 Structure of the Superframe Preamble

10.1.1 Acquisition Pilots

The last three OFDM symbols in the superframe preamble (the symbols indexed 5 through 7) are TDM pilots which are used for initial acquisition. These symbols will also be referred to TDM pilot 1, TDM pilot 2 and TDM pilot 3. The first two of these form the Acquisition Channel (F-ACQCH), while the last is reused in order to transmit the Other Sector Interference Channel (F-OSICH). The system supports a large number of unique PN codes (total number is 4096) in order to simplify PN planning requirements. Each sector is identified by a number between 0 and 4095 called the PilotPN, and the PN planning needs to be such that no two sectors with the same PilotPN are observable at any location.

In order to reduce acquisition complexity, while still maintaining a large number of PN codes, we use a hierarchical pilot structure. Therefore, TDM pilot 1 is scrambled using only 2 bits of information (i.e., takes on 4 possible values). TDM pilot 2 is scrambled using 8 bits of information (i.e., takes on 256 possible values), 2 of which are the same ones used to scramble TDM pilot 1. Finally, TDM pilot 3 is scrambled using 12 bits of information, which enables us to distinguish 4096 different PilotPN values. Moreover, TDM 1 is a perdiodic pilot (consisting of two periods), and can hence be detected using a low complexity delayed correlation. The delayed correlation can reduce the initial number of time-hypotheses dramatically, as well as aid in frequency synchronization.

The system supports two modes, namely Semi-synchronous and Asynchronous, and the TDM pilots are generated differently in the two cases.

10.1.1.1 Semi-synchronous Mode

For some applications, such as ranging, it is necessary to detect extremely weak sectors. For this reason it can be important to allow correlations over more than one superframe preamble. In order to get processing gains over more than one superframe preamble, it is necessary that the interfering signal (i.e., the TDM pilot of the neighboring sector) changes from one superframe preamble to the next. To enable this, we define an auxiliary quantity for each sector called the PilotPhase that changes from superframe to superframe. The PilotPhase is defined as PilotPhase = PilotPN + SuperframeIndex mod 4096, and is used to generate the TDM pilots. Here, SuperframeIndex is a counter that is incremented from one superframe to the next, and is defined globally across all sectors.

The PilotPhase is a 12 bit quantity. The 2 LSBs of PilotPhase are used to scramble TDM pilot 1, the 8 LSBs are used to scramble TDM pilot 2, while all 12 bits are used to scramble TDM pilot 3. This generates the hierarchical structure and enables a low-complexity search algorithm.

This pilot structure requires some level of synchronization between two sectors. To be more precise, if all possible values of PilotPN are possible, then this structure requires that any two sectors be synchronized to within half a superframe. Otherwise it is conceivable that two sectors with different PilotPNs will transmit the same acquisition pilots (same PN sequences) at the same time.

10.1.1.2 Asynchronous Mode

In some situations it is not possible to accurately synchronize two sectors. In order to support these scenarios, we have a mode in the system that has no synchronization requirements. In this mode, the TDM pilots are scrambled directly using the PilotPN instead of the auxiliary quantity PilotPhase. The two LSBs of the PilotPN are used to scramble TDM1, the 8 LSBs are used to scramble TDM2, while all 12 bits are used to scramble TDM3. Thus the hierarchical structure is still maintained, however the processing gain across different superframes is potentially lost.

10.1.2 Primary Broadcast Channels

The first five OFDM symbols in the superframe preamble are used to carry the two Primary Broadcast Channels, namely F-pBCH0 and F-pBCH1. These channels carry configuration information that the AT needs to have before it can demodulate the PHY Frames. In addition, the F-pBCH1 channel also carries paging information, which will be discussed in a separate section. A more detailed description of the different ways in which overhead information is transmitted will be provided in Section 10.2.

An F-pBCH0 packet is encoded over 16 superframes, and occupies ¼ of an OFDM symbol in each superframe preamble. This channel therefore has an extremely small overhead. An F-pBCH1 packet is encoded over a single superframe and occupies 4 ¾ OFDM symbols in each superframe preamble. The bandwidth overhead of this channel is approximately 2%.

10.1.3 Other Sector Interference Channel (F-OSICH)

This channel carries a three state quantity, that is modulated as a phase on TDM pilot 3. Since the TDM pilot waveform is known once acquisition is completed, the superposition causes no degradation to the performance of the OSICH. The function of this channel will be described in detail in the power control section, but one of the characteristics is that it is used by ATs in the neighboring sector, i.e., it should be decodable at extremely low SNRs. This is accomplished by providing an extremely large spreading gain for this channel, i.e., an entire OFDM symbols is used to transmit less than two bits of information.

10.2 Overhead Channel Structure

The overhead messages consist of the following blocks:

1. SystemInfo block: This block is transmitted over the F-pBCH0 channel in the superframe preamble. It contains information that is expected to be constant over a large group of sectors. The parameters it carries include the cyclic prefix duration, the number of guard carriers etc. In addition, it also carries the 12 LSBs of the superframe index. It is transmitted at a spectral efficiency of approximately 0.03. Since the F-pBCH0 is transmitted over 16 superframes, the latency in order to demodulate this channel can be as much as 0.4sec. However, in the SemiSynchronous mode, the F-pBCH0 needs to be demodulated only at initial wake-up time, since the parameters it carries are static over an entire deployment.

2. QuickChannelInfo block: This block is transmitted over the F-pBCH1 channel in the superframe preamble, in superframes with an odd superframe index. It contains configuration information that is required for two purposes:

   a. To enable an AT to demodulate other overhead channels that are transmitted in the PHY Frames.

   b. To enable an idle-mode AT to demodulate pages that may be transmitted in the PHY Frames. It is desirable that the AT be able to demodulate these pages even if it wakes up in a new sector, and hence this information needs to be transmitted with a high periodicity.

   Parameters contained in the QuickChannelInfo block include FL hopping parameters, configuration parameters for the SSCH etc.

3. ExtendedChannelInfo message: This message contains other configuration information related to FL and RL structure, which is not required in order to demodulate pages or in order to demodulate the ExtendedChannelInfo message itself. This includes parameters related to RL configuration as well as parameters related to transmission of FL power control bits. The ExtendedChannelInfo message consists of several groups like the PowerControl group, the SectorInformation group etc. This block is transmitted like a regular data channel in predefined superframes, using a broadcast MAC ID. The periodicity of this channel is 16 superframes or approximately 0.4sec.

4. SectorParameters message: This message carries information that the AT is not required to know in order to access the system. For example, it carries neighbor list information, which is used by the AT to optimize its handoff algorithm. It is transmitted as a regular data channel, similar to the ExtendedChannelInfo message. The periodicity of this message is 64 superframes, which is about 1.5sec.

10.3 AT wake-up procedure

On initial wake-up, the AT first detects a sector and achieves time and frequency synchronization using the TDM pilots. In SemiSynchronous mode, the AT knows the value of the PilotPhase variable at the end of this stage, while in Asynchronous mode, the AT knows the value of the PilotPN variable at the end of this stage. The AT then goes on to demodulate the F-pBCH0 and F-pBCH1 channels. The F-pBCH0 channel carries the lower 12 bits of the SuperframeIndex, which enables the AT to find the value of PilotPN in the SemiSynchronous case (PilotPN = PilotPhase – SuperframeIndex mod 4096). Therefore, in both SemiSynchronous as well as Asynchronous modes, at the end of this stage, the AT knows the

PilotPN and SuperframeIndex variables, which are together used to seed various random number generators (for hopping, scrambling etc) used in generating the FL waveform.

Moreover, the two F-pBCH channels provide the AT with sufficient configuration information to enable it to demodulate the remaining overhead channels. The AT can now read the ExtendedChannelInfo message as and when it is transmitted, at which point it knows the complete configuration of the FL and RL. At this stage, the AT is ready to access the system.

11 Access Channel Procedures

11.1 Introduction

This chapter describes the reverse link access channel (R-ACH) procedures used in the proposed system.

The access channel is used by the AT to gain initial access to the network. The access succeeds when the AT receives an acknowledgement sent by the AP upon successful detection of an access probe. The access probes are multiplexed with other reverse link control channels in a CDMA control channel segment as described in Section 14.

11.2 AT wake-up procedure

The AT wake-up procedure, as described in detail in Section 10.3 is briefly revisited here.

On initial wake-up, the AT first detects a sector and achieves time and frequency synchronization, finds the PilotPN and SuperFrameIndex variables, and the complete configuration of the FL and RL.

The AT accesses the system by sending successive access probes on the CDMA control segment. The access succeeds when the AT receives an AccessGrant message sent by the AP over the SSCH upon successful detection of an access probe. The AccessGrant message assigns to the AT a MACID and initial reverse link resources. The AccessGrant message also contains a timing adjustment command so that the AT can properly orthogonalize its transmission to all the other ATs.

11.3 Access Probe Structure

The Access Channel MAC Protocol transmits access probes by instructing the Physical Layer to transmit a probe. With the instruction, the Access Channel MAC Protocol provides the Physical Layer the power level, AccessSequenceID, PilotPN of the sector to which the access probe is to be transmitted, and other parameters.

An access probe is transmitted over the CDMA control segment and it is modulated as a 1024 Walsh sequence (given by the AccessSequenceID) scrambled with a PN sequence generated using the PilotPhase of the FL Serving Sector and the RL PHYFrame as seeds.

Some information is embedded in each access probe through the selection of the Walsh sequence. The space of the Walsh sequences is partitioned into a number of sets, each set corresponding to a possible value of the variable to be transmitted. For each access probe, the Walsh sequence is chosen randomly out of the sequences of the appropriate set. This method is used to send information about the received pilot level and to request different amount of resources for subsequent transmissions.

11.4 Access Probes Transmission Procedure

Figure 11-1 illustrates the grouping of access probes in access probe sequences. In the figure, Ns probe sequences are shown, where each probe sequence has Np probes.

The number of slots left empty between successive probes from the same access sequence is given by AccessCycleDuration. The term “slot” denotes a RL PHYFrame which contains the CDMA segment.

The probes of each sequence are transmitted at increased power. The ramping of power of successive access probes in one sequence is designed such that most of the time an AccessGrant is received before the complete transmission of the first access sequence (maximum Np probes). If after the transmission of one access sequence no AccessGrant is received, the AT starts the transmission of a new access sequence after waiting a random number of slots.

The number of slots before the retransmission of a new access sequence is determined in the same way as the number of slots before the transmission of the first access probe, as follows; we'll refer to this random number as persistence interval.

If the access attempt is made in response to a page and no other AT is paged simultaneously, then the transmission of each access sequence starts in the first available slot; no empty slot is left between successive access sequences (the persistence interval is 0). If multiple ATs are paged simultaneously, then the persistence interval is drawn from a uniform distribution on the integers between 0 and a maximum value (given by PageResponseBackoff*3).

For all access attempts that are not made in response to a page, the persistence interval is given by the value of a geometric random variable with parameter *p*. If this value exceeds a certain maximum, then the persistence interval is set to the maximum value.

The transmission of access probes stops (without completion of the current access sequence) if any of the following conditions are met:

- The access terminal receives an AccessGrant, or
- Transmission is aborted because the MAC protocol at AT received a *Deactivate* command.

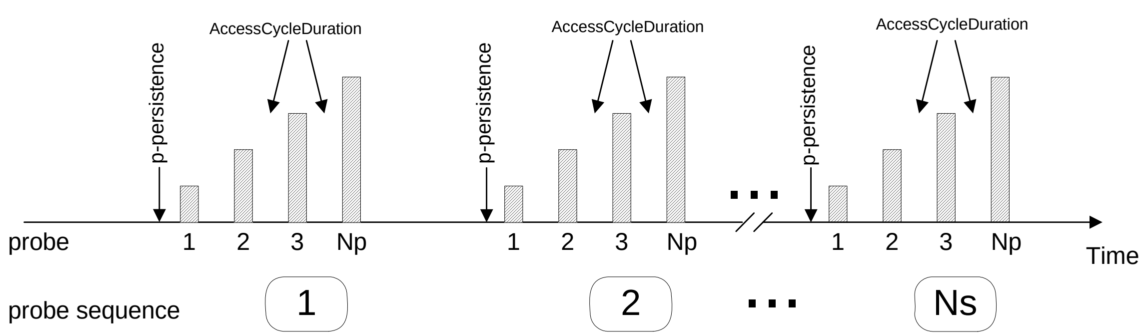

Figure 11-1 Access probe sequences. Ns sequences with Np probes per sequence

For each access probe, the MAC protocol chooses a different AccessSequenceID randomly as described before. This reduces the probability of successive collisions between the access probes of two different ATs.

12 Hopping Modes

12.1 Introduction

The proposed system supports two hopping modes on the forward link data channel, viz. symbol rate hopping and block hopping. The reverse link data channel employs only block hopping.

12.2 FL symbol rate hopping (SRH)

In SRH, the subcarriers allocated to a user are scattered across the band. The hop permutation that maps the assigned hop ports to frequency changes every two OFDM symbols. This permutation changes independently across sectors.

12.2.1 Common pilot channel

In this mode, every user in the sector uses a broadband common pilot channel for channel estimation. This channel is assigned variable power and bandwidth resources on every OFDM symbol. The pilots are equi-spaced in frequency and staggered over consecutive OFDM symbols. The staggering provides excess delay spread mitigation. The staggering phase on the pilot channel is randomized over every odd OFDM symbol. The pilot channel power is set so as to provide sufficient channel estimation resources to a cell edge user. An example scenario with the common pilot channel assigned 6.25 % bandwidth overhead for a carrier with 512 subcarriers is as depicted in Figure 12-1.

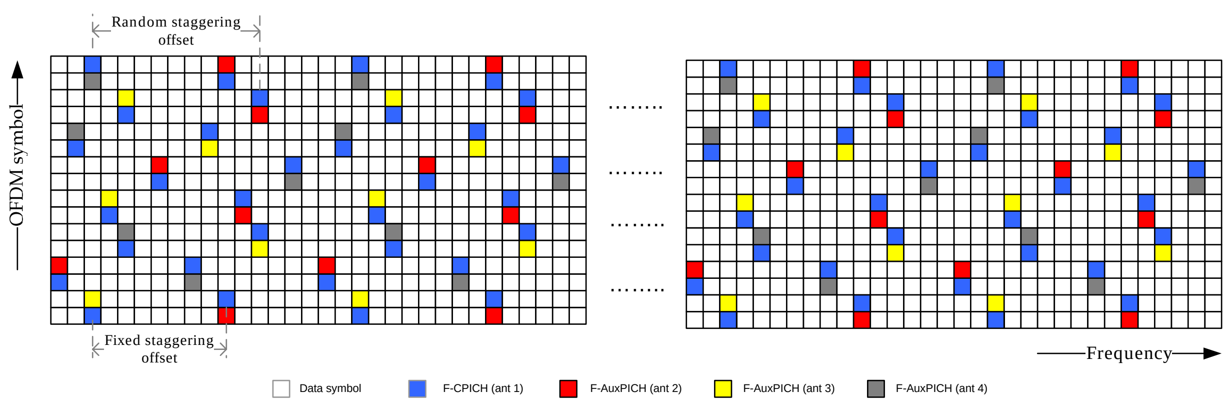

Figure 12-1 Symbol Rate Hopping, common (F-PICH)and auxiliary (F-AuxPICH) pilot channels

12.2.2 MIMO support

Channel estimation for multiple transmit antennas is supported by defining broadband auxiliary pilot channels (F-AuxPICH). The auxiliary pilot channels are time-multiplexed to support multiple transmit antennas. An example scenario where three antennas are time multiplexed over the F-AuxPICH is depicted in Figure 12-1.

12.3 FL block hopping (BH)

In BH (refer Figure 12-2), users are assigned sets of 16 contiguous carriers that are distributed randomly across frequency. The mapping between hop ports and frequency is kept constant through out the physical layer frame (8 OFDM symbols). Each set therefore defines a hop region consisting of 16 contiguous subcarriers and 8 contiguous OFDM symbols. The hop permutation changes every PHYFrame and is independent across sectors.

12.3.1 Dedicated pilot channel

The pilot symbols populate a pre-defined pattern within this two-dimensional grid. Three pilot patterns are depicted in Figure 12-2. Format 0 is the default pilot pattern and is used while transmitting to users capable of supporting only one spatial multiplexing layer. Format 1 is used to support users with high delay spread channels and Format 2 is used by MIMO users with four layers. The pilot symbols are used for interference estimation as well as to estimate the channel over every data symbol within a hop region using time and frequency interpolation.

12.3.2 MIMO support

The dedicated pilot channel consists of six strips of pilots placed as depicted in Figure 12-3. Each strip consists of 3 or 4 contiguous pilots. The former can support up to three effective antennas while the latter can support up to four effective antennas. (The concept of effective antennas is defined in section 20.)

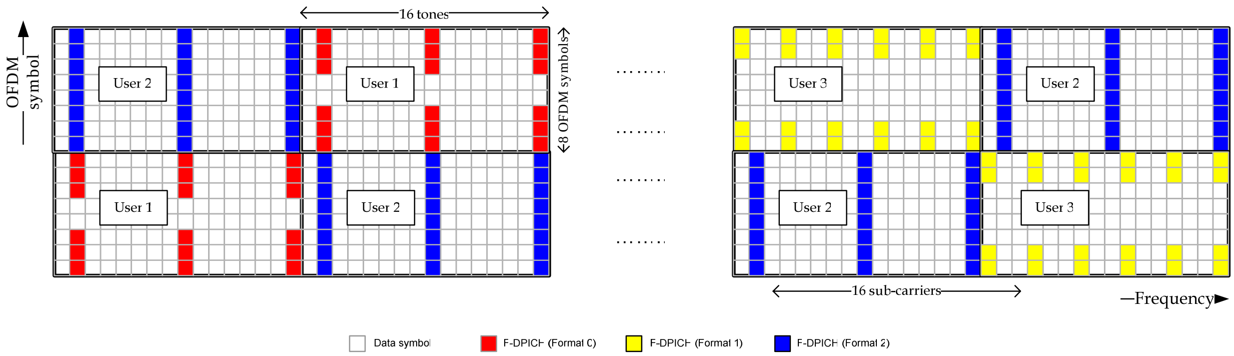

Figure 12-2 Block hopping

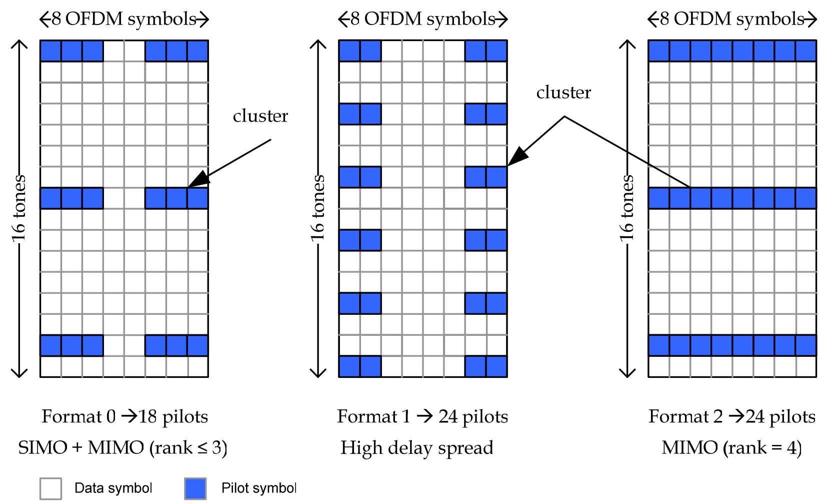

Figure 12-3 Pilot patterns (F-DPICH)

Symbol rate hopping maximizes channel and interference diversity and provides some advantage in channel estimation performance. Block hopping provides efficient support for interference estimation and nulling as well as multiple-antenna techniques such as beamforming, precoding, and SDMA.

12.4 RL block hopping

Every user is assigned sets of 16 contiguous carriers that are distributed randomly across frequency. The mapping between hop ports and frequency is kept constant through out the PHY Frame (8 OFDM symbols). Each set therefore defines a hop region consisting of 16 contiguous subcarriers and 8 contiguous OFDM symbols. The hop permutation changes every PHY frame and is independent across sectors.

12.4.1 Dedicated pilot channel

The pilot symbols populate a pre-defined pattern within this two-dimensional grid. Two pilot patterns are defined and are as depicted in Figure 12-4. Format 0 is the default pilot pattern. Format 1 is used by ATs with high delay spread channels. The pilot symbols are used for interference estimation as well as to estimate the channel over every data symbol within a hop region using time and frequency interpolation. Format 0 and Format 1 can support quasi-orthogonal reverse link transmissions (Section 17) with up to three and two overlapping users respectively.

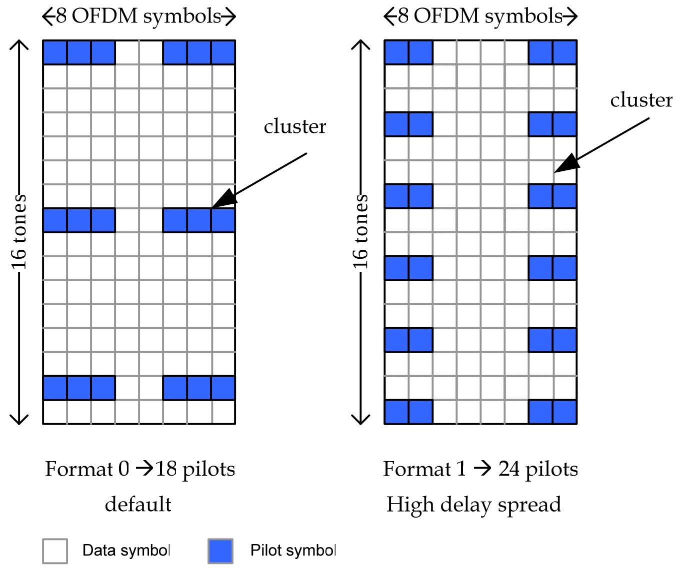

Figure 12-4 Pilot patterns (R-DPICH)

13 Forward Link Control Channels

13.1 Introduction

The FL control channels are used to assign and manage FL and RL resources and specify the respective packet formats, grant access to users in the idle state, acknowledge RL transmissions, send RL power control commands, and send other sector interference indications. These channels are combined in a single physical layer channel called forward link shared signaling channel (F-SSCH) which is described later in this section.

13.2 Forward link signaling messages

A selection of FL signaling messages (called "blocks") is shown in Table 13-1. Columns of this table indicate different fields while rows correspond to different signaling blocks. Every cell in the table indicates multiplicity of a given field. A 4-bit block type field allows the AT to identify the type of block and therefore interpret the subsequent fields. The set of information bits of every block is extended by a 16-bit CRC to enable reliable detection.

Table 13-1 Structure of different FL signaling blocks

Field	Block type	MACID	ChanID	PF	Duration	Ext. TX	Timing	Suppl.	Rank
# bits	**4**	**9-11**	**6-8**	**4-6**	**2**	**1**	**6**	**1**	**2**
NS-FLAB	0001	1	1	1	1	1	0	0	0
Access Grant	0000	1	1	1	0	0	1	0	0
FLAB	0101	1	1	1	0	1	0	1	0
RLAB	1001	1	1	1	0	1	0	1	0
MCW FLAB1	0110	1	1	1	0	1	0	1	0
MCW FLAB2	0111	1	0	3	0	0	0	0	0
SCW FLAB	1000	1	1	1	0	1	0	1	1
CCB	1011	1	Active Carriers (4 bits)				Request Carrier (2 bits)		

The first message is called non-sticky forward link assignment block (NS-FLAB). This message contains MACID of the target AT(s) which is possibly a broadcast MACID, channel ID (ChanID) that indicates the hop-ports assigned (via the channel tree in use), the packet format (PF) to use (specifies modulation, coding and dedicated pilot format in the block hopping mode), the number of PHY frames occupied by this assignment and an indication (denoted by Ext TX) of whether or not to use the extended transmission duration for the assignment (each H-ARQ transmission spans multiple PHYFrames as described in Sections 6 and 7). Unlike other assignments that last until explicitly de-assigned or lost due to packet failure, this assignment lasts for a predefined number of frames and is primarily used to transmit a broadcast or multicast message.

The Access Grant message is used to acknowledge an access attempt by an AT, assign a new MACID along with the initial RL channel assignment and PF, and supply a 6-bit timing adjustment for the AT to align its RL transmission with the RL timing of the AP. The sequence of modulation symbols corresponding to the Access Grant is scrambled according to the index of the preceding access probe transmitted by the AT, enabling the AT to respond only to Access Grant blocks that correspond to the probe sequence that it transmitted.

The forward link assignment block (FLAB) signals a FL resource assignment to an active AT (MACID) with resources assigned indicated by ChanID and spectral efficiency indicated by PF. The new field in this message is the supplemental assignment flag. Whenever set, this flag indicates an incremental assignment that takes effect starting from the new packet. Otherwise, the new assignment replaces the existing one. The reverse link assignment block (RLAB) signals RL resource assignments in a fashion identical to FLAB.

Note that any assignment message automatically de-assigns resources from the AT that is currently using resources corresponding to the ChanID indicated in the assignment message. Hence, assignment messages are often multicast since they target both the intended recipient of the assignment as well as any current owners of the resources specified by the assignment.

Note that Access Grant is an exception since it is scrambled according to the access probe index and therefore is only decodable by the AT that issued access probe. In the event that the assignment carried by the Access Grant yields de-assignment for other AT, a separate de-assignment message is issued.

The multi-codeword FLAB is a FL assignment block that can be used for ATs in the MIMO multi-codeword mode. Unlike other assignment messages, this one indicates four packet formats corresponding to (at most) four MIMO layers (codewords). This assignment message is split into two parts as shown in the Table. In the event when the number of layers in use is less than four, the remaining PF fields are set to zero. The single-codeword MIMO FLAB is similar to the FLAB, except that it also indicates rank of the MIMO transmission. For more details on various MIMO modes, see Section 20.

The last block is called change carrier block (CCB) and is used when the system is in multi-carrier mode (see Section 24) to indicate the 4-bit bit mask of 5MHz carriers with up to 20MHz total system bandwidth (hence up to four carriers) to be used by the AT as well as a 2-bit index of the carrier that is used to issue request (R-REQCH, see Section 14).

Summary of FL signaling messages

It is easy to see that most messages span between 37 and 43 bits including 16-bit CRC when ChanID spans 6 bits (between 39 and 45 bits when ChanID spans 8 bits), except for CCB message that spans only 33 bits. Based on this observation, we can pad all messages to the same maximum number of bits (43 or 45, depending on ChanID size) with a relatively low efficiency loss. While the loss is non-negligible for CCB, the use of this message is quite infrequent. Having a unified size for all signaling messages is convenient when all the messages are encoded and modulated separately since it removes the need to use extra overhead to indicate message sizes.

Modulation of FL signaling messages

All signaling messages will be independently encoded and modulated with the same spectral efficiency, hence resulting in the same number of modulation symbols. Rate 1/3 convolutional encoder with constrained length 9, appropriate puncturing and QPSK modulation will be used to achieve the desired spectral efficiency. The number of modulation symbols per message (hence spectral efficiency) is a quasi-static parameter specified through an overhead message. Spectral efficiencies on the order of 0.5-1 bps/Hz will be used. Furthermore, every message will be power controlled individually according to the FL channel strength of the target ATs. Such a design allows for a low bandwidth overhead with flexible power overhead that depends on the instantaneous F-SSCH load and can be adjusted within every FL PHY Frame.

13.3 Acknowledgement segment

This channel is used to acknowledge RL H-ARQ transmissions and therefore is present in every FL PHY frame to acknowledge the associated RL PHY frame. Each acknowledgement (ACK) is a one bit message indicating either positive or negative acknowledgement. The number of ACK bits required is equal to the total number of usable channel IDs since ACKs are linked to a channel assignment. Every base node of the channel tree is associated with a one bit acknowledgement message. A larger channel assignment corresponding to multiple base nodes uses the ACK bit associated with the base node with the lowest channel ID in the assignment.

ACK bits are encoded using on/off keying so that the absence of ACK transmission implies negative acknowledgement

13.4 Reverse link power control segment

The reverse link power control segment carries CQI erasure indicator bits that are used to indicate RL channel quality. Additionally, it can also carry optional power control bits for ATs that are being served on the reverse link by this sector. As explained in Section 15, the power level of the R-CQICH channel is based on the power control bits from the Reverse Link Serving Sector if these are present, and on the erasure indication bits (also from the Reverse Link Serving Sector) otherwise. BPSK modulation is used for both the power control bits and the CQI erasure indication bits. If both power control bits and CQI erasure indication bits are present for a given MACID, then the two are I/Q multiplexed with each other, thus resulting in QPSK modulation.

Power control frequency of 150Hz or higher has been found to be sufficient for most channel conditions. This implies that every active AT could receive a power control bit once every six PHY frames. Hence, RL power control segment of every PHY frame reports power control bits for a subset of active terminals.

13.5 Fast OSI Segment

This is an optional segment which is used to provide other sector interference indication to an AT from sectors in its active set. As compared to the F-OSICH physical layer channel, this segment is transmitted at a faster rate but with less coverage. When present, this segment carries a three-state OSI value and occupies 8 modulation symbols in each PHY Frame. This segment therefore has a very small overhead. The OSI value is used as part of the reverse link power control algorithm, which is described in 15.3.

13.6 F-SSCH Channelization

The F-SSCH is assigned a minimum of 3 base nodes on the channel tree via signaling on the FL primary broadcast channel. Thus, in a 5MHz deployment, this amounts to about 10% minimum bandwidth overhead with granularity of 3.3%.

The F-SSCH bandwidth is subdivided into four segments of predefined sizes. The segmentation information is signaled in an overhead channel. The first segment carries signaling messages and is zero-padded if not fully used. Modulation symbols of each message are interleaved across the entire F-SSCH assignment to ensure maximum diversity. With the present design, we achieve at least third order diversity with block hopping (achieved with the minimum F-SSCH assignment of three tiles) and higher diversity in symbol rate hopping. The second segment carries ACK bits with every bit mapped to three QPSK modulation symbols. The specific mapping is a function of the base nodes belonging to the assignment being ACK'ed. Modulation symbols of every ACK bit are interleaved over the F-SSCH resources to achieve third order diversity. The third segment carries modulation symbols corresponding to CQI erasure indications and RL power control bits (one modulation symbol per one erasure indication and one power control bit) , and the specific mapping is a function of the MACID of the terminal being power controlled. The fourth segment carries an OSI value and occupies 8 modulation symbols.

F-SSCH allocation parameters such as the total bandwidth allocation, sizing of the four F-SSCH segments (signaling messages, ACK, RL power control, and Fast OSI), power control interval for every MACID, number of modulation symbols per signaling message and a suitable packet format, are periodically broadcast along with other sector parameters.

14 Reverse Link Control Channels

14.1 Introduction

The reverse link control channels include the acknowledgement channel (R-ACKCH), the channel quality indicator channel (R-CQICH), the request channel (R-REQCH), the access channel (R-ACH), the beamforming feedback channel (R-BFCH), the subband feedback channel (R-SFCH), and the pilot channel (R-PICH).

The control channels other than the acknowledgement channel are transmitted in a CDMA control segment that occupies an integer number of contiguous subbands on a single RL traffic interlace. The use of CDMA for the control segment provides statistical multiplexing benefits since resources don't have to be reserved for all the channels. It also provides a broadband reference for power control, subband scheduling, and beamforming in TDD. The presence of request and access channels in the CDMA segment allows terminals to access the system and request resources with low latency and minimal extra overhead. In a synchronous network, multiple sectors can have a common CDMA control segment. This enables multiple sectors in the AT's active set to monitor its control channel transmissions, enabling a fast and efficient handoff mechanism as described in Section 16.

14.2 Acknowledgement channel

The purpose of R-ACKCH is to acknowledge FL H-ARQ transmissions. For FL transmissions in either SISO or single codeword MIMO mode (see Section 20), a single bit has to be transmitted per PHY frame in R-ACKCH while multi-codeword MIMO transmissions will be acknowledged by multiple bits.

A R-ACKCH ID will be associated with every valid base node of the FL channel tree. This allows us to support H-ARQ with the maximum number of channel assignments (i.e., when all assignments have the minimum possible assignment size). A larger channel assignment corresponding to a parent node located at some intermediate level of the channel tree will make use of the R-ACKCH ID associated with the base node with the lowest Channel ID under the said parent node. In multiple codeword MIMO mode, the acknowledgement message consists of multiple bits (one per layer). An MCW MIMO FL assignment is restricted to include a number of base nodes at least as large as the number of MIMO layers. This way, R-ACKCH IDs corresponding to different base nodes of such a MIMO assignment can be used to acknowledge different MIMO layers. Some orthogonal dimensions are set aside for every R-ACKCH ID as described below.

The R-ACKCH occupies N_t 8x8 time frequency tiles that will subsequently be referred to as R-ACKCH tiles (i.e. each R-ACKCH tile spans 8 subcarriers and 8 OFDM symbols) over each RL traffic interlace. Each R-ACKCH tile occupies the lower 8 subcarriers of some traffic tile. The number of R-ACKCH tiles scales (with a granularity of one R-ACKCH tile) as required by the number of traffic channels, with the minimum of N_t =4 R-ACKCH tiles in order to ensure channel and interference diversity. Each R-ACKCH tile is further split into four contiguous 8x2 groups referred to as R-ACKCH subtiles (i.e. each R-ACKCH subtile spans 8 subcarriers and 2 OFDM symbols). Each R-ACKCH subtile accommodates 8 R-ACKCH bits corresponding to 8 different traffic channels. These bits are transmitted using on-off keying (OOK) and spread over the R-ACKCH subtile with orthogonal spreading codes. Each code is given by a column of the DFT matrix of size 16. The remaining 8 codes of the DFT basis can be used by the AP to perform interference estimation. The bit corresponding to each R-ACKCH ID is transmitted over 4 R-ACKCH subtiles with every R-ACKCH subtile taken from a different R-ACKCH tile, thereby ensuring 4-th order diversity. Furthermore, each R-ACKCH bit will be multiplexed with different R-ACKCH bits on different R-ACKCH subtiles, in order to ensure some diversity w.r.t. multiplexing interference that is due to some loss of orthogonality in time/frequency selective channels. Finally, R-ACKCH tiles will be hopping randomly w.r.t. traffic tiles, to make sure that R-ACKCH uniformly punctures different traffic channels.

Such a design allows us to support up to $8 \cdot N_t$ different R-ACKCH IDs. In FDD and TDD with symmetric (1:1) partitioning between FL and RL, the equivalent bandwidth overhead is 1/16. Indeed, every R-ACKCH IDs occupies 8 complex dimensions while the total number of R-ACKCH IDs equals to the number of (minimum) channels, wherein each channel occupies 16 subcarriers and 8 OFDM symbols.

14.3 Channel quality indicator channel

The primary purpose of R-CQICH is to supply the AP with a FL channel quality measure that can be used for scheduling transmissions on the F-DCH. Hence, R-CQICH should be regularly transmitted by every AT. This fact makes R-CQICH a suitable reference for RL power control. Details of the power control scheme are given in Section 15.

The R-CQICH includes CQI reports to support channel quality feedback for SCW and MCW MIMO transmissions (including the number of MIMO layers). In addition, the R-CQICH includes a control CQI report to support channel quality feedback for SISO transmission, and indicating the desired FL serving sector for FL L1 handoff. The control CQI report can also be used by sectors in the active set for FL power control.

14.4 Request channel

The primary purpose of R-REQCH is to request RL traffic resource allocation from the RL serving sector. Furthermore, a R-REQCH sent to a sector different from the RL serving sector will be interpreted as RL handoff request. The R-REQCH transmission includes indications of the QoS flow associated with the request, and the amount of resources requested.

14.5 Feedback for pre-coding and SDMA

The reverse beamforming feedback channel (R-BFCH) provides feedback that allows for adaptive beamforming and spatial multiplexing (SDMA) of multiple ATs on the FL.

14.6 Feedback for sub-band scheduling

The reverse sub-band feedback channel (R-SFCH) provides feedback that allows for adaptive sub-band scheduling on the forward link.

14.7 Reverse link broadband pilot channel

In addition to the described control channels, the AT can transmit a pilot channel (R-PICH) within the CDMA control segment. Such pilot channel is defined as a random HPSK sequence based on the FL serving sector ID and MACID of the AT. This pilot sequence is used to enable adaptive transmission such as RL sub-band scheduling and FL beamforming in TDD where FL propagation channel is assumed to be reciprocal to the RL channel.

14.8 Control Segment Channelization

The control channels described in Sections 14.3 through 14.7 are transmitted in a CDMA control segment that occupies a fraction of the RL bandwidth on a single RL traffic interlace. With the minimum allocation and granularity of 128 subcarriers (about 1.28MHz), this amounts to 1/24 in terms of RL bandwidth overhead for a 5 MHz system. The CDMA control segment occupies a contiguous set of subcarriers (multiple of 128 subcarriers) at any point in time and will be hopping over the entire RL bandwidth.

Every control channel in the CDMA segment will carry at most 10 information bits. Messages with a smaller number of information bits will have the remaining bits set to '0' so that AN can make use of this information to reduce receiver complexity and improve detection performance. Every 10-bit message will be mapped to a Walsh space of size 1024. Furthermore, various control channels will undergo random HPSK scrambling which is defined by the channel type (e.g. R-CQICH), FL serving sector ID and MACID whenever applicable.

The reporting rates and power offsets of each control channel described in Sections 14.3 through 14.7 can be controlled on a per-terminal basis using higher layer messages.

14.9 Access channel

Reverse access channel (R-ACH) will be located in the CDMA control segment and modulated as a 1024 Walsh sequence with the target sector scrambling. The entire space of available probe sequences may be subdivided into a number of groups. AT selects an access probe sequence randomly from a group with the desired parameters such as buffer level, measured FL strength etc. thereby communicating these parameters to the AP though the access process. More details on the access procedure and logic can be found in Section 11. Unlike other control channels within CDMA segment, R-ACH will have an extended guard band and guard time in order to prevent intra-sector interference caused by a misalignment of the access probe with CDMA segment boundaries resulting from the fact that AT in the access phase does not have accurate RL timing information.

15 Reverse Link Power Control

15.1 Introduction

This section discusses how access terminal transmit power is controlled in the proposed MBWA system. Fast closed loop power control is used to set the transmit power levels on the reverse link control channels that are transmitted periodically. The traffic channel power level is set at an offset relative to the control channel power level; this offset is adjusted based on interference indications received from neighboring sectors.

15.2 Reverse link Control Channel Power Control

In the proposed MBWA system, the reverse link channel quality indicator (R-CQI) channel serves as a reference power (power spectrum density) level for RL traffic transmission. The use of control channel as reference power level is necessary to control the inter-carrier-interference and to ensure appropriate power level for minimum RL date rate.

A closed loop power control algorithm has been designed for the R-CQI channel, which is periodically transmitted on the reverse link to convey a limited number of information bits in each codeword regarding down link channel quality. The closed loop algorithm may be based on erasure indications or on up-down commands. The power control mode to be used is determined using a bit in the overhead message protocol.

Each sector in the active set sends a one bit erasure indication to indicate whether the CQI has passed the erasure decoding or not. In the erasure-indication based algorithm, the erasure indication from the serving sector is interpreted as up/down power control command by the AT. When the AT receives an erasure indication from the serving sector, it increases the CQI channel PSD by “PowerControlUp”; otherwise, it reduces the PSD by “PowerControlStepDown”. The up and down step sizes are determined by the target erasure rate and are communicated to the AT using the ActiveSetUpdate message.

In the up-down command based algorithm, the serving sector sends an explicit up/down indication in addition to the CQI erasure indication. When the AT receives an up indication, it increases the CQI channel PSD by “PowerControlStepUp”; otherwise, it reduces the PSD by “PowerControlStepUp”. Note that equal up and down step sizes are used for this algorithm. The step size is communicated to the AT using the ActiveSetUpdate message.

15.3 Reverse Link Traffic Channel Power Control

While traffic channel transmissions from different terminals occupy different dimensions in time and frequency, it is not desirable to have a large difference in received power across subcarriers since this will increase receiver dynamic range requirements and also cause loss of orthogonality with time and frequency errors. Also, in a multi-sector layout, high inter-sector interference can drastically reduce the network capacity. Hence, it is necessary to tightly control the transmit power levels.

Since the proposed system uses orthogonal multiple access for traffic channels, the serving sector lacks information regarding the inter-sector interference caused by RL traffic originated from this sector. Hence, it is desirable to have an interference control algorithm implemented in an AT where the interference information could be made readily available from other sectors. In the proposed system, a load indication is broadcasted every superframe over the forward link other sector interference channel (F-OSICH) from each sector when the average interference over thermal level exceeds a target threshold. The load indicator takes on one of three values (0, 1 and 2) to control the interfering AT’s power level. The coverage of F-OSICH is the same as the acquisition pilots, which penetrate far into neighboring sectors. In addition, the load indicator can also be transmitted over the Fast OSI segment of the F-SSCH. This segment is transmitted at a faster rate, namely once every PHY Frame, however it has more limited coverage. The purpose of this segment is to reach nearby ATs who can potentially cause high levels of interference.

The amount of inter-sector interference per-subcarrier caused by a given AT is determined by the transmit power level used by that AT and the location of the AT relative to the neighbor sectors. For the traffic channels, power control may be performed such that each AT is allowed to transmit at a power level that is as high as possible while keeping intra-sector and inter-sector interference to within acceptable levels. An AT located closer to its serving sector may be allowed to transmit at a higher PSD level since this AT will likely cause less interference to neighbor sectors. Conversely, an AT located farther away from its serving sector and toward a sector edge may be restricted transmit at a lower power level since this AT may cause more interference to neighbor sectors. Controlling transmit power in this manner can potentially reduce the total interference observed by each sector while allowing “qualified” ATs to achieve higher SNRs and thus higher data rates.

The transmit Power Spectral Density (PSD) (defined as the transmit power per assigned subcarrier) for a traffic channel for a given AT may be expressed as:

$$\mathrm{P}_{\mathrm{dch}}(n) = \mathrm{P}_{\mathrm{ref}}(n) + \Delta\mathrm{P}(n),$$

where $\mathrm{P}_{\mathrm{dch}}(n)$ is the transmit PSD for the traffic channel for update interval *n*;

$\mathrm{P}_{\mathrm{ref}}(n)$ is a reference PSD level for update interval *n*; and

$\Delta\mathrm{P}(n)$ is a transmit PSD delta for update interval *n*.

The PSD levels $\mathrm{P}_{\mathrm{dch}}(n)$ and $\mathrm{P}_{\mathrm{ref}}(n)$ and the transmit power delta $\Delta\mathrm{P}(n)$ are given in units of decibels (dB/Hz). The power control algorithm described in this document will be called Delta-based power control because of the transmit power delta $\Delta\mathrm{P}(n)$. The reference PSD level is the amount of transmitted PSD needed to achieve a target SNR for a designated transmission, which is provided by the RL CQI channel in our design.

The transmit PSD for the traffic channel is set based on (1) the amount of inter-sector interference the AT may be causing to other ATs in neighbor sectors, (2) the amount of intra-sector interference the AT may be causing to other ATs in the same sector, (3) the maximum power level allowed for the AT.

In order to keep the intra-sector interference at acceptable levels, the transmit PSD delta, $\Delta\mathrm{P}(n)$, is constrained to be within a range as follows:

$$\Delta\mathrm{P}(n) \in [\Delta\mathrm{P}_{\mathrm{min}},\ \Delta\mathrm{P}_{\mathrm{max}}],$$

where $\Delta\mathrm{P}_{\mathrm{min}} / \Delta\mathrm{P}_{\mathrm{max}}$ is the minimum/maximum transmit PSD delta allowable for a traffic channel.

Each sector can estimate the average amount of interference relative to thermal noise power (referred to as IOT) experienced by that sector from terminals in other sectors, where the thermal noise level is measured during the Reverse Link Silence Interval. Each sector broadcasts an indication of its interference measurements for use by ATs in other sectors. For simplicity, the following description assumes the use of a single load indicator bit to provide interference information. Each sector may set its other sector indication (OSI) as follows:

$$\mathrm{OSI}_{\mathrm{m}}(n) = \begin{cases} \text{'1' or '2'}, & \text{if } \mathrm{IOT}_{\mathrm{meas},\,m}(n) \geq \mathrm{IOT}_{\mathrm{target}}\text{, and} \\ \text{'0'}, & \text{if } \mathrm{IOT}_{\mathrm{meas},\,m}(n) < \mathrm{IOT}_{\mathrm{target}}\text{,} \end{cases}$$

where $\mathrm{IOT}_{\mathrm{meas},\,m}(n)$ is the measured IOT for sector *m* in time interval *n*; and $\mathrm{IOT}_{\mathrm{target}}$ is the desired operating point for the sector. OSI value of '2' is used to indicate excessive IOT level. Here IOT refers to the interference-over-thermal , which is a ratio of the total interference power observed by the sector to the thermal noise power.

Each AT can estimate the channel gain (or propagation path gain) for each sector that may be interfered by the AT. For a user, the channel gain ratio between the serving sector and a neighbor sector may be viewed as a "relative distance" that is indicative of the distance to a neighbor sector relative to the distance to the serving sector.

In one setup, each AT monitors the OSI broadcast by neighbor sectors and only responds to the OSI of the strongest neighbor sector, which has the smallest channel gain ratio. If the OSI from that sector is set to '1' or '2' (due to the sector observing higher than nominal inter-sector interference), then the AT adjusts its delta downward. Conversely, if the OSI is set to '0', then the AT adjusts its delta upward.

The load indicator bit thus determines the direction in which to adjust the transmit power. The amount of transmit power adjustment for each AT may be dependent on the current transmit power level (or the current transmit power delta) of the AT and the channel gain ratio for the strongest neighbor sector.

The AT sends the transmit PSD delta and the maximum number of subcarriers that the AT can support at the current transmit PSD delta via in-band signaling. This information is used by the AN for making reverse link assignments. Thus, an AT having a low delta may be assigned a large number of subcarriers so that it can use all its transmit power to achieve a higher data rate.

The proposed power control mechanism has the following advantages.

1. Explicit interference control leads to tight interference tail distribution.

2. Delta-based power control naturally shapes the PSD of users that cause high and low interferences.

16 Handoff

16.1 Introduction

In this section, we describe the techniques used to handle handoff between access network sectors in the MBWA system. The system has been designed to meet the following handoff requirements:

- Support for fast handoff
  - to minimize the handoff impact on latency-sensitive traffic,
  - to minimize the response time of the system to rapid variations in the path loss and shadowing components of the channels of serving and interfering sectors at vehicular speeds, especially with frequency reuse of one, and to get fast fading gains and improve diversity at pedestrian speeds.
- Support for disjoint links (i.e. an access terminal can be served by different sectors on the forward and reverse links)
  - to achieve best cell site selection gains in situations where the best serving sectors for forward and reverse links are different.
- Low signaling overhead
  - this becomes especially more important when fast handoff is used to obtain fast fading gains at pedestrian speeds, or when the handoff rate is high due to high vehicular speeds.

The sector from which the Access Terminal (AT) received the last Forward Link Assignment Block (FLAB) is referred to as the FL serving sector, and the sector from which the AT received the last Reverse Link Assignment Block (RLAB) is referred to as the RL serving sector. The FL and RL serving sectors of an AT can also be assigned by the Access Network (AN) through special assignment blocks for handoff called FLAB-HO and RLAB-HO. The active set of an AT is the set of sectors that have allocated MAC IDs and dedicated control resources to the AT. The AT monitors the SSCH from the FL and RL serving sectors, as well as the desired FL and RL serving sectors. The term handoff is used to refer to a change in the AT's FL or RL serving sector, which occurs when the AT receives a link assignment block from its desired serving sector and the desired serving sector is different from the current serving sector, or when the AT receives a link assignment block for handoff (FLAB-HO or RLAB-HO) from its current serving sector. The AT can perform a handoff to any member of the active set.

On a handoff, a new sector in the active set becomes the serving sector. One objective of fast switching is to not introduce any packet loss at the higher layers while allowing for uninterrupted transmission to the access terminal. A second objective is to minimize the backhaul communication required between the AP's in the active set. There is a trade off between the amount of back haul communication required and the handoff latencies that can be achieved.

In order to continue uninterrupted data transmission, the new serving sector needs to know the forward looking RLP state for the forward link, where the forward looking RLP state is defined as the data received at the anchor AP and not yet transmitted, and the data that needs to be retransmitted based on ReceiverStatus messages from the access terminal. This RLP state is transferred to the new serving sector as part of the L2 handoff negotiation.

16.2 Active Set Management

In this subsection, we briefly describe the active set management protocol.

- The AT performs SINR measurements on the pilots suggested by the AN and on pilots it autonomously searches. The AT filters the measured values to remove measurement noise and fast fading components, and reports the filtered values to the AN through its RL serving sector.
- Based on the pilot SINR reports from the AT, the AN determines the AT's active set. Members of the active set are indexed using a three-bit field that is used in the reverse link Channel Quality Indicator (CQI) and REQuest (REQ) channels.
- The AN assigns dedicated control resources to the members of the active set for the AT, assigns a MAC ID for the AT on each member of the active set and transmits an active set assignment message to the AT through its FL serving sector.
- The AT monitors assignment channels from the members of the active set that are currently serving the AT on forward and reverse link, as well as the desired serving sectors on both links.

The handoff procedures described in the following two sub-sections apply to the AT initiated handoffs, when the different sectors are synchronous with each other. Procedures for handing off between asynchronous sectors are described in 16.5.

16.3 Forward Link Handoff

The forward link handoff is performed in the following steps:

- **Selection:** Forward link serving sector selection is primarily based on the FL pilot strength measurements. The AT can use the acquisition pilots (similar to the active set management algorithm), common pilot channel (if present), or the pilots on the shared signaling channel, to select its desired FL serving sector.
- **Indication:** The AT requests a forward link handoff by computing and targeting the CQI to its desired FL serving sector (setting the target sector field in the CQI to the active set index of its desired FL serving sector), and setting the one-bit Desired FL Serving Sector flag (DFLSS bit).
- **Signaling and Detection:** All CQI channels (regardless of the target sector) are scrambled using the FL serving sector scrambling code. All the members of the active set, including the desired FL serving sector, know the FL serving sector of the AT, and are able to decode the CQI values that are sent to them. A successful decoding of the CQI carrying handoff request at the desired FL serving sector indicates a FL handoff request.
- **RLP State Transfer:** When the desired FL serving sector decodes a CQI from the AT targeted to it and with the DFLSS bit set, it sends backhaul messages to indicate the handoff and request the RLP state of the AT from its current serving sector.
- **Completion:** Once the desired serving sector receives the RLP state of the AT, it sends a Forward Link Assignment Block (FLAB) to that AT on the shared signaling channel, by which it indicates the completion of the FL serving sector switch or L1 handoff.

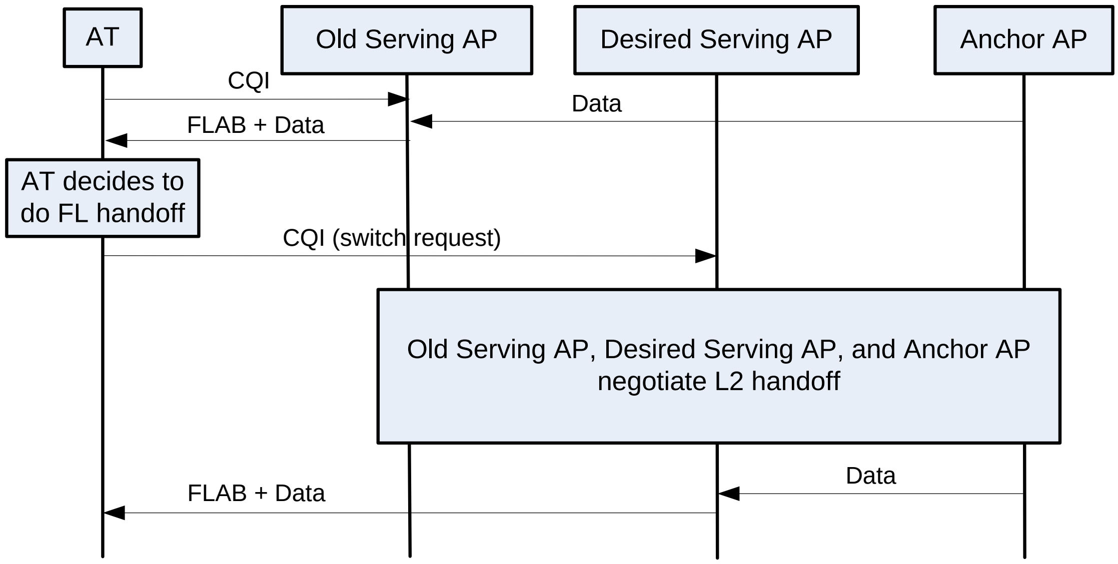

Figure 16-1 FL Handoff Call Flow

Figure 16-1 shows the call flow of a forward link handoff event. Notice the following points in the proposed FL handoff algorithm:

- Even though the FL serving sector selection is primarily based on the FL pilot strength, it needs to take the corresponding RL channel quality into account as well. Sectors with very poor RL channel quality will not be included in the list of considered sectors, so that the new FL serving sector will have sufficient RL quality. The RL quality can be estimated based on the CQI erasure rate on that link, which itself can be measured using the up/down power control commands received from the corresponding sector. More details on this procedure are provided in the description of the RL design.
- AT ignores any FLAB received from the old serving sector after it receives an FLAB from the desired serving sector, because after receiving the FLAB from the desired serving sector, the serving sector switch has completed and the old serving sector is neither the serving sector, nor the desired serving sector.
- The AT only requests a server switch (as opposed to a server switch command). This allows for the old serving sector to continue serving the AT even after the AT has requested a handoff, and minimizes the service outage duration. The old serving sector stops scheduling new packets when it receives the RLP state request from the desired serving sector. Once all the HARQ processes to the AT are terminated, the old serving sector sends the RLP state to the new serving sector.
- All the signaling is done using only the CQI channel, and there is no need for additional channels for advance indication of handoff.

16.4 Reverse Link Handoff

On the reverse link, handoff is performed in the following steps:

- **Selection:** RL serving sector selection is primarily based on the CQI erasure rate estimate. All members of the active set to which AT sends a CQI value, send up/down power control commands to the AT based on the CQI erasures. Therefore, the up commands received from any member of the active set can be considered as CQI erasure indicators, and AT can use them to obtain an estimate of RL channel quality between the AT and that sector. The AT then selects the RL serving sector based on the lowest filtered percentage of erasures.
- **Indication:** The RL handoff is indicated in the desired RL serving sector field of the REQ channel.
- **Signaling and Detection:** AT sends the REQ to the desired RL serving sector. Each AP in the active set only decodes the requests targeted to it. A successful decoding of a REQ at an AP other than the current RL serving sector indicates a RL handoff request.
- **Completion:** The desired RL serving sector completes the handoff by sending RL Assignment Block (RLAB) to the AT.

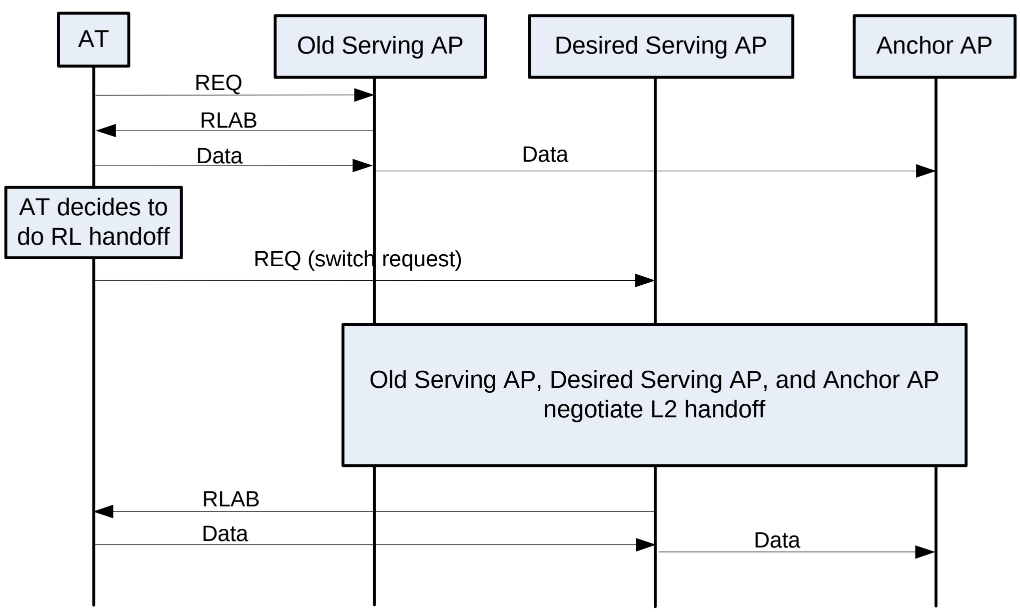

Figure 16-2 RL Handoff Call Flow

Figure 16-2 shows the call flow of a RL handoff. Notice that the AT ignores any RLAB received from the old RL serving sector after it receives an RLAB from the desired serving sector, because after receiving the RLAB from the desired serving sector, the serving sector switch has been completed and the old serving sector is neither the serving sector, nor the desired serving sector.

16.5 Comments on Asynchronous Deployment

The proposed handoff algorithms in the previous sections can be used in both synchronous and asynchronous deployments. In the case of asynchronous deployment, the active set may consist of multiple synchronous subsets. The AT, depending on its baseband processing capabilities, can send RL control channels to more than one synchronous subset of the active set, and decode FL SSCH from the desired FL and RL serving sectors, even if they do not belong to the same synchronous subset as the current FL and RL serving sectors. Alternatively, it may send the RL control channels to only the synchronous subset to which the current FL and RL serving sectors belong, and decode the SSCH from the desired FL and RL serving sectors, only if they belong to the same synchronous subset as the current FL and RL serving sectors. In the latter case, handoff decisions across synchronous subsets will be made based on the filtered ACQ pilot SINR. AT will terminate CQI transmission on the control segment of the old serving sector, and will send a special access probe, which is scrambled by the MACID of the AT, to the desired serving sector. The access grant will provide a timing correction and also will complete the handoff.

17 Quasi-Orthogonal Reverse Link

17.1 Introduction

While orthogonal multiple access schemes such as OFDMA benefit from the elimination of intra-sector interference, they have the disadvantage of becoming dimension limited as the number of receive antennas at the AP increases. This is in comparison to non-orthogonal multiple access schemes like DS-CDMA which while intra-cell interference limited (unless interference cancellation is used at the AP), benefit from a linear increase in capacity with the number of receive antennas. As the number of receive antennas increases, a non-orthogonal scheme like CDMA can provide higher capacity than an orthogonal multiple access scheme. A simple rationale for this behavior is as follows. A non-orthogonal scheme can operate in the interference limited regime where signal to noise ratio (E_S/N_0) per modulation symbol is low. In this regime, the system capacity scales linearly with E_S/N_0. Hence the capacity increases linearly with the number of receive antennas. In an orthogonal RL, the total load is limited by the number of orthogonal dimensions (system bandwidth) while spectral efficiency per dimension scales logarithmically with E_S/N_0. Hence, the capacity of an orthogonal RL scheme scales logarithmically with the number of antennas.

In the proposed system, this fundamental limitation of orthogonal multiple access is mitigated by using a quasi-orthogonal multiplexing scheme where multiple ATs of the same sector are assigned the same bandwidth resources. Spatial processing with multiple antennas is used to recover signals from different ATs. Thus, we have a design that retains the benefits of an orthogonal design when the number of receive antennas is small, and offers improved capacity scaling with the number of antennas.

17.2 Quasi-orthogonal Reverse link with random hopping

The proposed quasi-orthogonal scheme achieves intra-sector interference diversity through random hopping. Specifically, AP assignment to each AT consists of a set of time-frequency blocks that hop in frequency over time. When such hopping sequences assigned to multiple ATs overlap, every AT will overlap with a set of ATs on every time-frequency block. The sets of ATs will be different for different blocks, hence providing co-channel interference diversity which is advantageously used by H-ARQ to terminate packet transmission at an appropriate rate. The concept of random hopping should be implemented to support different values of the multiplexing factor (e.g. the number of ATs assigned the same time-frequency blocks). It is also important to ensure co-existence of quasi-orthogonal with orthogonal assignments which may be needed to support high QoS requirements. A general approach is explained in the following section.

17.3 Multiplexing factor control through scheduling

As described in section 9.2, the channel assignment structure is defined by a channel tree, each base node of which maps to a time-frequency block. The mapping from base nodes to time-frequency blocks is randomized in time, thereby resulting in frequency hopping for each assignment. In the quasi-orthogonal mode, one possibility is to use a channel tree that contains Q identical sub-trees as illustrated in Figure 17-1 such that base nodes of each sub-tree are randomly mapped to the same set of time-frequency blocks. Within each sub-tree, the base nodes map to disjoint resources.

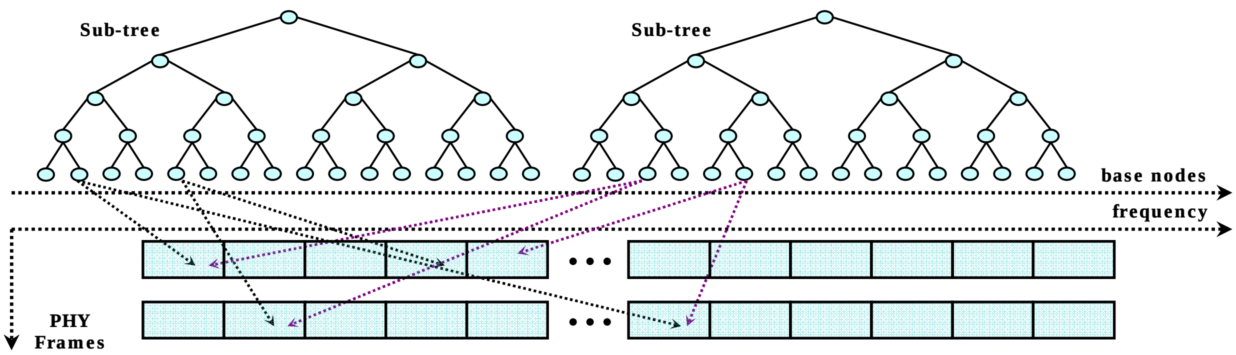

Figure 17-1 Multiple channel trees to support quasi-orthogonal operation

For orthogonal operation, ATs are scheduled on a single sub-tree. In quasi-orthogonal mode, an integer multiplexing factor Q can be achieved by loading exactly Q sub-trees. A fractional value of Q can be achieved through a symmetric partial loading of $\lceil Q \rceil$ sub-trees, where $\lceil . \rceil$ is the integer ceiling operation. Alternatively, fractional Q can be achieved through asymmetric partial loading so that the RL scheduler starts assigning resources to an AT on a new sub-tree, in addition to the existing fully loaded sub-trees.

17.4 Orthogonal pilot multiplexing

In order to optimize demodulation performance for the quasi-orthogonal scheme, pilots should be designed so as to enable accurate estimation of Q channels corresponding to the ATs multiplexed over a block. These channel estimates will be used to set the parameters of a linear receiver (such as MMSE) or a non-linear (e.g. successive cancellation) receiver at the AP. In OFDMA, RL traffic resources can be assigned in units of time-frequency blocks with local (dedicated) pilots placed in every block for channel and interference estimation. Our approach consists of defining a few contiguous clusters of pilots with cluster location optimized to minimize channel estimation error in orthogonal mode and to multiplex pilot dimensions of different ATs over these clusters, by using some orthogonal codes. Different sequences of orthogonal codes will be implicitly assigned to different ATs if every channel tree described in the previous section is associated with an orthogonal sequence. An example of time-frequency block and pilot design that supports quasi-orthogonal operation with Q=3 is shown in Figure 17-2. Here pilot symbols are arranged in six clusters with three "strips" in frequency and two "strips" in time. In quasi-orthogonal mode, these clusters can be used to orthogonally multiplex pilots of different ATs. This multiplexing can be achieved by assigning e.g. different vectors of a 3x3 orthogonal basis (such as three-dimensional DFT basis) to different channel sub-trees.

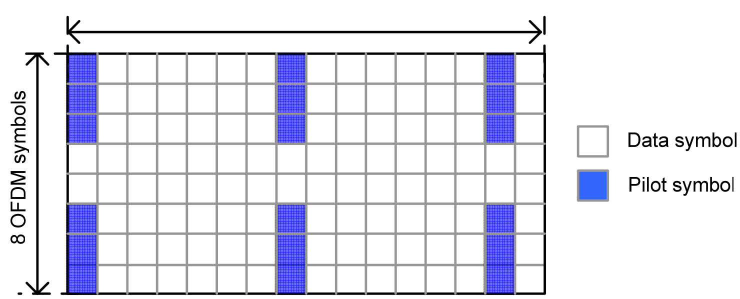

Figure 17-2 Pilot and data symbol placement in a time-frequency traffic block

This pilot structure can be also used to facilitate decoding of RL transmission by different sectors of the serving cell. This mode of operation known as softer handoff helps to improve system coverage and link budget for ATs located close to the sector boundaries. In orthogonal mode, different sectors of the same cell will use the same pilot/data symbol multiplexing with orthogonal pilot multiplexing for different sectors of the same cell. Such an orthogonal multiplexing allows any sector to accurately extract channel state corresponding to RL transmission within this sector as well as RL transmissions taking place in the adjacent sectors. Hence, traffic demodulation within a given sector can be assisted by a neighbor sector that will use appropriate receiver architecture to separate data symbols transmitted within the two sectors. Performance of softer handoff can be enhanced substantially when multiple receive antennas are used in every sector AP to enable spatial separation of data symbols transmitted in different sectors.

18 Fractional Frequency Reuse

18.1 Introduction

The proposed system has been designed to be robust to interference so that it can be deployed with universal frequency reuse across all sectors. This section discusses how frequency planning could be used to enhance coverage and QoS. Frequency reuse is often used in interference limited systems to improve channel C/I, hence improving link reliability at sector edge. The resulting channel quality improvement, however, comes at a cost of bandwidth reduction, which is not necessarily a good capacity tradeoff. For example, a system with 1/3 frequency reuse needs to improve an AT's spectral efficiency by 200% to achieves the same throughput as a 1/1 reuse system. According to the AWGN capacity formula, a 200% spectral efficiency gain requires at least 5 dB gain in C/I in the linear (low SNR) regime, and each bps/Hz improvement in the nonlinear (high SNR) regime asymptotically requires 3 dB gain in C/I.

18.2 Fractional Frequency Reuse Concept

A fractional frequency reuse (FFR) scheme is a frequency reuse scheme with reduced bandwidth overhead compared to traditional frequency reuse schemes. Unlike the case of traditional reuse schemes, where the same frequency is only used in 1 out of 3, 7 or 12 sectors (sectors), fractional frequency reuse allows ATs in different channel condition to enjoy different frequency reuse factor.

Fractional frequency reuse can be implemented in different forms. One method is to disallow transmission on some set of subcarriers in each sector. A more general method is to have different transmit power restrictions for different sets of subcarriers.

In this section, we describe a sample scheme for implementing fractional frequency reuse.

In a static fractional frequency reuse scheme, each AT is associated with a particular frequency reuse plan that corresponds to a frequency "reuse set". Sectors are colored such that no neighboring sectors share the same color. Note that the same set of subcarriers is used in sectors of the same color. A reuse set is defined as the set of sector colors associated with the strongest neighboring sectors. In Figure 18-1, the reuse sets of a few ATs are illustrated for a 3 sectors per base station deployment. In Figure 18-1, the AT denoted by a blue dot has reuse set (**1**) since its serving sector of color **1** is much stronger than any other sectors. The two ATs denoted by squares at the edge of sectors **1** and **3**, have reuse sets (**1**, 3) and (**3**, 1), respectively since both sector 1 and 3 are received strongly by these ATs. The AT denoted by triangle has reuse set (**1**, 2, 3) since sectors of all three colors are received strongly.

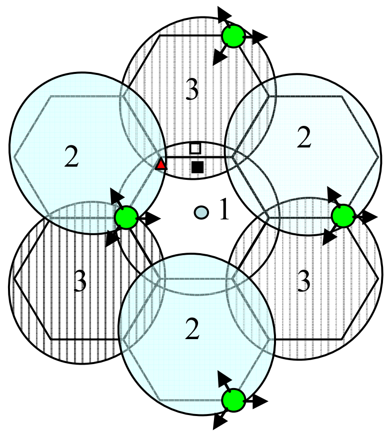

Figure 18-1 Sector layout and FFR reuse set assignment where ○ denotes an AT with reuse set (1), ■ denotes an AT with reuse set (1,3), ▲ denotes an AT with reuse set (1, 2, 3) and □ denotes an AT with reuse set (3,1). Note that the bold-faced font denotes serving sector color.

The goal of FFR design is to deploy frequency patterns such that an AT can avoid interfering or being interfered by non-serving sectors in its reuse set. Since sectors in the reuse set are those that contribute most significantly to the overall interference on FL, avoiding interference from these sectors is expected to effectively reduce the interference.

One FFR frequency plan with 3 color sectorization is as the following:

1. Let Ω denote the overall spectrum and ϕ denote the empty set. Define three over lapping frequency sets, F_1, F_2 and $F_3 \subset \Omega$, where $|F_i \cap F_j| \neq \phi$, for $i,j \in \{1, 2, 3\}$ and $|F_1 \cap F_2 \cap F_3| = \phi$. The spectrum allocation is illustrated in Figure 18-2.

2. For ATs served by sector of color *i*, only subcarriers that belong to the frequency set $\Omega \backslash F_i$ are used. If a neighboring sector of color *j* is added to the reuse set of an AT, the AT's subcarriers are further restricted to $(\Omega \backslash F_i) \cap F_j$. If a third sector is added to the reuse set, say *k*, then the subcarriers are further restricted to $(\Omega \backslash F_i) \cap F_j \cap F_k$. The mapping from the reuse set to the allowed frequency set is shown in Table 18-1.

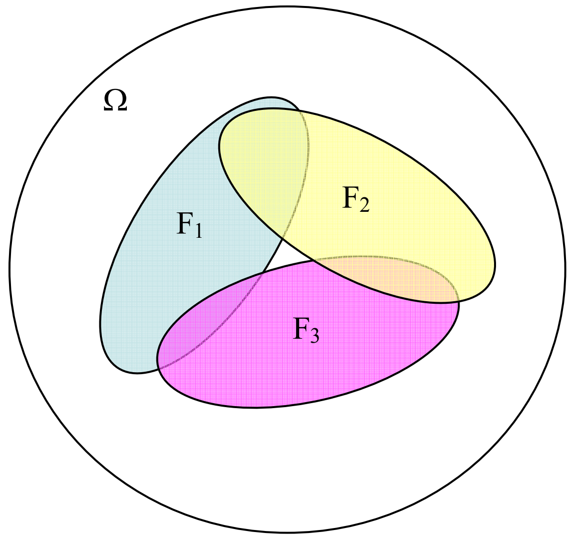

Figure 18-2 Frequency allocation for FFR scheme with maximum reuse set size 3

Note that Ω denote the overall spectrum and F_i denote the unused frequency set in sector *i*.

Table 18-1 Mapping between reuse sets and the usable frequency set

Reuse Set	Usable Frequency Sets	Description
(**<u>1</u>**)	$U_1 = \Omega \setminus F_1$	Main/unrestricted usable set for sector 1
(**<u>1</u>**, 2)	$U_{1\text{-}2} = U_1 \cap F_2 = F_2 \setminus F_1$	Restricted usable set with no interference from sector 2
(**<u>1</u>**, 3)	$U_{1\text{-}3} = U_1 \cap F_3 = F_3 \setminus F_1$	Restricted usable set with no interference from sector 3
(**<u>1</u>**, 2, 3)	$U_{1\text{-}23} = U_1 \cap F_2 \cap F_3 = F_2 \cap F_3$	More restricted usable set with no interference from sectors 2 & 3
(**<u>2</u>**)	$U_2 = \Omega \setminus F_2$	Main/unrestricted usable set for sector 2
(**<u>2</u>**, 1)	$U_{2\text{-}1} = U_2 \cap F_1 = F_1 \setminus F_2$	Restricted usable set with no interference from sector 1
(**<u>2</u>**, 3)	$U_{2\text{-}3} = U_2 \cap F_3 = F_3 \setminus F_2$	Restricted usable set with no interference from sector 3
(**<u>2</u>**, 1, 3)	$U_{2\text{-}13} = U_2 \cap F_1 \cap F_3 = F_1 \cap F_3$	More restricted usable set with no interference from sectors 1 & 3
(**<u>3</u>**)	$U_3 = \Omega \setminus F_3$	Main/unrestricted usable set for sector 3
(**<u>3</u>**, 1)	$U_{3\text{-}1} = U_3 \cap F_1 = F_1 \setminus F_3$	Restricted usable set with no interference from sector 1
(**<u>3</u>**, 2)	$U_{3\text{-}2} = U_3 \cap F_2 = F_2 \setminus F_3$	Restricted usable set with no interference from sector 2
(**<u>3</u>**, 1, 2)	$U_{3\text{-}12} = U_3 \cap F_1 \cap F_2 = F_1 \cap F_2$	More restricted usable set with no interference from sectors 1 & 2

Subcarriers assigned to an AT according to FFR are never used by non-serving sector(s) in this AT's reuse set. According to Table 18-1, an AT with reuse set (**1**,3) is allocated subcarriers in $F3\backslash(F1\cap F3)$ and all ATs with sector 3 as serving sector only use subcarriers in $\Omega\backslash F3$. These are mutually disjoint frequency sets, hence ATs with reuse set (**1**,3) will not be interfered by sector 3. In general, ATs whose subcarriers are allocated according to FFR are immune to inter-sector interference from the strongest interfering sectors on FL.

The interference experienced by an AT on FL decreases as the size of the reuse set increases. Note that a larger reuse set size implies a higher bandwidth partial loading factor, hence smaller usable bandwidth given fixed amount of total resources. As shown in Figure 18-3, an AT *x* with reuse set size 3 is not interfered by any sectors in the first tier neighbors, i.e., 1/3 reuse; an AT with reuse set size 2 is not interfered by the dominant interfering sectors in the first tier neighbors, i.e., 2/3 reuse that avoids the most dominant interfering sector; an AT *x* with reuse set size 1 is not interfered by transmissions to ATs handing off to the serving sector of AT *x*, i.e., 1/1 reuse that avoids high power transmission to edge AT in a power controlled system.

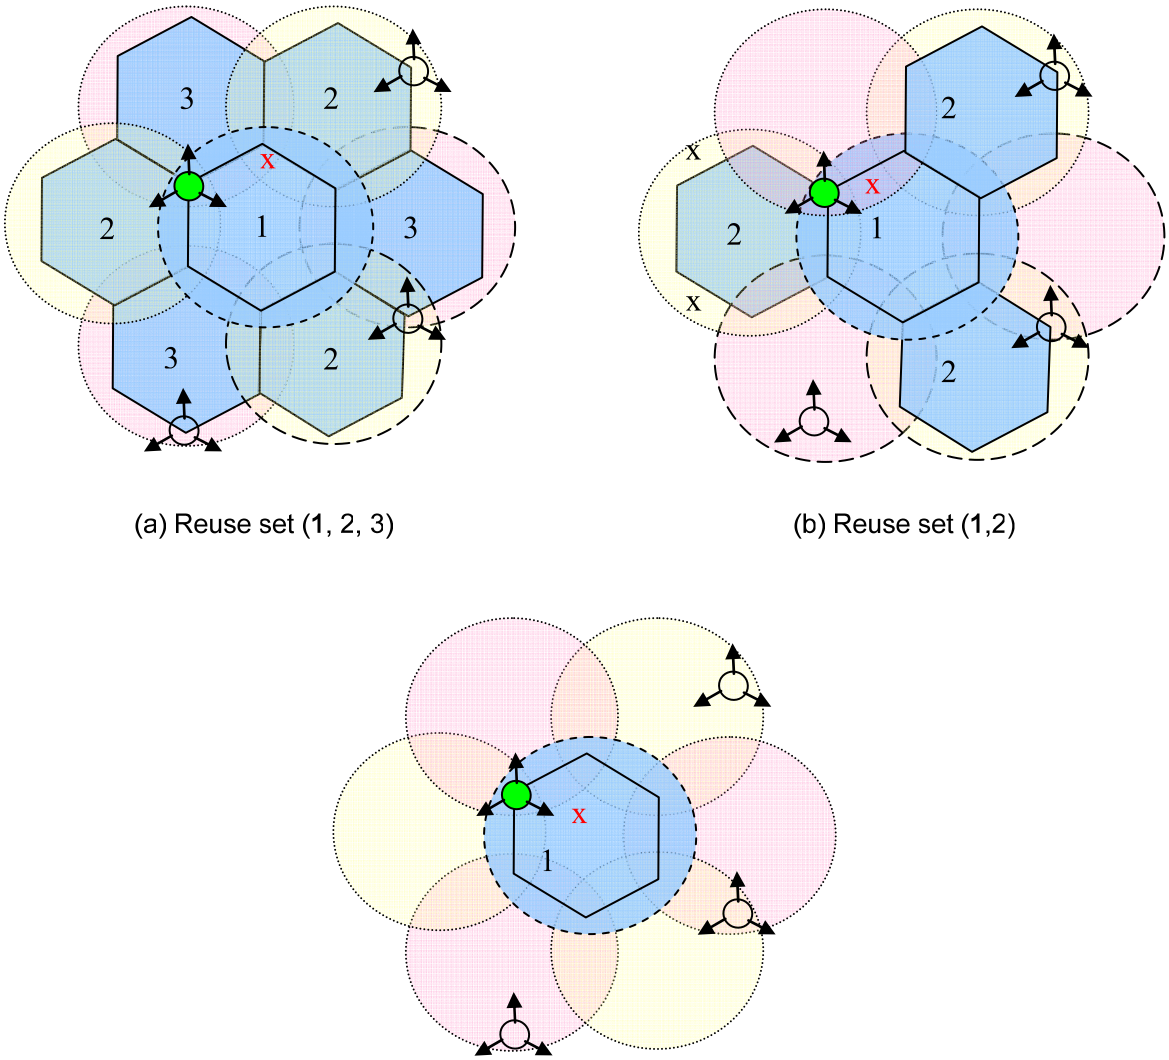

(a) Reuse set (**1**, 2, 3) (b) Reuse set (**1**,2)

(c) Reuse set (**1**)

Figure 18-3 Examples of interference experienced by ATs with different reuse set size, where other ATs in the blue shaded areas are served over subcarriers orthogonal to AT *x*

18.3 Static FFR Reuse Set Management

Static FFR algorithm is defined as a version of FFR where each AT is associated to a fixed reuse set for a duration beyond the transmission of one physical layer packet. In the proposed system, ATs report FL pilot measurements from neighboring sectors to the serving sector either periodically or when the strength of some pilot changes significantly. The scheduler at the AP can use these pilot measurements to schedule the ATs on specific reuse sets depending on the interference levels seen from neighboring sets.

One limitation of the static FFR algorithm is the limited number of carriers in each reuse set, hence limiting the peak rate in each reuse set and reducing the trunking efficiency of overall system. An alternative approach involves the use of a dynamic FFR scheduler.

18.4 Dynamic FFR

The basic idea behind dynamic FFR is to schedule ATs on different reuse sets on a packet-by-packet basis, while also enforcing fairness and QoS of different users and flows. Dynamic FFR requires fast reuse set specific channel quality information over all reuse sets to take advantage of the time domain multi-user diversity gain and FFR interference avoidance gain simultaneously. Reuse set specific instantaneous channel quality information is obtained by combining AT feedback of an instantaneous non-reuse set specific CQI and a slow update of interference difference between reuse sets through vector CQI reporting.

18.5 Discussion

FFR is also applied on RL to reduce the interference from the most dominant interferers to a sector.. The transmission of handoff ATs could be orthogonalized with the sector being interfered.

MIMO users could potentially benefit more from the SNR improvements provided by FFR, since improved SNR translates to larger capacity gains over channels of higher dimensions.

FFR could be implemented as an intra-base station interference avoiding technique without global frequency planning, since ATs at sector edges often experience low SNR. FFR could potentially alleviate the intra-base station interference by orthogonalizing AT transmissions that were served by co-located sectors.

19 Subband scheduling

19.1 Introduction

When scheduling multiple terminals in frequency selective channels, system capacity can be increased by scheduling each terminal in a preferred subband based on its current channel frequency response. This section describes how the proposed design supports subband scheduling on the forward and reverse links.

19.2 Local hopping and channel trees

The subband size should provide enough frequency diversity to prevent performance degradation for fast moving ATs. Another implication of having too narrow subbands is a loss in trunking efficiency since there would be fewer candidate ATs to be scheduled per subband. Based on these considerations, the proposed system uses a subband size around 1.25 MHz. Furthermore, when an AT scheduled within a subband with a bandwidth assignment less than the entire subband, its assignment will hop 'locally' across this subband in order to maximize channel and interference diversity.

Based on the above guidelines, we can define channel trees with local hopping in the following manner. All parent nodes that host 8 base nodes (hence 128 subcarriers) will map to fixed subbands of 128 contiguous subcarriers. However, channels corresponding to the base nodes underneath this parent node will be hopping across the subband. This concept is illustrated in Figure 19-1.

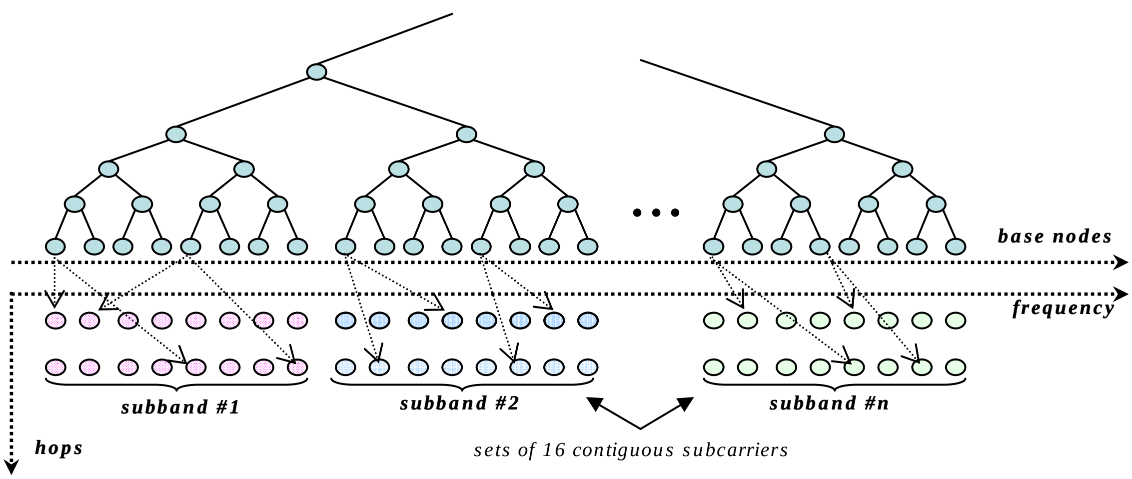

Figure 19-1 Channel tree structure to support subband scheduling.

In addition to the subband scheduling mode, the system can be operated in a diversity mode where base nodes of the channel tree hop across the entire band. The diversity mode may be preferred e.g. in sectors that serve predominantly fast moving users. In such cases, subband scheduling yields no gains while local hopping reduces the amount of channel and interference diversity. In the diversity mode,. The choice of diversity mode versus subband scheduling mode is implicit, through the choice of the appropriate channel tree advertised by the AP.

19.3 Subband feedback

In order to support subband scheduling, AT needs to provide feedback about FL channel properties relative to different subbands. The amount of feedback should balance gains in FL performance due to subband scheduling versus the RL overhead caused by feedback channels.. In the present design, we introduce a control channel called reverse subband feedback channel (R-SFCH). The minimum version of this channel includes a 4-bit subband index which is used by AT to indicate the preferred subband (with up to 16 subbands in 20 MHz). Based on this indication, the AP may schedule the AT over this subband with rate determined from the reverse channel quality indicator channel (R-CQICH).AP. An extended version of the R-SFCH will include a subband channel quality indicator in addition to the subband index.

20 MIMO

20.1 Introduction

The proposed system supports multiple-input multiple-output (MIMO) techniques that increase spectral efficiency through spatial multiplexing. This section describes the MIMO design used in the proposed system. It covers the different MIMO schemes, pilot structures, and overhead channels to support the design.

20.2 Data Channel Structure

The system uses the concept of effective antenna signaling at the base-station, i.e., the AP creates multiple beams using the set of physical antennas, each beam being referred to as an effective antenna. The different beams are generated to preserve the channel statistics as well as to transmit the same power from all physical antennas. For the first constraint, unitary matrices are needed to generate the spatial signatures. For the second constraint, the sum of the absolute values of the entries for every row is constant. One such signaling matrix would be a diagonal matrix, whose entries are phasors with random angels, multiplied by a DFT matrix. The beams are generated by multiplying the transmitted vector, of modulation symbols, by the signaling matrix. These beams can change slowly over frequency (using delay-diversity or cyclic delay-diversity techniques, for example) and over time[3]. These changes are such that the effect on the AT channel estimation algorithm is the same as that of channel variations. Moreover, each of these beams utilizes the different physical transmit antennas equally, thus ensuring that all the power amplifiers at the base station are used equally. SIMO transmissions use only the first effective antenna, while MIMO transmissions can use a subset or all of the available effective antennas. Since the number of effective antennas can be made smaller than the number of physical antennas, the AP can pick any trade-off between channel estimation overhead and available transmit diversity. Basically, the total number of effective antennas created dictates the maximum transmit diversity order that can be exploited. This number also dictates the amount of overhead needed to estimate the spatial channels.

In symbol rate hopping mode, the number of effective antennas is a sector wide parameter that can be adapted (long term) according to SNR and channel conditions. In block hopping mode, in general, a subset of the total number of effective antennas is used on any given tile. In the specification, the term “tile antenna” is used to denote an effective antenna on a particular tile. In this document, for simplicity, we only use the term effective antenna with the understanding that it refers to a tile antenna when we describe a process that occurs on a tile. The set of effective antennas in a given tile can be adapted (short term) to channel conditions since dedicated pilots are used. In both hopping modes, the number of modulation symbols simultaneously transmitted for a given packet is adapted (short term) to channel conditions.

[3] For the given example, changes in time amounts to changing the phasors’ angles in the diagonal matrix.

20.3 Pilot Structure

In this section we describe the pilot structure for symbol rate hopping and block hopping.

20.3.1 Symbol Rate Hopping

In symbol rate hopping, a common broadband pilot is transmitted from each effective antenna. Pilot subcarriers are present in every OFDM symbol, and the set of pilot sub-carriers in each OFDM symbol are spaced equally over the entire bandwidth to enable efficient channel estimation. A common pilot channel (F-CPICH) is transmitted from the first (SIMO) effective antenna and is used for SIMO demodulation. The F-CPICH subcarriers are staggered over groups of 2 OFDM symbols, i.e., if the F-CPICH occupies subcarriers {0, 16, 32, 48, 64, …} in the first OFDM symbol of a TTI, it occupies subcarriers {8, 24, 40, 56, 72, …} in the second OFDM symbol. The F-CPICH location changes in a pseudo-random fashion (subject to the constraint that it occupies an equally spaced set of subcarriers) every two OFDM symbols. This ensures that the pilot subcarriers from one sector do not always collide with the pilot subcarriers used by another sector.

An Auxiliary Pilot Channel (F-AuxPICH) is transmitted from the remaining effective antennas. The F-AuxPICH occupies subcarriers that are not occupied by the F-CPICH and equally spaced across the band. F-AuxPICH subcarriers from the different effective antennas are TDM'd with each other. For instance, if staggering is not employed, , if the F-AuxPICH subcarriers from effective antenna 2 occupy OFDM symbol 1, F-AuxPICH subcarriers from effective antenna 3 occupy OFDM symbol 2, F-AuxPICH subcarriers from effective antenna 4 occupy OFDM symbols 3, F-AuxPICH subcarriers from effective antenna 2 occupy OFDM symbol 4 (assuming that the AP has 4 effective antennas), and so on. This particular example is shown in Figure 20-1. The pilots from the non-SIMO antennas thus occupy a smaller bandwidth than the pilot from the SIMO antennas.

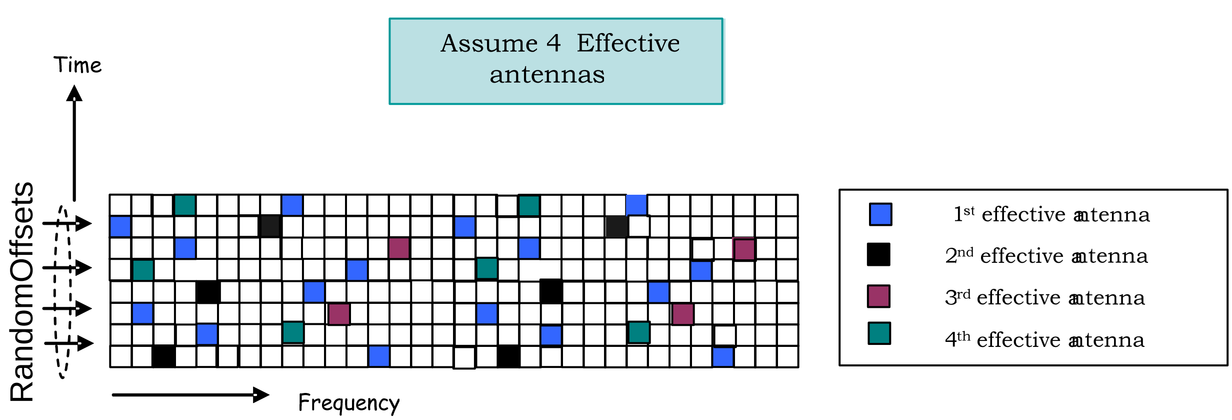

Figure 20-1 Pilot structure for the case of 4 effective antennas

20.3.2 Block Hopping

In block hopping mode, the assignment for a MIMO transmission consists of one or more tiles. Each tile is 16 contiguous subcarriers by 8 OFDM symbols. The pilot patterns, used to estimate the MIMO channel, are described in Section 12.3.2 . These patterns allow multiplexing pilots that correspond to different effective antennas. The channel estimates on the pilots are interpolated over the tile to obtain the estimates for every subcarrier, OFDM symbol, and effective antenna.

20.4 STTD Mode

The space time transmit diversity (STTD) is a diversity mode that is allowed by the specifications. The spatial multiplexing modes are described in Sections 20.5 through 20.7.

In STTD mode the AP is assumed to employ only two effective antennas. The STTD block code is signaled at the same subcarrier and two consecutive OFDM symbols. It is then required to have the same hopping pattern across two consecutive OFDM symbols. That is mandated by the specifications for the two hopping modes.

20.5 MIMO Design

Let the number of transmit antennas be M_t , the number of receive antennas be M_r , the number of effective antennas be M_e , and the number of modulation symbols simultaneously transmitted on a given subcarrier and OFDM symbol (a.k.a. spatial multiplexing order) be $M \le \min(M_e, M_r)$ [4].

The design supports two main MIMO modes that work for FDD and TDD, namely single codeword (SCW) and multiple codeword (MCW) designs. In TDD, a third mode namely pseudo-Eigen beamforming (p-EBF) is supported that exploits reciprocity of the forward and reverse link channels. The p-EBF scheme can be supported with both SCW or MCW modes and is discussed in Section 21

In the SCW mode, one codeword is transmitted in the frequency-space domain. A simple linear receiver is used to decouple the multiple transmitted modulation symbols, thus, it is not capacity achieving. More sophisticated receiver for SCW can significantly improve performance, especially at high SNR, on the expense of increasing complexity at AT.

In the MCW mode, multiple encoded streams of data are simultaneously transmitted. A successive interference cancellation (SIC) receiver is adopted. MCW with SIC is capacity achieving, hence optimal in performance. On the other hand, the SIC process can increase AT complexity as well as memory requirements. In addition, this mode requires extra overhead compared to SCW mode as will be described later.

In general, when transmitting M modulation symbols over M_e effective antennas, the M modulation symbols are cyclically shifted every subcarrier as shown in Figure 20-2. That is, the first M modulation symbols are transmitted on the first M effective antennas and first subcarrier in the assignment. The second M modulation symbols are transmitted on effective antennas $2 \text{ to } (M+1)$ and the second subcarrier in the assignment, and so on.

It remains to say that the design mandates the number of effective antennas used in a tile to be equal to the spatial multiplexing order, i.e., $M_e = M$, in block hopping.

[4] The constraint of $M \le M_r$ is necessary if linear receivers are used to decouple the incoming MIMO sub-streams.

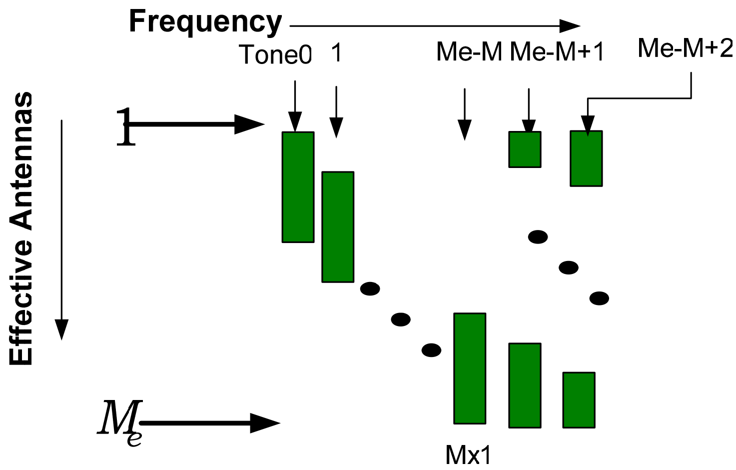

Figure 20-2 MIMO Layer Cycling

20.6 SCW Design

In this section we describe the SCW MIMO scheme.

20.6.1 Rate and Rank Prediction

The code rate, the constellation size, and the spatial multiplexing order (also denoted by rank) M are adapted to the channel. The AT runs a rank prediction mechanism and feeds back the rank value together with the corresponding channel quality indicator (CQI) value to AP. AP adjusts the transmitted power level, based on the power control loop and rank, and runs a rate prediction algorithm by which it chooses the packet format (PF). The PF, in addition to the fed back rank, define the data rate transmitted.

20.6.2 Transmitter Structure

The transmitter structure is shown in Figure 20-3. The input data stream is Turbo encoded using the selected code rate, and mapped to the selected QAM constellation. The stream of modulation symbols is then de-multiplexed to M parallel sub-streams. The M sub-streams are mapped to the physical antennas using the effective antenna signaling described in Section 20.2.

20.6.3 Receiver Structure

The receiver presented in this contribution runs a linear MMSE filter on the received samples to decouple the incoming M sub-streams[5]. The soft estimates of the modulation symbols are then fed to an LLR computer and the output is fed to a Turbo decoder. The receiver structure is shown in Figure 20-4.

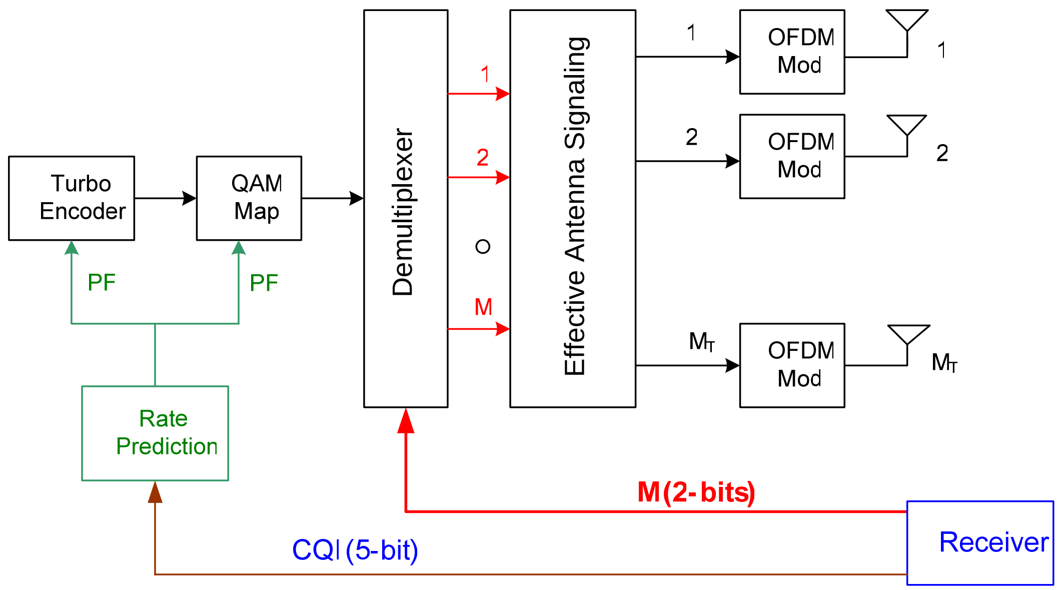

Figure 20-3 Transmitter structure for SCW MIMO

[5] As mentioned before, more sophisticated receiver structures can be used.

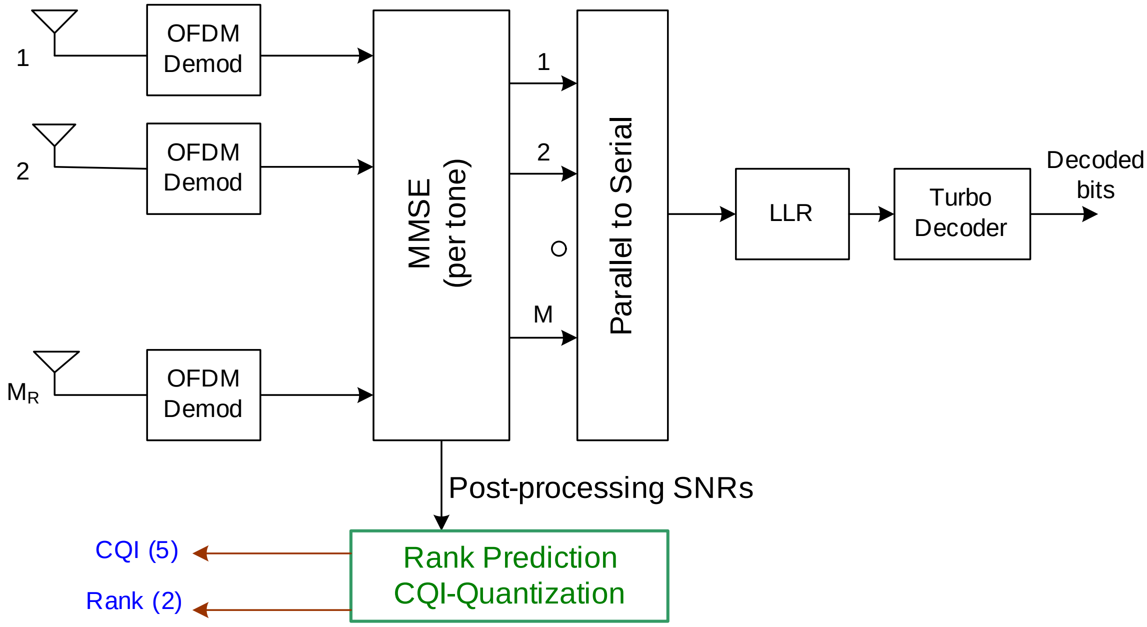

Figure 20-4 Receiver structure for SCW MIMO

20.6.4 SCW HARQ

H-ARQ for MIMO SCW is as described in Sections 6 and 7.

20.6.5 Feedback Channels

20.6.5.1 CQI Channel

The CQI and rank are reported to AP via a RL control channel denoted by the R-CQICH that is periodically transmitted. The CQI value is quantized to 5 bits and the rank to 2 bits.

20.6.5.2 ACK-NACK Channel

The receiver sends back a 1 bit ACK, on the R-ACKCH, if the packet is decoded and nothing if it is not (i.e. ON-OFF Keying).

20.7 MCW MIMO Design

In the MCW scheme M codewords (or packets) are transmitted in parallel. Each codeword is transmitted form all effective antennas to exploit the available spatial diversity.

20.7.1 Rate and Rank Prediction

The code rate and the constellation size on each of the M data sub-streams are adapted to channel. The AT runs a rank prediction algorithm by which it determines the value of M to be used. The AT also computes M CQI values, one for each data sub-stream, and feeds them back to the AP. The AP adjusts the transmitted power level on each data sub-stream, based on the power control loop and rank, and runs a rate prediction algorithm by which it chooses the code rate and constellation (the PF) for each data sub-stream.

20.7.2 Transmitter Structure

The transmitter structure is shown in Figure 20-5. Basically, M data packets are transmitted in parallel. The $m^{th}, m = 1, \cdots, M$, data packet is Turbo encoded using the m^{th} selected code rate, and mapped to the m^{th} selected QAM constellation. The modulations symbols corresponding to the m^{th} data packet is denoted hereafter by a layer. The M layers are then mapped to the physical antennas using the effective antenna signaling described in Section 20.2.

20.7.3 Receiver Structure

In this contribution, we employ a SIC receiver with linear MMSE filter to decouple the incoming M layers. The receiver attempts first to decode the first layer. The linear MMSE filter generates the soft estimate of the modulation symbols corresponding to the first layer and all subcarriers in the user's assignment. The different soft estimates are sent to an LLR computer, and the resultant LLRs are fed to the Turbo decoder. If the first layer is decoded properly (passes the CRC), the receiver regenerates a clean version of the modulation symbols corresponding to the first layer, multiplies each modulation symbol by the corresponding channel coefficient, and subtracts the contribution of the first layer from the received signal. The receiver then attempts to decode the second layer, if decoded, the receiver subtracts its contribution from the received signal, and so on. If at any point, one of the layers is not decoded, the receiver stops the decoding process and sends an R-ACKCH indicating the layers that got decoded. The receiver structure is shown in Figure 20-6.

20.7.4 BL HARQ

The HARQ scheme used is denoted by blanking layers (BL). Basically, if N_{dec} layers have been decoded at some transmission q , on the $(q+1)^{th}$ transmission the AP does not transmit any new codewords on the successfully decoded layers and only sends redundancy information on the $M - N_{dec}$ layers that have not yet been decoded. In doing so, the AP equally divides the available power per subcarrier on the outstanding layers.

20.7.5 Feedback Channels

20.7.5.1 CQI Channel

The CQI channel is send on the CQI RL control channel. AP feeds back M_e CQI values, one for each layer. In symbol rate hopping, if the rate prediction suggests supporting only $M \le M_e$ layers, then the AT sends zero CQI for the last $M_e - M$ layers. The CQI value per layer is quantized to 4 bits.

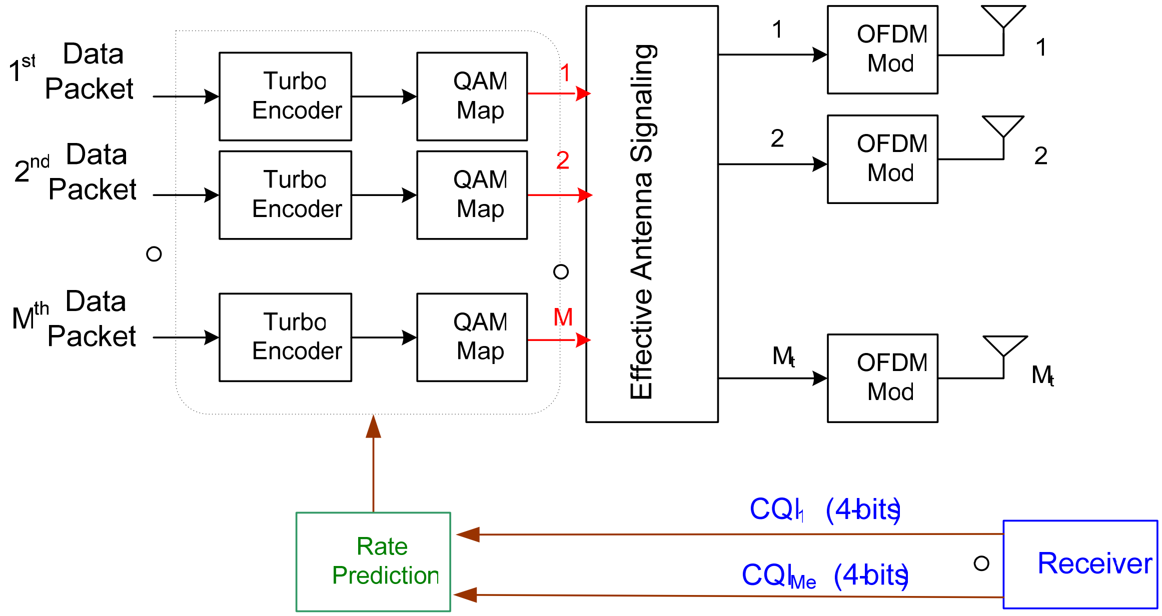

Figure 20-5 Transmitter structure for MCW MIMO

20.7.5.2 ACK-NACK Channel

M_t bits are used for R-ACKCH in MCW. Each bit corresponds to one of the transmitted layers and indicates whether or not the layer is decoded. ON-OFF signaling is used on each bit.

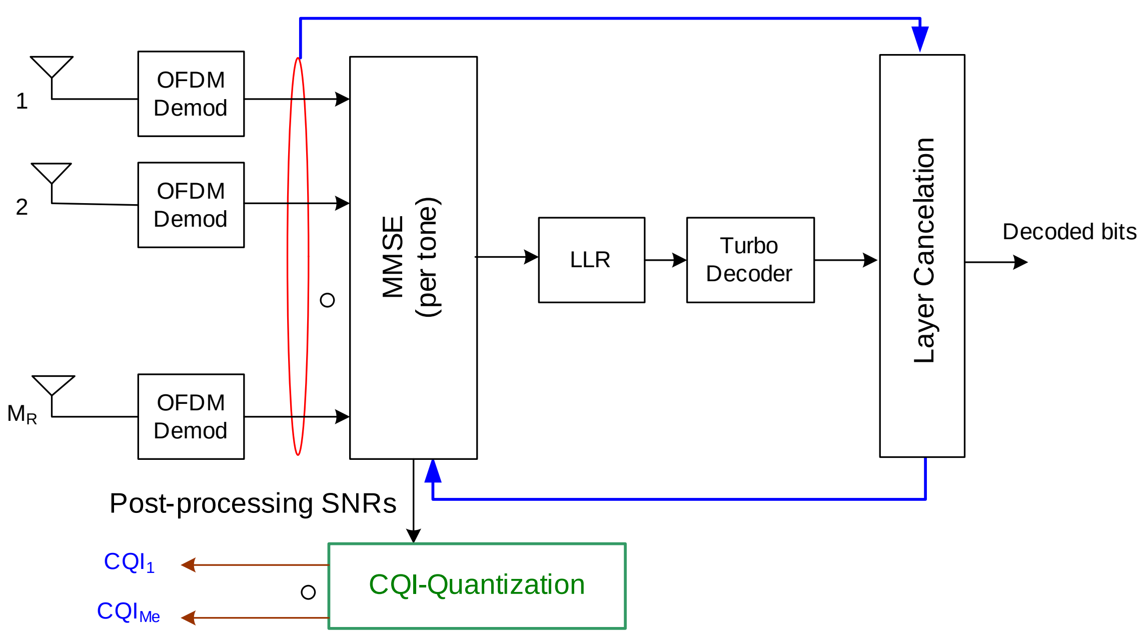

Figure 20-6 Receiver structure for MCW MIMO

21 Precoding

21.1 Introduction

In a TDD system, the forward and reverse link transmissions are on the same frequency region so that the reciprocity principle allows the estimation of the forward link channel from the reverse link channel. This enables the AP to extract transmit beamforming gain on the forward link when multiple antennas are available at the AP.

However, in a FDD system, the forward and reverse link transmissions are on widely separated frequencies. As a result, the forward link channel and the reverse link channel may fade independently. A direct consequence is that the reverse link channel estimates do not provide instantaneous channel knowledge of the forward link. Transmit beamforming gains, in an FDD system are possible by explicitly feeding back FL channel information over the reverse link and then using this information to transmit data in the preferred direction to every user. We refer to this technique as precoding. Precoding has value in a TDD system as well where the reverse link feedback may be utilized to enhance the forward link channel information available through reverse link pilots.

Precoding is used in the system in block hopping mode on the forward link data channel.

21.1.1 Precoding gains

In systems with multiple transmit and single receive antennas, precoding provides beamforming gain which manifests as improved capacity and coverage. In systems with multiple transmit and receive antennas (MIMO), precoding enables eigen-beamforming where multiple layers (refer Section 20.5) being transmitted to a user can be signaled on the eigen-beams of the channel. The resulting gains are especially noticeable when the number of layers (rank) transmitted is less than the number of transmit antennas at the AP. This is always true in asymmetric antenna scenarios, where the number of transmit antennas at the AP is larger than the number of receive antennas at the AT.

21.1.2 Design challenges

The precoding gains depend on timely, accurate FL channel information being fed-back on the reverse link. Given that the channel is frequency selective (delay spread) and space selective (multiple transmit antennas), the information fed-back also has to be frequency and space selective. Furthermore, there has to be a means to estimate the frequency selective forward link channel for all users over all transmit antennas even when a user is not scheduled for transmission. The resulting overheads to support precoding can be excessive.

The latter challenge is addressed by using a low overhead broadband pilot channel on the forward link, that enables accurate estimation of the channel information required for precoding. The reverse link overhead requirement is mitigated by trading off precoding gain with reduction in overhead through quantization techniques. We include details of the proposed solution below.

21.2 Proposed architecture

21.2.1 Pilot design and feedback frequency

Broadband CQI pilots are transmitted on the forward link approximately every 5.5 ms and across all available physical antennas. The pilot channel allows frequency and space selective channel estimation at the AT that provides the information necessary for the AT to compute the FL channel information to be fed-back to the AP. The AT transmits this information over the CQI channel on the reverse link control segment.

21.2.2 Feedback and transmission

A key design requirement is to properly trade-off the precoding gains with the feedback overhead on the RL control segment. The overhead is chiefly driven by the need to capture frequency and space selectivity. Feeding back the exact coefficients of the beams on every transmit antenna is quite expensive. Some sort of quantization is therefore necessary and is achieved using a pre-defined codebook that contains entries that define the spatial processing to be applied at the transmitter. Currently up to 16 codebooks may be defined. The codebooks may be sector specific and selected based on antenna configuration, terrain profile, etc. Each codebook contains up to 64 entries that provide support for diversity transmission, beamforming to line-of-sight users (vectors), MIMO precoding (matrices), SDMA and MIMO-SDMA. The AT computes the preferred transmission mode from among this pre-defined set of entries over each subband. A sub-band index and precoding index for that subband are signaled to the AP using the R-SFCH and R-BFCH channels on the reverse link control channel.

The desired precoding index is utilized while transmitting to the AT. Some of the codebook entries may be matrices that are defined to support MIMO precoding. In MIMO mode, the AT performs rate and rank prediction using this set of pre-defined matrices. The desired matrix index (6 bit quantity) and the rank (2-bit quantity allowing rank up to 4 to be signaled) are fed back on the reverse link control segment. Together, the two quantities provide the information required by the AP for precoding. Consider for example a sector with four transmit-antennas and a precoding and MIMO capable user with four receive antennas. The user indicates the preferred precoding matrix, the rank ‘r’ and CQI through the reverse link control channel. The AP receives this information and adapts the transmission to the AT by using the precoding matrix signaled by the AT. The ‘r’ layers transmitted to the user are modulated on the first ‘r’ columns (beams) of the preferred precoding matrix.

21.2.3 Transparent operation in block hopping

The precoding directions employed by the AP do not have to be signaled to the AT. This is because precoding mode is allowed only with block hopping mode (refer Section 12.3) which employs dedicated pilots for channel estimation. Since the pilots and data are both modulated on the same precoding directions, the choice of these directions is transparent to the AT.

22 Beamforming for TDD

22.1 Introduction

In a TDD system, the forward and reverse link transmissions are on the same spectrum so that the reciprocity principle allows the estimation of the forward link channel through pilot transmissions on the reverse link. These transmissions could be explicit pilots, i.e., known pilot sequences, or reverse link control channels that may be used as pilot symbols post demodulation and allow the estimation of the space-selective (over multiple receive antennas) channel.

In systems with multiple transmit and single receive antennas, beamforming provides increased SINR. When multiple transmit and receive antennas (MIMO), are available, eigen-beamforming can be used where multiple layers (refer Section 20.5) being transmitted to a user can be signaled on the eigen-beams of the channel. The achievable gains (over non-beamformed transmission) are especially noticeable when the number of layers (rank) transmitted is less than the number of transmit antennas at the AP. This is always true in asymmetric antenna scenarios, where the number of transmit antennas at the AP is larger than the number of receive antennas at the AT.

However, availability of forward link channel information may limit the ability to realize eigenbeamforming gains. The AP obtains forward link channel information from reverse link pilot transmissions. With only one transmit chain at the AT, the pilot transmissions allow channel estimation to only one AT antenna. When the AT has multiple receive antennas, the information available at the AP is therefore incomplete. Secondly, in frequency selective channels, the reverse link pilot transmissions need to provide broadband channel information. In the following we describe the system solutions to these challenges.

22.2 Proposed architecture

22.2.1 Beamforming transmission

When the AT is capable of demodulating only a single layer, or when only one layer is transmitted to a MIMO capable AT, the AP may use a transmit beam that is matched to the available forward link channel estimate.

22.2.1.1 Pseudo-eigenbeamforming for MIMO transmission

As noted previously in this section, the reverse link pilots can be transmitted through only one transmit antenna. The assumption is that in most scenarios the AT has only one transmit PA available. Therefore, even under ideal conditions only partial information is available for determining the beamforming weights. When multiple layers are transmitted to the AT, the first layer is signaled using the beamforming described above. The remaining layers are signaled on beams that are in the subspace orthogonal to the beam used on the first layer. We refer to this form of MIMO transmission as pseudo-eigenbeamforming.

22.2.1.2 Multiplexing support

Efficient channel estimation support for demodulating beamformed transmission is provided by beamforming the forward link pilot transmissions as well.

The block hopping mode (Section 12.3) with dedicated pilot patterns is ideally suited to support beamformed transmission and is the mode used with TDD beamforming. This mode provides flexibility (for example frequency and user selective beamforming), robustness and transparency (AP does not need to signal which beams are being used) of operation between the AP and the AT when using beamforming.

22.2.2 Feedback

The reverse link control channel carries the CQI channels (R-CQICH, R-SFCH and R-BFCH) as well as an explicit pilot channel (R-PICH) described in Section 14 All channels are transmitted using CDMA sequences over the sub-bands allocated to the RL control channel segment which is available approximately every 5.5 ms. The first set of channels may be used as pilot transmissions.

22.2.2.1 Frequency selectivity

The reverse link control channel occupies a number of contiguous sub-bands, each sub-band being of size 1.25 MHz. These sub-bands hop over the carrier, in a pre-defined fashion so that broadband information over the entire frequency range is possible by collecting the reverse link pilots over successive transmissions.

22.2.3 Precoding

Previously we noted that the reverse link pilot transmission does not provide complete forward link information when multiple receive antennas are available at the AT. The precoding support provided through the R-BFCH and R-SFCH can be used to improve performance in this scenario. Since the full forward link channel, and hence also the eigen-beams can be estimated at the AT, the R-BFCH may be utilized to feedback desired precoding directions (refer Section 21). This information along with that provided by the reverse link pilots can be utilized by the AP to properly adapt forward link beamforming. This may be achieved, for example through intelligent (instead of random) selection of the subspace basis vectors when the number of transmitted layers is smaller than the number of transmit antennas. Since this is nearly always the case with a large number of transmit antennas, precoding support provides gains even in a TDD environment.

23 Space Division Multiple Access

23.1 Introduction

Space-division multiple access (SDMA) on the forward link is an advanced transmission technique where multiple users are signaled on the same time-frequency resources. A key characteristic of SDMA is the opening up of new dimensions at the expense of reduced signal to interference and noise ratio (SINR).

23.2 Intra-sector interference management

When multiple transmit antennas are available at the AP, they can be used to mitigate intra-sector interference in SDMA mode. This is achieved by transmitting simultaneously to the overlapping users using properly defined beams for each user. These beams are dependent on the set of overlapping users and in particular are dependent on the spatial channels (across multiple transmit antennas) of the overlapping users. The intra-sector interference in turn is dependent on the beam(s) of the overlapping user(s). As an example, consider two overlapping users, each with a single receive antenna and orthogonal spatial channel (across multiple transmit antennas) vectors. A simple matched filter transmission at the AP, i.e., the beam for each user is the complex conjugate of their respective channel vector, suppresses intra-sector interference. This is an ideal scenario because there is no intra-sector interference at all and full SDMA gain may be leveraged. On the other hand, if the user channels (spatial signatures) are similar then the beams used to transmit to the users are likely to be similar as well causing high intra-cell interference.

23.2.1 User grouping

The above discussion suggests that user grouping would play an important factor in realizing SDMA gain, i.e., it is preferable to overlap users with sufficiently different spatial signatures to mitigate intra-sector interference. This in turn suggests the need for availability of forward link channel information at the AP. In an FDD system, similar to the requirement for precoding support, this is provided through feedback channels on the reverse link, while in a TDD system this information is available through reverse link pilots.

23.3 Proposed architecture

In a WAN type environment, SDMA gain is achievable with proper intra-cell management techniques. A fundamental requirement is the provisioning of forward link channel information so as to enable the AP to group users into multiple sets. Additionally, the AP also needs to compute and schedule overlapping users on appropriate sets of beams. Before outlining the solution to these challenges, we first identify a third problem that has bearing on the design, namely the need to support *hybrid operation.* In any WAN type environment it is necessary to support coexisting users in different modes, i.e., the need is to simultaneously support users in SDMA mode, precoding mode (refer Section 21), MIMO mode (refer Section 20), etc., while also providing support for various broadcast control channels. Furthermore, the bandwidth requirements of these modes may change frequently. The challenge therefore is to enable design scalability to provide flexibility in supporting the bandwidth requirements of the scheduled modes. Thus, we identify the three main challenges to be user clustering, beam computation and feedback and support of hybrid modes.

23.3.1 Feedback and user clustering

In TDD, user clustering may be based on FL channel information derived from the reverse link pilots. In FDD, the clustering is based on information provided by a reverse link feedback mechanism. In the following, we describe the feedback and user clustering approach in FDD.

The same mechanism defined for precoding support may be reused for SDMA support. As outlined in the precoding section, the codebook has 64 entries indexed using a 6 bit vector. The codebook entries may be segmented into multiple sets, with one set for precoding transmission, another set for SDMA transmission and remaining entries indicating default transmission (more in this later in the section). The set of entries corresponding to SDMA transmission are segmented further into distinct clusters. Therefore each SDMA cluster corresponds to a set of matrices that can be used for transmission by users in SDMA mode. Only users corresponding to different SDMA clusters are allowed to be overlapped.

23.3.1.1 Feedback

In SDMA mode, the AT is required to indicate the preferred beam (from an SDMA cluster) as well as feedback the CQI associated with the preferred beam.

The CQI pilot channel (F-CPICH), scheduled every 5.5 ms in block hopping mode, is utilized to estimate the broadband frequency domain channel response on all physical transmit antennas. The signal quality from the entries of the codebook can now be computed. This procedure is very similar to the signal quality computation in the precoding mode.

The one key difference between precoding CQI computation and CQI computation in SDMA mode is in the computation of the interference level. The interference is driven by both inter-sector interference as well as intra-sector interference from a simultaneously scheduled overlapping user. The exact intra-cell interference depends on the beams and power on the users within the overlapping SDMA clusters. The AT has no knowledge of an interfering user (or the beams allocated to that user) since the users are scheduled at the AP and has to make a best effort guess as to the potential interference from overlapped users. The preferred beam index and SDMA CQI is signaled through the R-BFCH . Similar to the precoding case, the reporting is based on a particular subband.

23.3.1.2 User clustering

Every user in SDMA mode, reports a preferred beam index that is contained within a particular SDMA cluster within the codebook. All users corresponding to the same SDMA cluster are placed into the same group. Users within a group are scheduled so that they are always orthogonal to each other, i.e. they are not allowed to overlap. The rationale for this is the fact that the beams within the same SDMA cluster have similar spatial characteristics; therefore, users preferring these beams are also likely to have similar spatial characteristics and should not be overlapped. This is achieved by scheduling all users within a group on the same tree (refer Section 9.2)

23.3.2 Hopping

SDMA scheduling is based on a channel tree (section 9.2) containing identical sub-trees that map to the same set of resources. Users within the same SDMA cluster are grouped together and a sub-tree is assigned per group. Hop-ports within each sub-tree map to disjoint subcarriers. Since all sub-trees map onto the same frequency region, users scheduled across sub-trees may overlap with each other. The permutation mapping hop ports to frequency is assumed to be the same across all sub-trees, i.e., if the hop ports assigned to a user on one sub-tree correspond to the hop ports assigned to a different user on another sub-tree then they are scheduled on the same frequency resource. One of these sub-trees is the primary sub-tree and is used to schedule users that are not in SDMA mode. The hop ports corresponding to those allocated to non-SDMA users on the primary sub-tree are not used in any of the other sub-trees. Constraining the permutations to be the same allows flexibility in scheduling and in particular enables support of the hybrid scenario.

23.3.2.1 Hybrid mode support

The design allows for a multiplicity of modes, SDMA mode where users are overlapped, and non-SDMA modes such as precoding, MIMO, and diversity mode, where users are not overlapped. The last may be used to support channels with enhanced robustness requirements, for example broadcast control channels. Furthermore, since the bandwidth requirements of these modes could change rapidly, there is a need for a high degree of scalability. This is achieved using the tree and hopping structure defined above and is the key reason for preferring the same permutation mapping hop ports to subcarriers across all sub-trees.

Channels that prefer non-SDMA mode may be assigned hop ports (and hence a frequency range) on the primary sub-tree. The hop ports on all other sub-trees that correspond to this frequency range are removed from the available scheduling resource. Since there is a one-one mapping between the channel IDs on different sub-trees, any users scheduling thereafter in SDMA mode will not overlap with the channels in the non-SDMA mode. Given that the granularity of each region is the same as that of a base node (16 carriers), the bandwidth allocated to each mode may be increased or decreased as required. This provides the granularity requirements required by the design.

24 Scalable Bandwidth

24.1 Introduction

For the case of large deployment bandwidths, it is desirable to support ATs that are not capable of demodulating the entire bandwidth, while maintaining the advantages of a large bandwidth deployment such as high peak rates for ATs that are capable of demodulating the entire bandwidth. The MBWA system supports such a feature through the MultiCarrierOn mode. In this mode, the entire bandwidth is divided into multiple carriers, each spanning 512 subcarriers. A 20MHz deployment, which spans 2048 subcarriers, would therefore contain 4 carriers. Each carrier has guard subcarriers at its edge, which are referred to as quasi-guard subcarriers.

Each carrier in this mode behaves like an independent 512 subcarrier system. Independent overhead channels are transmitted over each carrier, and each carrier contains its own forward and reverse link control resources (SSCH and CDMA segment respectively). We define different classes of terminals, each of which is capable of demodulating (and transmitting on) different numbers of carriers. Note that when we say an AT is capable of demodulating only one carrier, we mean this not only in terms of baseband processing power but also in terms of the RF front end. For example, the AT could have an ADC that only samples 5MHz of bandwidth.

An AT that is capable of simultaneously demodulating all the carriers can be dynamically scheduled on one or more of the carriers on a PHY Frame by PHY Frame basis. Since an AT can be simultaneously scheduled on all the carriers, the peak rates in this system are similar to those obtained in the MultiCarrierOff mode. ATs that are capable of demodulating a subset of the carriers can be switched from one subset to another as fast as one PHY Frame. This mode thus supports very fast switching times, which are useful for load-balancing purposes and in order to get statistical multiplexing gains. The modulation symbols received from all of the carriers are combined into a single MAC layer frame.

24.2 Acquisition design

Each carrier carries its own acquisition preamble in MultiCarrierOn mode. The overhead channels F-pBCH0 and F-pBCH1 are also carried independently on each carrier. The total number of carriers in a multi-carrier group (denoted by NumCarriers) and the index of a given carrier within the group (denoted by CarrierIndex) are carried in the F-pBCH0 channel of that carrier. Here, a multi-carrier group refers to the entire FFT bandwidth.

In MultiCarrierOff mode, the acquisition PN sequences in a given superframe are based on the quantity PilotPN in Asynchronous mode, and on the quantity PilotPhase = PilotPN + SuperframeIndex mod 4096 in SemiSynchronous mode. In MulticarrierOn mode, we modify this setup so that the acquisition PN sequences in a given superframe are based on the quantity PilotPN + CarrierIndex mod 4096 in Asynchronous mode, and on the quantity PilotPhase + CarrierIndex mod 4096 in MultiCarrierOn mode.

Since the set of acquisition PN codes in a superframe is the same as that for a 512 subcarrier deployment (and moreover they cycle the same way in SemiSynchronous mode), the acquisition algorithm (including search complexity and search time) for an AT capable of demodulating only one carrier is the same as that for a 512 subcarrier deployment. (An AT capable of demodulating the entire bandwidth can improve its performance using the PN codes in all the carriers.) Once the acquisition PN codes have been detected, the AT then goes on to read the F-pBCH0 overhead channel. This channel carries the quantities SuperframeIndex and CarrierIndex, using which the AT can determine the PilotPN of the sector.

The AT can then go on to read the remaining overhead channels in the carrier it has acquired, following which it is ready to access the system.

24.3 Access design

One of the parameters that is advertised on the overhead channels is whether or not ATs are allowed to access the system on that carrier. The AT picks one of the allowed carriers autonomously and sends an access probe on that carrier. Prior to doing this, the AT has to read the overhead parameters on the chosen carrier. Note that even an AT that is capable of demodulating all carriers will access the system on only one carrier.

On receiving the access probe, the AN initially assumes that the AT is capable of demodulating only one carrier. Therefore it sends an access grant on the same carrier that the AT sends its access probe on. Meanwhile, the AN locates the AT's session which contains information on how many carriers the AT is capable of demodulating. (If a session is not found, it will be negotiated assuming that the AT is restricted to demodulating only one carrier.) Once the AN knows the AT capability, it can start scheduling the AT on multiple carriers.

24.4 FL signaling design

As mentioned previously, each carrier in MultiCarrierOn mode contains bandwidth reserved for the Shared Signaling Channel (SSCH). This channel carries assignments (both FL and RL) and acknowledgements (for RL traffic) for channels in the same carrier. The assignment management logic works independently on each carrier. If an AT receives assignments in multiple carriers, its assignment is simply the union of its assignments on all carriers. This enables the AT to potentially be assigned on all carriers, thus achieving high peak rates.

If an AT assignment on the RL spans multiple carriers, then the FL ACK for that channel is carried on one of these carriers, picked in a deterministic manner.

The AN can ask the AT to change the set of carriers it is demodulating using the Change Carrier Block (CCB), which is carried as part of the SSCH. This block is only defined in MultiCarrierOn mode. On initial access, the AT is required to monitor only the carrier on which it sends its access probe. This initial carrier can then be modified or expanded to a target set of carriers by the AN using the CCB. Note that the set of carriers that the AT listens to is a network-wide parameter and is not restricted to a single AN (i.e., this set is assumed not to change on handoff). The CCB becomes active one PHY Frame after it is transmitted, thus achieving very fast switching times between the carriers.

Prior to sending the CCB, the AN has to ensure that the AT has the overhead parameters for the carrier that it is switching to. This can be ensured by sending these parameters to the AT using a unicast message. This needs to be done very infrequently if the AN has long expiry timers on the overhead messages.

24.5 RL signaling design

As in the FL, each carrier has bandwidth reserved for the RL control channels, namely the R-ACKCH and the CDMA control segment. The R-ACKCH in each carrier is used to acknowledge FL channels in that carrier. If an AT is scheduled on multiple carriers on the FL, one of them is deterministically picked for R-ACKCH transmission.

The CDMA control channel is assigned independently and the AT transmits the CQI channel on all carriers that it is demodulating. The parameters governing CQI transmission (e.g., CQI reporting interval) are determined independently on each of the carriers. The same holds for the R-BFCH and the R-SFCH as well. The R-REQCH is transmitted on only one of the carriers, which is determined by the CCB.

Power control also runs independently on each carrier. The CQI in each carrier is power controlled using the power control bits transmitted on that carrier. Moreover, the OSICH transmission is also independent on each carrier, which results in the data being independently power controlled on each carrier. If an AT is assigned channels in multiple carriers, it shall use the power spectral density determined by the power control procedure in each carrier, in order to modulate subcarriers in that carrier.

25 Inter-Frequency and Inter-Radio Access Technology Handoff

25.1 Motivation

The MBWA system may be deployed in more than one frequency. At the frequency boundaries of a multi frequency deployment, there is a need to perform handoff during Idle and Connected Mode to other frequency system. Furthermore, the MBWA system may be deployed along with other Radio Access Technologies (RAT). The other RATs include but are not limited to cellular systems such as CDMA2000, WCDMA, GSM/GPRS/EDGE and WLAN systems like 802.11. There is a need to perform Idle and Connected Mode handoffs to these other RATs. In addition, there is a need to receive Page messages for other RAT while in Idle or Connected Mode in the MBWA system. This may be needed if other RAT provides a service that is unavailable in MBWA system. The other RAT may or may not be synchronous to MBWA system.

25.2 Inter-Frequency Handoff

The MBWA system supports both Idle and Connected Mode Inter-Frequency Handoff. The solution uses concept of Active Set Management. The L1 Handoff section describes the Active Set Management. To facilitate Inter-Frequency Handoff, the Active Set concept is extended to include members from one or more frequencies. This means that Active Set consists of Sectors from one or more frequencies. The Sector from different frequencies may be synchronous or asynchronous with respect to each other. To facilitate adding another frequency Sector into the Active Set, MBWA provides:

- Ability for AN to specify other frequency neighbors in the SectorParameters message (specified in Overhead Messages Protocol)
- Ability for AT to report other frequency Sector Pilot strength in the PilotReport message (specified in the Active Set Management Protocol)
- Ability for AN to specify other frequency members in ActiveSetAssignment message (specified in the Active Set Management Protocol)

In order to report other frequency Sector Pilot strength, the AT needs to take measurements. In Idle mode, it is easy to do so as the receiver is available for other frequency measurement (assuming slotted operation in Idle Mode). In order to report Pilot strength in Connected Mode, either dual receivers or temporary tune-away mechanism is needed. Since one can not always assume availability of dual receivers, a tune away mechanism is provided. Furthermore, similar tune-away mechanism is also needed for inter-RAT handoff and listening for Pages for another RAT that may be asynchronous to the MBWA system. The following section describes a generic tune away mechanism that provides solution to these requirements.

25.2.1 Tune Away Mechanism

The Default Connected State Protocol provides a tune-away mechanism. The tune away mechanism consists of a tune away schedule and a tune away message. The TuneAwaySchedule*N* attribute provides a way to negotiate tune away schedule(s) between AT and AN. The tune away schedule concept is depicted in Figure 25-1.

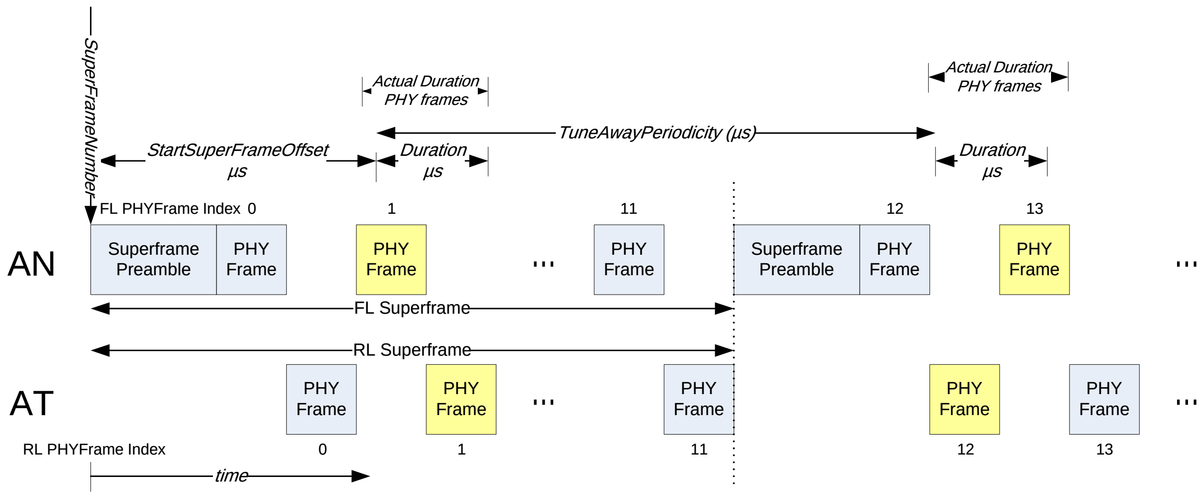

Figure 25-1 MBWA TDD TuneAway timeline

In tune away schedule, it is assumed that the first tune away occurred during the Superframe defined by SuperframeNumber provided in the TuneAwayScheduleN attribute. Furthermore, refined time of first tune away is StartSuperframeOffset microseconds from the beginning of the Superframe identified before. The TuneAwayDuration is how long in microseconds the AT and AN tune away for measurement/monitoring of other systems. The TuneAwayPeriodicity determines the time between the start of successive tune away in units of microsecond. The AT can negotiate one or more tune away schedules. More than one schedule may be needed, for example, to monitor pages of other RAT and at the same time tuning away for inter-frequency handoff in MBWA system. The detailed specification of the TuneAwayScheduleN attribute is in the MBWA system specification.

The tune away control mechanism provides two functions:

- Enable/Disable tune away
- Provides time correction to the tune away schedule

The AT can enable or disable tune away at any time. Furthermore, AT can enable or disable one or more schedules at the same time. The time corrections are needed for a time critical tune away to receive Page messages for a system that is asynchronous to MBWA system. For example, whenever a new Sector is added to Active Set, the AT has to provide a time correction called SectorOffset in units of microseconds to correct time such that it tunes away at the right time in other system to receive a Page.

The TuneAwayRequest and TuneAwayResponse messages in the Default Connected State Protocol provides a mechanism to reliably enable/disable tune away and provide time correction for any Sector in the Active Set. The detailed specification of the TuneAwayRequest and TuneAwayResponse messages is in the MBWA system specification.

At the beginning of a TuneAwayDuration, the AT and AN implicitly expire all RL/FL assignments, and suspend transmission on both links. The AT tunes away to measure/listen to the other system/frequency. At the end of the TuneAwayDuration, if FL/RL resources are needed, they are explicitly assigned to the AT. Note that the tune away mechanism is the same for both TDD and FDD designs.

25.3 Inter-RAT Handoff

The MBWA system supports both Idle and Connected Mode Inter-RAT Handoff. Mechanisms are provided to facilitate handoff from MBWA system to other RAT.

The solution uses the same tune away mechanism as described for the Inter-frequency handoff to measure other RAT Pilots. In addition the SectorParameters message in the Overhead Messages Protocol provides ability to send other technology Neighbor List. These two mechanisms provide AT with ability to find other RAT access network in the neighborhood, and measure the Pilots for the same. The detailed specification for the SectorParameters message is in the MBWA system specification.

25.4 Reception of Pages for other RAT

The MBWA system supports reception of Page messages for other RAT. There are two distinct mechanisms provided. 1) Tune away mechanism to receive pages for other systems. 2) InterRATProtocol in the Session layer that provides sending InterRATBlob message from AT or AN. The first mechanism is useful when MBWA system has no integration in the core network with the other RAT. Hence the only way to get a Page message from other RAT is by listening to its Paging Channel. As described before, the tune away mechanism supports tuning away for listening to paging channels at very specific time in other RAT, either synchronous or asynchronous to the MBWA system. The InterRATBlob message provides ability to encapsulate L2 signaling for other technologies within MBWA L2. This is useful when the two technology core-networks have established communication link such as L2TP.

26 Embedding Other PHY

The MBWA system has a provision for reserving some bandwidth in order to support a different physical layer on the forward link. This physical layer could, for example, be a broadcast system utilizing single-frequency network technology.

The F-pBCH0 channel specifies the number of interlaces as well the number of subbands per interlace that are reserved for the new physical layer. Thus the granularity of reservation is quite low (about 4% in the case of 5 MHz FDD). The interlaces as well as the sub-bands reserved for the new physical layer are chosen in a contiguous manner. The hopping pattern on the FL is chosen so as to avoid the reserved sub-bands.

At least one sub-band on each interlace is left unreserved, so that control signaling can be maintained in order to support the reverse link. The control signaling that is required consists of assignments, acknowledgements and power control bits.

27 Conclusion

In this document, we have presented an overview of a system designed for mobile broadband wireless access. The PHY and MAC have been designed to provide true broadband experience in the challenging mobile wireless environment. The system supports different levels of QoS and seamless connectivity for users across cells in a wide area environment. The PHY employs advanced transmission techniques and features a system design that supports implementation of those features with low system overheads. A key feature is the ability to simultaneously support terminals employing techniques such as MIMO, beamforming, precoding, SDMA, and subband scheduling. This enables the system to adaptively choose the appropriate transmission techniques for different terminals based on their propagation environments. All the system features described in this document make the current proposal a good choice for the new mobile broadband wireless access standard.