Project	**IEEE 802.16 Broadband Wireless Access Working Group**	
Title	**MAC Proposal for IEEE 802.16.1**	
Date Submitted	**1999-12-24**	
Source	Dr. James F. Mollenauer Technical Strategy Associates 37 Silver Birch Road Newton, Massachusetts, 02468	Voice: 617-244-0077 Fax: 617-244-0077 E-mail: jmollenauer@TechnicalStrategy.com
	Ken Stanwood Ensemble Communications, Inc. 6256 Greenwich Dr., Ste 400 San Diego, CA 92122	Voice: 858-404-6559 Fax: 858-458-1401 E-mail: ken@ensemblecom.com
	Jay Klein Ensemble Communications, Inc. 6256 Greenwich Dr., Ste 400 San Diego, CA 92122	Voice: 858-404-6544 Fax: 858-458-1401 E-mail: jay@ensemblecom.com
	Brian Petry 3Com Corp. 12230 World Trade Dr. San Diego, CA 92128	Voice: 858-675-8533 E-mail: brian_petry@3com.com
Re:	This is a MAC protocol proposal for "Phase II" of the 802.16 MAC protocol selection process. It is in response to the call for contributions as input to 802.16 Session #5.	
Abstract	The proposal described herein describes a MAC protocol that: • Supports the transport of diverse traffic types simultaneously (TDM, variable- and fixed-length PDU) • Maximizes capacity of the air link • Provides a commercially viable network for system operators • Uses well-understood technology • Supports FDD (both full and half-duplex) and TDD • Is responsive to varying bandwidth demands The MAC protocol closely resembles cable modem MAC protocols (e.g., DOCSIS and 802.14) yet addresses issues important to BWA systems.	
Purpose	The 802.16 Working Group should consider this MAC protocol proposal at Session #5.	
Notice	This document has been prepared to assist the IEEE 802.16. It is offered as a basis for discussion and is not binding on the contributing individual(s) or organization(s). The material in this document is subject to change in form and content after further study. The contributor(s) reserve(s) the right to add, amend or withdraw material contained herein.	
Release	The contributor acknowledges and accepts that this contribution may be made public by 802.16.	
IEEE Patent Policy	The contributor is familiar with the IEEE Patent Policy, which is set forth in the IEEE-SA Standards Board Bylaws <http://standards.ieee.org/guides/bylaws> and includes the statement: "IEEE standards may include the known use of patent(s), including patent applications, if there is technical justification in the opinion of the standards-developing committee and provided the IEEE receives	

assurance from the patent holder that it will license applicants under reasonable terms and conditions for the purpose of implementing the standard."

Table of Contents

Chapter

1

1 Introduction

1.1 Scope

The purpose of this document is to specify the TC and MAC layers of the 802.16.1 Broadband Wireless Access System. This document:

- Provides a general description of the Broadband Wireless Access System architecture.

- Describes the aspects of the physical layer (PHY) pertinent to access control.

- Describes the Transmission Convergence (TC) protocols and data exchange sequences between the Base Station (BS) and the Customer Premise Equipment (CPE).

- Defines the services offered by the TC to the Media Access Control (MAC).

- Describes the addressing scheme for mapping services destinations.

- Describes the uplink and downlink framing and formatting.

- Describes the MAC packet format and the mapping of MAC entities to the physical layer.

- Describes the services offered by the MAC to higher layers.

- Describes the MAC protocols and data exchange sequences between the BS and the CPEs.

- Defines the MAC messages, and their content, necessary to implement the MAC protocols.

The purpose of this document is not to define the design of the TC or MAC hardware or software. Nor is it, therefore, the purpose to map the elements of the TC or MAC to any specific hardware implementation. Additionally, this specification does not specify all the convergence subprocesses, which sit on top of the MAC. Some are specified in Chapter 4.

1.2 Definitions

HL-MAA MAC Domain

The portion of the media under the control of an instance of the HL-MAA sublayer of the MAC. Contains one or more LL-MAA MAC domains.

LL-MAA MAC Domain

The portion of the media under the control of an instance of the LL-MAA sublayer of the MAC

Physical Channel

For TDD: a frequency, sector pair
For FDD: a downlink frequency, uplink frequency, sector trpilet

1.3 Acronyms

TDD	Adaptive Time Division Duplexing
BS	Base Station
CG	Continuous Grant
CPE	Customer Premise Equipment
CS	Convergence Subprocess
DAMA	Demand Assign Multiple Acess
DES	Data Encryption Standard
DL	Down Link
FDD	Frequency Division Duplex
H-FDD	Half-duplex FDD
HL-MAA	High Level Media Access Arbitration
LL-MAA	Low Level Media Access Arbitration
MAC	Media Access Control
MTG	Modulation Transition Gap
PI	PHY Information element (600 bits)
PHY	Physical layer
PM	Poll Me bit
PPM	Priority Poll Me bit
PS	PHY Slot
QoS	Quality of Service
RS	Reed-Solomon
SAP	Service Access Point
TC	Transmission Convergence
TDD	Time Division Duplexing
TDM	Time Division Multiplex
TDU	TC Data Unit (55 bytes)
TRGT	Tx/Rx Transmission Gap
UL	Up Link

1.4 Applicable Documents

Date	Document No.	Name

Chapter
2

2 General Description

2.1 Overview

The function of the MAC layer in a shared-medium network is to deal with the fact that the physical medium is shared. All stations cannot transmit at the same time successfully, as they could in a dedicated-medium situation such as pertains with a switch and point-to-point wiring. The MAC layer determines who transmits when, and if contention is allowed, the MAC controls the contention process and resolves any collisions that occur.

The base station in the 802.16.1 network controls the transmission time for all user stations. It receives requests for transmission rights and grants these requests within the time available, taking into account the priorities of the various types. These services range from carrying TDM information such as voice trunks from a PBX to very bursty but delay-tolerant computer data.

2.1.1 Tunneling other Layer 2 protocols

In a typical LAN situation, the existence of the LAN is known by all user stations; each station on the LAN has appropriate hardware and software drivers for this particular type of LAN. Typically this is Ethernet at 10 or 100 Mbps, but token ring, FDDI, and Appletalk are also widely deployed.

However, beyond the corporate LAN, the situation is more complicated. Different types of communication devices have their own conventions for attachment to wide-area networks, and changing all customers to a new type of external network is not a practical scenario. Instead, 802.16.1 radio systems support the user's current connection technologies, providing a "tunnel" to carry the traffic over the air from the customer premises to the WAN carrier's point of presence.

This means that traffic from an existing PBX can be sent over the 802.16.1 Air Interface to an interexchange carrier, without any modification to the PBX. The output from the PBX is put in appropriate "containers" for transmission over the wireless link and then reconsituted at the other end. Likewise the output from a router can be sent from an Ethernet port to the 802.16.1 unit and delivered to another router at the other end of the wireless link. Alternatively, the router output can be sent over a V.35 physical interface; in this case also, the original format is reconstituted at the other end of the wireless tunnel.

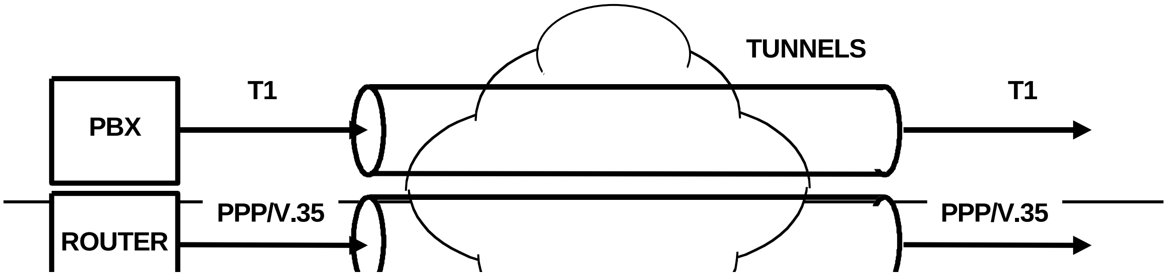

- Figure 2.1-1: Tunneling various traffic types through the network cloud.

Various types of traffic move across the wireless link in tunnels: the native format is converted to one suitable for the wireless link. At the termination of the wireless link, the original format is reconstituted, in terms of both physical and link layer components.

Another way to look at the process is in terms of protocol stacks. In Figure 2.1-2, we use the example of a router which has an Ethernet output port. That protocol is converted to the 802.16.1 wireless protocol, encapsulating the original Ethernet frame, and delivered across the wireless link under control of the wireless MAC. At the other end the process is reversed.

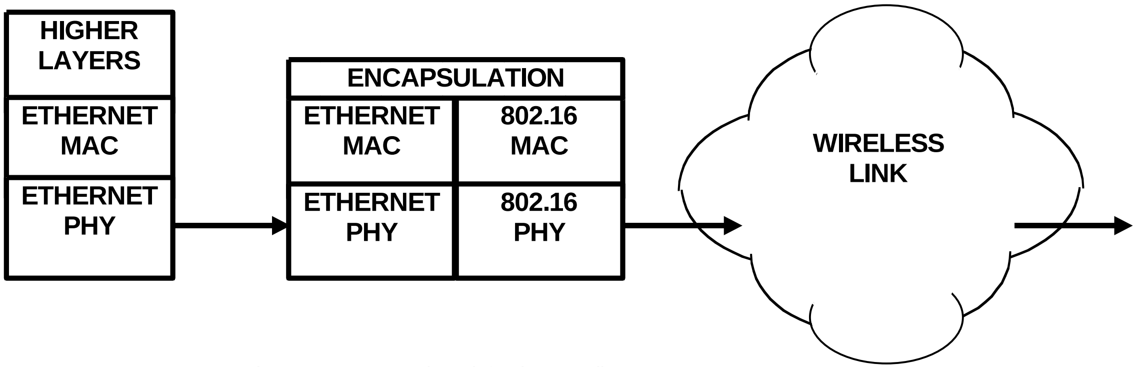

- Figure 2.1-2: Protocol stack for the tunneling process.

The conversion from Ethernet, used in the example here, to the 802.16.1 over-the-air protocol involves encapsulation, retaining the original Ethernet headers. On the far side of the cloud (not shown), the process is reversed to reconstitute the original Ethernet packets. Similar processes are used for other traffic types.

2.2 Architectural View

Figure 2.2-1 shows the architecture of the MAC and the service access points (SAP) provided by the MAC to higher layers.

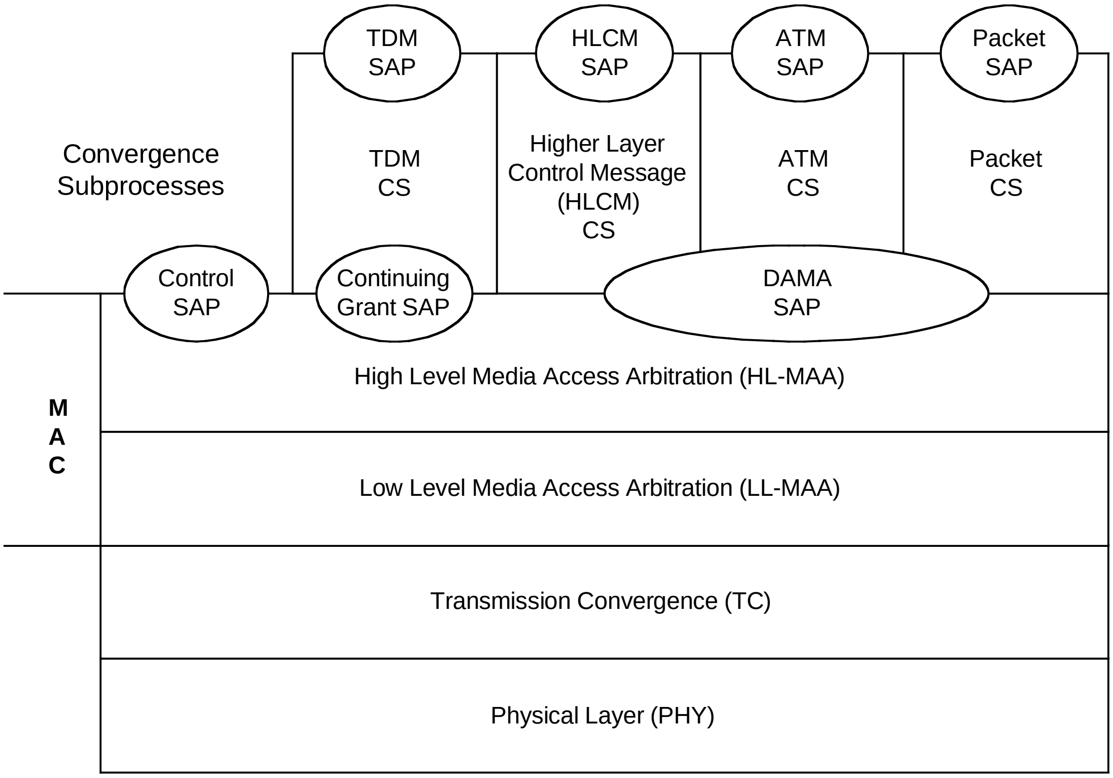
TDM SAP
HLCM SAP
ATM SAP
Packet SAP
Convergence Subprocesses
TDM CS
Higher Layer Control Message (HLCM) CS
ATM CS
Packet CS
Control SAP
Continuing Grant SAP
DAMA SAP
MAC
High Level Media Access Arbitration (HL-MAA)
Low Level Media Access Arbitration (LL-MAA)
Transmission Convergence (TC)
Physical Layer (PHY)

- Figure 2.2-1: MAC Services

The Convergence Subprocesses and their Service Access Points provide the interfaces to the higher layers for service specific connection establishment, maintenance and data transfer.

Due to implementation issues it would be advantageous to partition the MAC into 2 parts: A lower level (LL-MAA) and a higher level (HL-MAA). The HL-MAA has the following purposes:

- Interfacing with higher layers for the establishment and maintenance of data connections.
- Interfacing with higher layers for BS control, CPE registration, etc.
- Load leveling across physical channels.

Through the convergence sublayers, the BS HL-MAA interacts with the higher layers in the BS, accepting or rejecting requests for provisioned connections at varying levels of service based upon both bandwidth availability and connection specific bandwidth limits.

The BS HL-MAA sublayer of the MAC also controls bandwidth allocation and load leveling across physical channels. The BS HL-MAA is cognizant of the loading on all physical channels within this MAC domain. Existing connections may be moved to another physical channel to provide a better balance of the bandwidth usage within a sector.

The LL-MAA performs the bandwidth allocation on an individual physical channel. There is an instance of the BS LL-MAA for each physical channel and an instance of the CPE LL-MAA for each CPE. The LL-MAA is more tightly coupled with the TC and the physical layer than the HL-MAA is. The BS LL-MAA works with the BS HL-MAA, in determining the actual amount of bandwidth available at any given time, based upon bandwidth requests, control message needs, and the specific modulation used to communicate with each CPE. The BS LL-MAA packages downlink data for transmission to the CPEs. The CPE LL-MAA packages uplink data using the same bandwidth allocation algorithm as the BS LL-MAA except limited in scope to the CPE's allocate bandwidth. The LL-MAA may fragment messages across multiple frames.

The TC packages MAC messages into packets compatible with the air interface, distributing MAC messages across TDUs, as necessary.

2.3 BS / CPE Communications and System Management

Communication between the user station and the base station involves four basic types of information exchange:

- User registration and authentication
- Establishment of connection IDs to enable the user to access various services
- Request/grant procedures for transient requirements
- Actual transmission of data

2.3.1 Access and CPE Registration

As part of the access process, the CPE must achieve downlink synchronization with the BS. This process is described in section 3.4.1.2. Once downlink synchronization is achieved, the CPE must go through the ranging process, as described in section 3.4.1.3, to obtain the correct timing advance for uplink transmissions. In addition to affecting the Tx timing advance, the distance from the CPE to the base station affects the power level at which the CPE must transmit to ensure the BS sees all CPEs at similar power. It also affects the choice of modulations to ensure a certain quality of transmission. Power leveling is described in section 3.4.1.4.

2.3.2 Data Connections

For the purposes of mapping to services on CPEs and associating varying levels of QoS, all data communications are in the context of a connection. These connections are provisioned when a CPE is installed in the system, and set up over the air at CPE registration to provide a reference against which to request bandwidth. Additionally, new connections may be established when customer's service needs change.

Once connections are established they must be maintained. The maintenance requirements vary depending upon the type of service connected. For example, unchannelized T1 services require virtually no connection maintenance since they have a constant bandwidth allocated every frame. Channelized T1 services require some maintenance due to the dynamic (but relatively slowly changing) bandwidth requirements if compressed, coupled with the requirement that full bandwidth be available on demand. IP services may require a substantial amount of ongoing maintenance due to their bursty nature and due to the high possibility of fragmentation across frames. As with connection establishment, modifiable connections may require maintenance due to stimulus from either the CPE or the network side of the connection.

Finally, connections may be terminated. This generally occurs only when a customer's service contract changes. The termination of a connection is stimulated by the BS.

The procedures for handling data connections are described in section 3.4.2.

2.3.3 Bandwidth Allocation

One of the differentiating qualities of the 802.16.1 Broadband Wireless Access System is its efficiency in use of the physical channel and the resultant increase in useable bandwidth. In order to efficiently allocate the physical channel, the LL-MAA, in both the BS and the CPE, must be aware of the physical characteristics of the physical channel, and the BS HL-MAA must be aware of the topology of the network. Additionally, while the LL-MAA and TC are limited in its scope to a single physical channel (actually allocated uplink bandwidth in the CPE case), the BS HL-MAA requires visibility of all physical channels within a base station unit that are dedicated to a particular HL-MAA MAC domain.

There are a number of factors which affect the way the BS MAC allocates for a sector. Some of these are:

- Duplexing Scheme
  - Time Division Duplexing
  - Frequency Division Duplexing (half and full duplex)
- Number of physical channels available within the sector.
- Number of CPEs within the sector.
- CPE service requirements – CG vs. DAMA, Quality of Service (QoS), etc.
- CPE modulation capabilities and their affect on available bandwidth.

Due to the need to communicate with different CPEs at different modulation rates, the BS LL-MAA must allocate and map not only the uplink bandwidth, but the downlink as well. While the uplink bandwidth is assigned for each CPE separately the downlink bandwidth allocation can be either jointly (time division multiplexed) or similar to the uplink, separately. In the case of a multiplexed downlink a user is required to decode the whole data stream within a frame and identify its own data by addressing. The need for a mapping technique similar to the uplink is due to efficiency and latency presented in the case of half duplex FDD.

In TDD mode, the percentage of the TDD frame allocated to downlink versus uplink is a system parameter which may change with time. In FDD mode it is fixed. The MAC must meet the requirements of the constant bit rate (CBR) traffic (T1, E1, etc.), and it must also allocate the remaining bandwidth across the lower priority traffic, distributing bandwidth amongst the services using appropriate fairness algorithms such as fair weighted queuing. A certain amount of bandwidth must also be allocated, periodically for station registration and for control messages such as polling and requests for additional bandwidth.

The bandwidth allocation process is described in section 3.4.4.

CPEs are allocated bandwidth for CPE originated connections by making bandwidth requests to the BS. In order to minimize the bandwidth allocated for making bandwidth requests, the opportunities for requesting bandwidth are tightly controlled. There are two general situations where a CPE is allowed to make bandwidth requests, when it is polled or by piggybacking a request on bandwidth already allocated. Polling of CPEs by the BS may be in response to the “poll me” bit in the MAC packet header (See section 3.2.2.7) or periodic as available bandwidth allows. Periodic polls may be to individual CPEs, multicast to groups of CPEs, or broadcast to all CPEs on a physical channel. The polling/bandwidth request mechanism is described in detail in section 3.4.3.

Once a CPE has been allocated bandwidth, the CPE LL-MAA uses the same allocation algorithm as the BS LL-MAA except limited in scoped to the segment of uplink bandwidth it was allocated.

When a sector contains multiple physical channels, an individual channel may become congested due to the requests by CPEs on that channel. To alleviate this situation, traffic may be moved to another physical channel in the sector. All traffic on a given CPE must be moved simultaneously to a different physical channel. The handover process is described in section 3.4.5.

2.3.4 Privacy

Because the 802.16.1 Broadband Wireless Access system operates through the air, there is a need for measures to provide privacy for the network.

It is emphasized that the system is not intended for providing security functions on its air interface. It is expected that if there is an application requiring security (i.e., secured transactions) this will be supplied by the application (layer) itself. This follows the same philosophy as in the case of ordering a telephone line (or a full T1) from the local PTT. The connection itself maintains some privacy, (there is still a chance that eavesdropping will occur) but it is up to the user to either buy special security equipment or some software to increase the security level of the connection even if eavesdropping is possible.

Several types of security need to be provided:

- Privacy: message contents cannot be read by other than the intended destination.
- Authentication: assurance that the user station is the one it purports to be.
- Integrity: messages are delivered complete with no unauthorized additions.
- Availability: unauthorized usage attempts will not result in denial of service to authorized users.

These goals are met by a combination of an authentication process at registration time, secure key exchange, and the use of encryption on all user messages.

Message headers are carried in the clear to enable stations to easily recognize downstream transmissions sent to them. Additionally, MAC control messages are sent in the clear to ensure known data is not transmitted encrypted.

Chapter

3

3 Media Access Control

3.1 Overview

In a network that utilizes a shared medium, there must be a mechanism to provide an efficient way to share the medium. A two-way point-to-multipoint wireless network is a good example of a shared medium: here the medium is the space through which the radio waves propagate.

The downlink, from the base station to the user operates on a point-to-multipoint basis. The 802.16.1 wireless link operates with a central base station and a sectorized antenna which is capable of handling 6 independent sectors simultaneously. Within a given frequency channel and antenna sector, all stations receive the same transmission. The base station is the only transmitter operating in this direction, hence it can transmit without having to coordinate with other stations, except for the overall time-division duplexing that divides time into upstream and downstream transmission periods. It broadcasts to all stations in the sector (and frequency); stations check the address in the received messages and retain only those addressed to them.

However, the user stations share the upstream period on a demand basis. Depending on the class of service utilized, the CPE may be issued continuing rights to transmit, or the right to transmit may be granted by the base station after receipt of a request from the user.

In addition to individually-addressed messages, messages may also be sent to multicast groups (control messages and video distribution are examples of multicast applications) as well as broadcast to all stations.

3.2 Principles of MAC Operation

Within each sector, users must adhere to a transmission protocol which minimizes contention between users and enables the service to be tailored to the delay and bandwidth requirements of each user application.

This is accomplished through polling, with contention procedures as backup should unusual conditions make polling of all users unfeasible within appropriate delay constraints. Contention can also be used to avoid the individual polling of CPEs which have been inactive for a long period of time.

The use of polling simplifies the access operation and guarantees that applications receive service on a deterministic basis if it is required. In general, data applications are delay tolerant, but real-time applications like voice and video require service on a more uniform basis, and sometimes on a very tightly-controlled schedule.

3.2.1 Addressing

Addressing is done at two levels: each user station has a unique 48-bit permanent address, following normal LAN practice. This address is used in the registration process, during which the user station is given a 16-bit basic connection ID and a 16-bit control connection ID. The 16 bit basic connection ID is truncated to 14 bits to provide a short CPE ID. Each service provisioned for a CPE is also assigned a connection ID. Connection IDs are generated by the BS HL-MAA and most are unique across a HL-MAA MAC domain. There is no internal structure to connection IDs.

3.2.1.1 Temporary Registration Connection ID

Connection ID 0x0000 is a multicast connection ID reserved for communicating with CPEs entering the registration process, until they receive a basic connection ID from the base station.

3.2.1.2 PHY Connection ID

Connection ID 0xFFFE is a broadcast connection ID reserved for PHY messages such as the PHY Control message. The MAC Control message is also transmitted on this CID to guarantee synchronization between the two messages.

3.2.1.3 Basic Connection ID

The basic connection ID, assigned to a CPE at registration, is used by the BS MAC and the CPE MAC to exchange MAC control messages. The lower n (where n is base station implementation dependant) immediately following the Temporary Registration connection ID 0x0000 are reserved for basic connection IDs. Of this set, connection IDs 0x0001 and 0x0002 are reserved, and not used. Each physical channel, and therefore each LL-MAA MAC domain, use the same set. Since these basic connections are unique only within an individual LL-MAA MAC domain, they are allocated by the LL-MAA rather than the HL-MAA.

3.2.1.4 Control Connection ID

A control connection ID, also assigned to a CPE at registration, is used by the BS higher layers and the CPE higher layers to exchange control and configuration information at the higher layers. Control connection IDs are unique across the HL-MAA MAC domain.

3.2.1.5 Multicast and Broadcast

The connection ID indicating a broadcast to all stations is indicated by setting all 16 bits to 1 (0xFFFF).

Of the overall total of 64K addresses, 62 are reserved for multicast polling group use. If a connection ID is in the range 0xFFC0-0xFFFD, then the address is for multicast polling group use. Unlike most other connections, the set of connection IDs reserved for multicast polling groups are not unique across an HL-MAA MAC domain. The entire group is allocated to each LL-MAA MAC domain. During handover, a CPE must leave any multicast polling groups it is a member of on the original physical channel. The LL-MAA instance for the new physical channel may assign the CPE to a multicast polling group based on the polling policy currently in effect for the physical channel.

3.2.1.6 Null Connection

Since the TC layer uses 0x55 as a stuff byte, the connections in the range 0x5500 through 0x55FF are reserved. The connection ID 0x5600 is reserved for situations where a fake connection ID is needed. This connection is used when padding a Reed-Solomon block at the end of a modulation on the downlink or at the end of a CPE's transmission on the uplink. Any data received on this connection is always discarded.

3.2.1.7 Additional Connection IDs

For additional services, the higher layers of the BS set up connections via requests to the BS HL-MAA. These connections are set up based upon the provisioning information distributed to the base station. The registration of a CPE, or the modification of the services contracted at a CPE, stimulates the higher layers of the BS to initiate the setup of the connections.

The connection ID can be considered a connection identifier even for nominally connectionless traffic like IP, since it serves as a pointer to destination and context information. The use of a 16-bit connection ID permits a total of 64K connections within the sector.

Requests for transmission are based on these connection IDs, since the allowable bandwidth may differ for different connections, even within the same service type. For example, a CPE unit serving multiple tenants in an office building would make requests on behalf of all of them, though the contractual service limits and other connection parameters may be different for each of them.

Many higher-layer sessions may operate over the same wireless connection ID. For example, many users within a company may be communicating with TCP/IP to different destinations, but since they all operate within the same overall service parameters, all of their traffic is pooled for request/grant purposes. Since the original LAN source and destination addresses are encapsulated in the payload portion of the transmission, there is no problem in identifying different user sessions. When grants are made, they are made by CPE only; it is up to the CPE to organize the messages and transmit them within the designated time period.

The type of service is implicit in the connection ID; it is accessed by a lookup indexed by the connection ID.

3.2.2 Framing and Formatting

The uplink and downlink can be multiplexed in a TDD fashion as described in section 3.2.2.1 or in an FDD fashion as described in section 3.2.2.2. Each has a standard unit of a 1 msec frame. Within this frame are a downlink subframe and an uplink subframe. In the TDD case, the downlink subframe comes first, followed by the uplink subframe. In the FDD case, the downlink and uplink subframes occur simultaneously on their respective frequencies. The downlink subframe is prefixed with information necessary for frame synchronization.

Due to the coupling necessary between the MAC and the physical layer to accommodate such features as adaptive modulation, an example physical layer is used for reference. Other physical layers may be accommodated by adjusting the parameters important to bandwidth allocation.

The available bandwidth is allocated in symbol granularity. The PHY frame is a 1mSec frame sub-divided into 800 PHY Slots (PS). Forward Error Correction uses 600 bit units, called PHY Information Elements (PI). The modulation within the frame may vary, and determines the number of symbols required to transmit a PI. Transmission of one PI requires:

- 4 PSs for QAM-64
- 6 PSs for QAM-16
- 12 PSs for QAM-4 (QPSK, downstream only)
- 15 PSs for CQPSK (upstream only)

Each PI provides 55 bytes to the TC for transport of MAC messages, control information, and data. This 55 byte block is called a TC Data Unit (TDU).

3.2.2.1 Time Division Duplex and Support of Variable Traffic Asymmetry Conditions

In this mode of operation the downlink and uplink are on the same carrier frequency. The uplink and downlink share the same frequency in a TDM fashion. A TDD frame has a 1 msec duration, and contains 800 PS as shown in Figure 3.2-1. The TDD framing is adaptive in that the number of PS allocated to downlink versus uplink can vary. The split between uplink and downlink is a system parameter and is controlled at higher layers within the system.

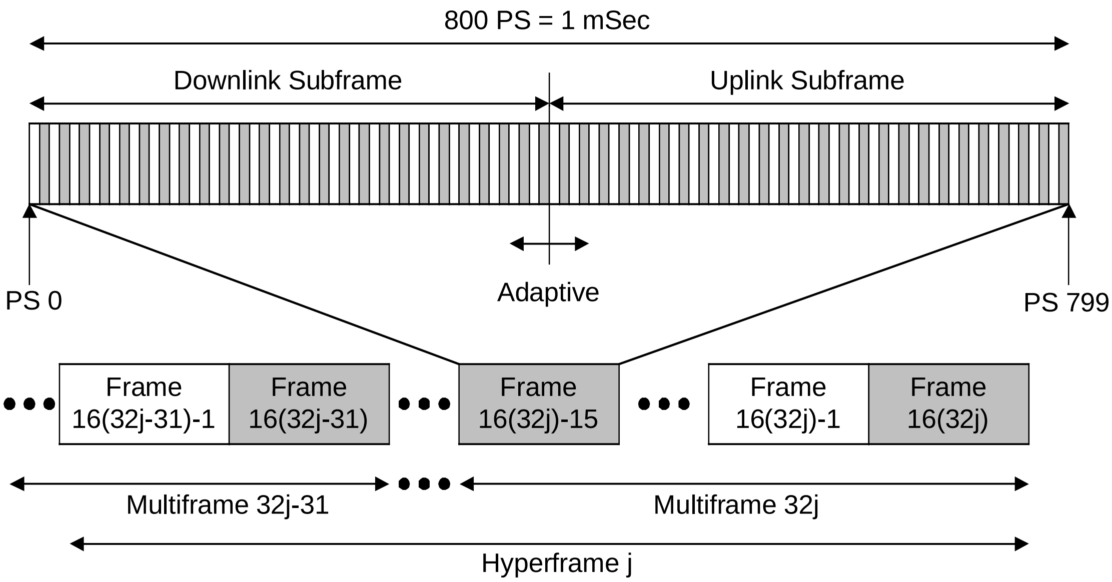

- Figure 3.2-1: TDD and Multiframe Structure

To aid periodic functions, multiple frames are grouped into multiframes and multiple multiframes are grouped into hyperframes. There are 16 frames per multiframe and 32 multiframes per hyperframe. Hyperframes rollover to zero after 32767. Since the TDD frames are 1 millisecond in duration, 1 multiframe is 16 milliseconds in duration and 1 hyperframe is 512 milliseconds in duration.

Frame numbering is synchronized with Universal Time. The frame, multiframe, and hypeframe are defined to be all 0 at Julian Date 2451179.0 (noon Jan. 0, 1990). The frame, multiframe, and hyperframe return to all 0 with a period of 16777216 milliseconds which equals 4 hours 39 minutes 37 seconds 216 milliseconds.

The base station MAC must be provided with time, in some form, synchronized to some synchronization signal with at least 1millisecond accuracy.

3.2.2.2 FDD and H-FDD Operation

In this mode of operation the downstream and upstream are using 2 different carrier frequencies. Both carriers are equal in channel bandwidth *and* instantaneous baud rate. The frequency separation between carriers is set either according to the target spectrum regulations or to some value sufficient for complying with radio channel transmit/receive isolation and desensitization requirements. In the time domain both upstream and downstream are frame *synchronized*.

A subscriber capable of full duplex FDD operation, meaning it is capable of transmitting and receiving at the same instant, imposes no restriction on the base station controller regarding its upstream bandwidth allocation management. On the other hand, a subscriber that is limited to half duplex FDD operation imposes a restriction on such a controller not to allocate upstream bandwidth for the subscriber, which may force it to instantaneously transmit and receive. It is mandatory that both types of subscribers could co-exist in a FDD deployment, meaning that radio channels could address both type of subscribers instantaneously.

Figure 3.2-2: FDD User Mapping) describes the basics of the FDD and H-FDD based operation. Frames are either even numbered or odd numbered. A subscriber limited to H-FDD operation is designated to operate either on even frames or odd frames. Those that are receiving downstream on even frames are using odd frames for upstream and vice versa. A user that is capable of full duplex FDD ignores the even/odd structure and may utilize the system on both even and odd frames.

In order to increase statistical gain a user may change its even-odd frame relationship according to traffic requirements. When a user has no upstream bandwidth it is required to receive all frames. When bandwidth is being allocated for it then the user limits itself by the frame assigning its bandwidth. If the frame assigning bandwidth on the downstream is even numbered than its upstream frames would be odd numbered and vice versa.

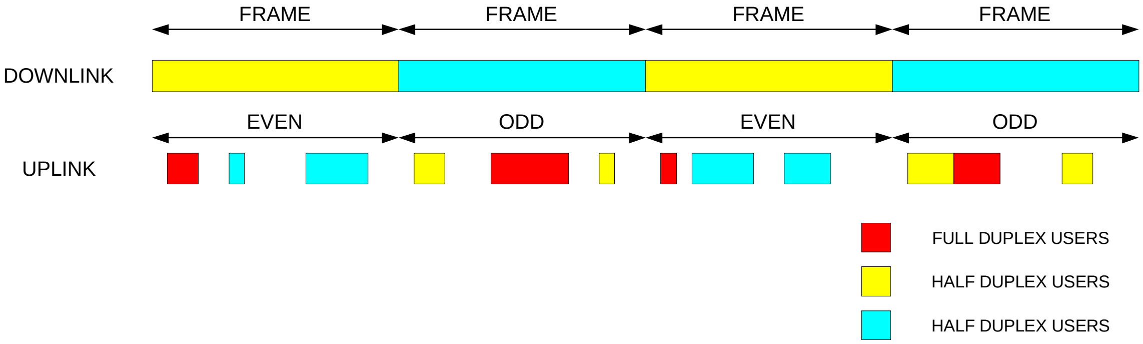

• Figure 3.2-2: FDD User Mapping

3.2.2.3 Downlink Subframe

The structure of the downlink subframe used by the BS to transmit to the CPEs is shown in Figure 3.2-3. It starts with a Frame Control Header, which is always transmitted in QAM-4. This frame header contains a preamble used by the PHY for synchronization and equalization. It also contains control sections for both the PHY and the MAC. Within the downlink subframe, transmissions are grouped by modulation type. They may optionally be grouped by CPE. Preambles are not RS coded, but all other downlink traffic is FEC coded. In the TDD case there is a Tx/Rx Transmission Gap (TRTG) separating the downlink subframe from the uplink subframe.

Note that any one or more of the 3 differently modulated data blocks may be absent.

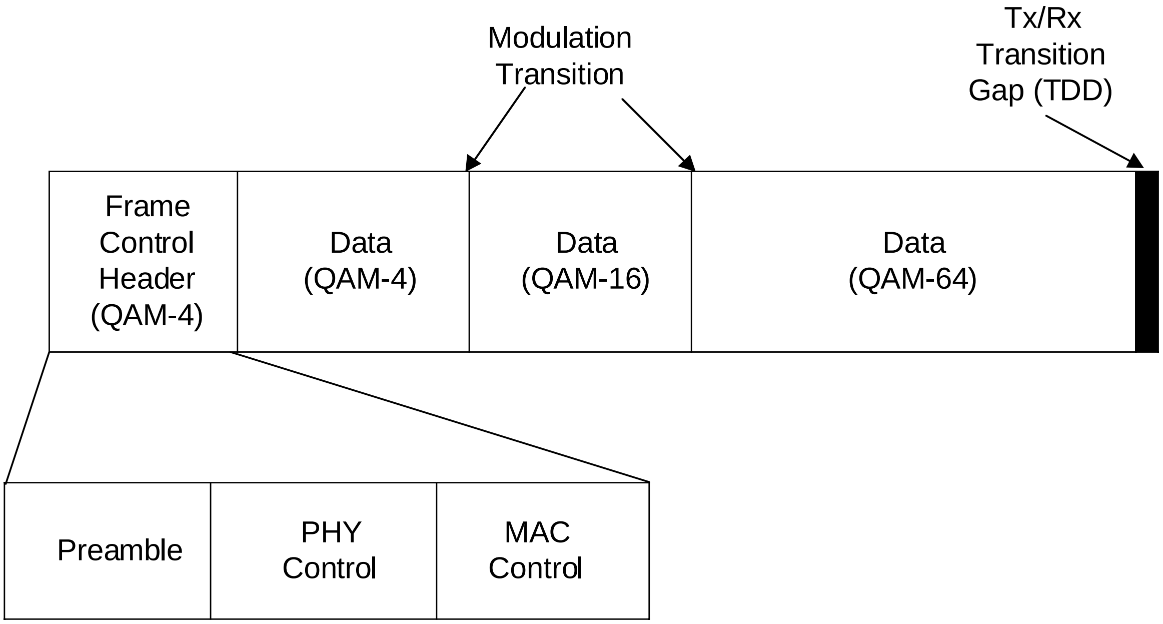

• Figure 3.2-3: Downlink Subframe Structure

3.2.2.3.1 PHY Control

The PHY Control portion of the downlink subframe is used for physical information destined for all CPEs. The PHY Control information is FEC encoded, but is not encrypted. The information transmitted in this section is always transmitted in QAM-4 and includes:

- Broadcast physical layer information
- Maximum Tx timing advance
- DL modulation transition points
- End of DL, this frame
- Frame/multiframe/hyperframe numbering

See section 3.6.1 for the format of the PHY Control message.

3.2.2.3.2 MAC Control

The MAC Control portion of the downlink subframe is used for MAC messages destined for multiple CPEs. For information directed at an individual CPE, MAC messages are transmitted in the established control connection at the operating modulation of the CPE to minimize bandwidth usage. The MAC Control messages are FEC encoded, but are not encrypted. The information transmitted in this section is always transmitted in QAM-4 and includes:

- MAC Version Identifier

- Uplink Map (CPE/PS pairs)
- Whether any bandwidth request contention periods (see section 3.2.2.4.2) are included the frame (in UL Map)
- Starting point and length of bandwidth request contention period, if any (in UL Map)
- Whether registration is allowed on this physical channel
- Whether a registration contention period is included the frame (in UL Map)
- Starting point and length of registration contention period, if any (in UL Map)
- Optional DL user map to support TDMA on the upstream

See section 3.6.2 for the format of the MAC Control message.

3.2.2.3.3 Downlink Data

The downlink data sections are used for transmitting data and control messages to the CPEs. This data is always FEC coded and is transmitted at the current operating modulation of the individual CPE. Message headers are sent unencrypted. Payloads of user data connections are encrypted. Payloads of MAC control connections are not encrypted. Data is transmitted in modulation order QAM-4, followed by QAM-16, followed by QAM-64. The PHY Control portion of the Frame Control Header contains fields stating the PS at which modulation will change.

If the downlink data does not fill the entire downlink subframe, the downlink subframe is padded with fill data (0x55). If one or more TC data units (TDUs) remain to be filled, the MAC performs the fill on a specific connection ID. If less than one TDU remains to be filled, the TC performs the fill. In the case of H-FDD filling is replaced by transmitter shut-down in order to allow parallel uplink allocations. If one or more complete TDUs remain to be filled, they may be filled using connection 0x5600.

3.2.2.4 Uplink Subframe

The structure of the uplink subframe used by the CPEs to transmit to the BS is shown in Figure 3.2-4. There are three main classes of MAC/TC messages transmitted by the CPEs during the uplink frame:

- Those that are transmitted in contention slots reserved for station registration.
- Those that are transmitted in contention slots reserved for response to multicast and broadcast polls for bandwidth needs.
- Those that are transmitted in bandwidth specifically allocated to individual CPEs.

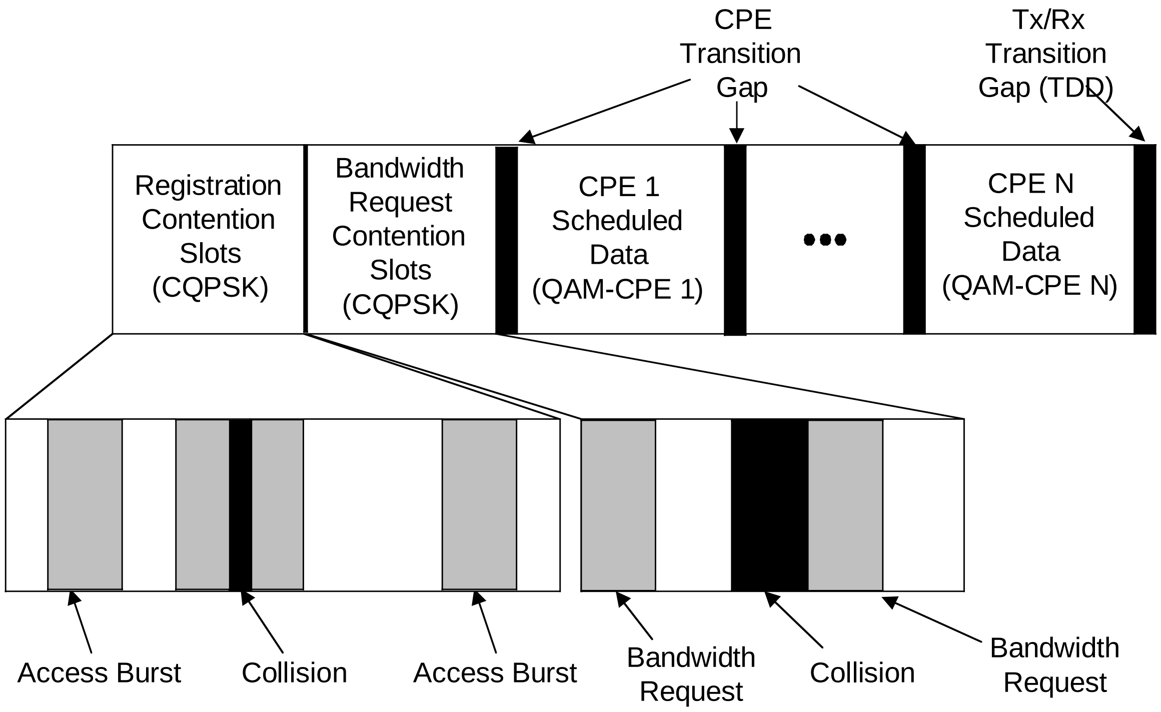
CPE Transition Gap
Tx/Rx Transition Gap (TDD)
Registration Contention Slots (CQPSK)
Bandwidth Request Contention Slots (CQPSK)
CPE 1 Scheduled Data (QAM-CPE 1)
•••
CPE N Scheduled Data (QAM-CPE N)
Access Burst
Collision
Access Burst
Bandwidth Request
Collision
Bandwidth Request

- Figure 3.2-4: Uplink Subframe Structure

The bandwidth allocated for contention slots is grouped together and is transmitted using CQPSK modulation. The remaining, scheduled, bandwidth is grouped by CPE. During its scheduled bandwidth, a CPE transmits with a fixed modulation, determined by the effects of environmental factors on transmission to and from that CPE. CPE Transition Gaps (CTG) separate the transmissions of the various CPEs during the uplink subframe. The CTGs are 2 PS in length. The transmitting CPE transmits a 1 PS preamble during the second PS of the CTG allowing the BS to synchronize to the new CPE.

3.2.2.4.1 Registration Contention Slots

A portion of the uplink bandwidth may periodically be allocated for registration contention slots. Registration contention slots are used to allow CPEs to register with the BS and to perform ranging. Additionally, they may be used for physical layer maintenance for CPEs that have been idle for some period of time. CPEs wishing to register and range must have acquired downlink synchronization with the BS, but do not know their Tx timing advance or an appropriate power level. Additionally, they do not yet have a basic connection ID assigned for direct communication with the BS. The registration contention slots allow access under these conditions, allowing CPEs to finalize their uplink physical synchronization with the BS and to establish a logical connection for control communication. The registration process is described in section 3.4.1.

Multiple CPEs may transmit in the registration contention period simultaneously, potentially causing collisions. When a collision occurs, the BS does not respond. If the BS successfully receives a registration message from a CPE, it responds with a registration results message in the QAM4 portion of the downlink subframe. Messages sent in the registration contention period are not encrypted. They are proceeded by a preamble.

The round trip delay for a 3 km cell causes a CPE with no Tx timing advance to transmit up to 16 PS late, not including delays through the modem. More PS may be allocated to reduce the likelihood of collision or to allow larger cells. Figure 3.2-5 shows the relationship between the registration contention slot window and the various parameters governing the timing of messages within the window.

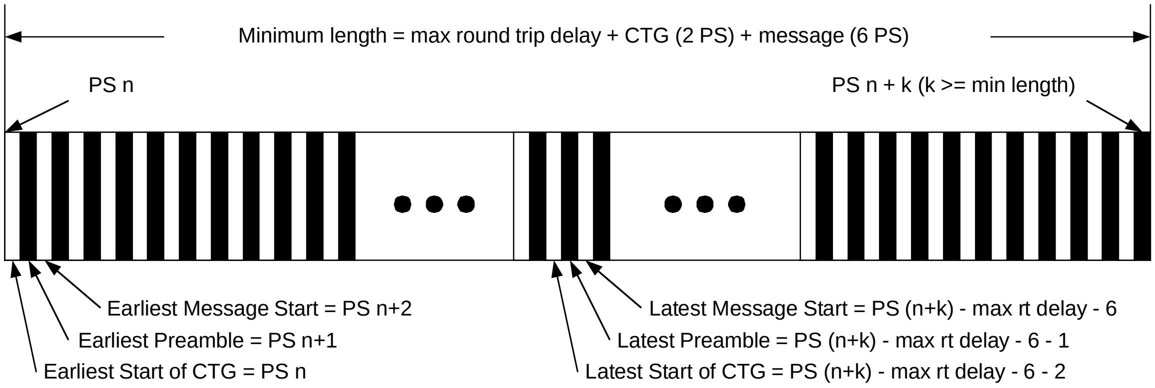
Minimum length = max round trip delay + CTG (2 PS) + message (6 PS)
PS n
PS n + k (k >= min length)
Earliest Message Start = PS n+2
Earliest Preamble = PS n+1
Earliest Start of CTG = PS n
Latest Message Start = PS (n+k) - max rt delay - 6
Latest Preamble = PS (n+k) - max rt delay - 6 - 1
Latest Start of CTG = PS (n+k) - max rt delay - 6 - 2

- Figure 3.2-5: Registration Contention Slot Usage

3.2.2.4.2 Bandwidth Request Contention Slots

A portion of the uplink bandwidth may periodically be allocated for bandwidth or connection requests. Since a CPE must be registered and must have achieved uplink synchronization with the BS before it is allowed to request bandwidth, there is no Tx time uncertainty to be allowed for in the length of the bandwidth request contention period. As with registration requests, if a collision occurs, the BS does not respond. If the BS successfully receives a bandwidth request message, it responds by allocating the CPE (additional) bandwidth in the Uplink Map. Polling and piggybacking help to minimize the need to use bandwidth request contention slots. The polling and bandwidth request process are described in detail in section 3.4.3.

3.2.2.4.3 Scheduled Uplink Traffic

Scheduled uplink traffic is bandwidth allocated to specific CPEs for the transmission of control messages and user data. The CPEs are ordered by modulation. The bandwidth is requested by the CPE and granted by the BS. All bandwidth within a given frame, allocated to an individual CPE, is grouped into a contiguous block. The 2 PS for the CTG are included in the allocation to the CPE in the Uplink Map. The CPE transmits a preamble in the second PS of the CTG at the start of its allocated block. The preamble is neither RS coded nor encrypted. The TDUs transmitted are always RS coded.

3.2.2.5 Time Relevance of PHY and MAC Control Information

The information in the PHY Control portion of the Frame Control Header pertains to the current frame (i.e., the frame in which it was received). The information in the Uplink Subframe Map in the MAC Control portion of the Frame Control Header pertains to the following frame (i.e., one frame after it is received). This timing holds for both the TDD and FDD variants of the system. The TDD variant is shown in Figure 3.2-6. The FDD variant is shown in Figure 3.2-7.

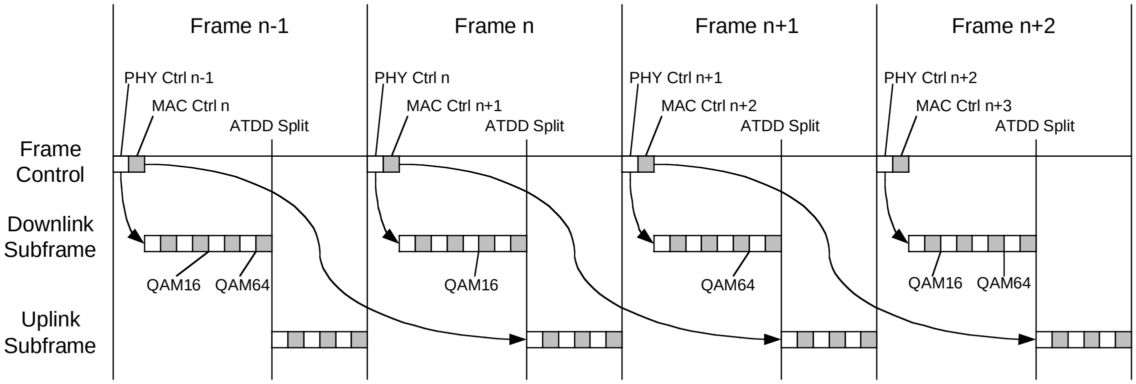

• Figure 3.2-6: Time Relevance of PHY and MAC Control Information (TDD)

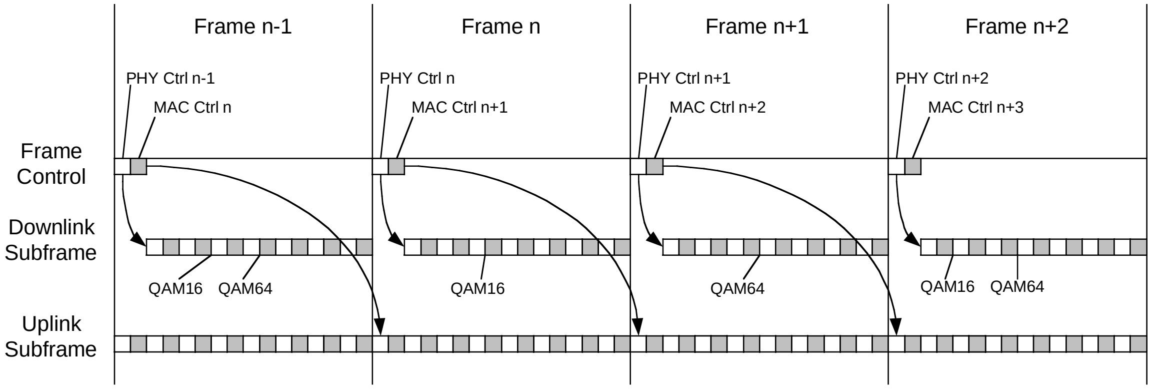

• Figure 3.2-7: Time Relevance of PHY and MAC Control Information (FDD)

3.2.2.6 TC/PHY Packet Unit (TDU) Format

Figure 3.2-8 shows the format of the TDU.

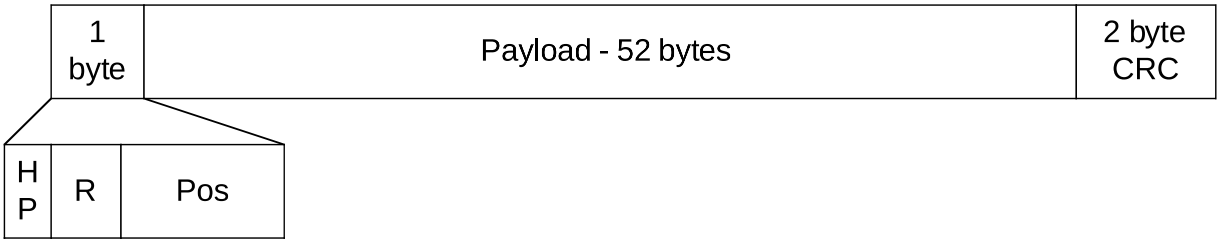

- HP - Header Start Present Bit
- R - 1 reserved bit
- Pos - 6 bits for position (0-based) of MAC header within TC/PHY payload

• Figure 3.2-8: TDU Structure

The TDU has a payload of 52 bytes. If a MAC packet (see section 3.2.2.7) is longer than 52 bytes, that packet must be fragmented over more than 1 TDU. The resultant packets are transmitted, back to back, within the same frame. If a MAC packet ends in the middle of a TDU, the next MAC packet, if one exists, is started immediately. Exceptions to this are:

- At change of modulation, on the downlink, the first packet at the new modulation starts in a new TDU following the modulation transition.

- At change of CPE, on the uplink, the first packet from the next CPE starts in a new TDU following a CTG.

In order to increase efficiency a smaller TDU with less payload bytes is allowed. For example if an allocation for a user on the uplink requires from the MAC 80 bytes then 2 packets would be used. One full TDU is sent with one shortened which contains 38 bytes. The scheduler has taken into account that the PHY resources would handle the shortened packet correctly by not transmitting any zero padding necessary to complete the FEC operation.

The TDU has an 8-bit header and a 2 byte CRC. The header starts with the Header Present bit. This bit is 1 if a MAC header starts somewhere in the TDU. The next bit is reserved. The last 6 bits indicate the byte position with in the payload at which the MAC header, if present, starts.

Encryption is a MAC issue. So, when encryption is performed, the TC/PHY header and CRC are always left in the clear. The payload follows the encryption rules for the MAC layer.

3.2.2.7 MAC Packet and Header Format

Data for all applications is sent in packets prefaced with a header containing the connection ID and a variety of status bits and other fields. User stations recognize data sent to them by the base station by the connection ID. They then process the packets appropriately based on information referenced by the connection ID.

There are two distinct forms of MAC header: the standard MAC header and the abbreviated MAC header. The two header types are mutually exclusive. A particular network of base stations and CPEs will use one or the other, but never both. The standard MAC header is the normal case, for support of variable length packets over the air interface. The abbreviated MAC header allows for fixed size packets over the air interface. It is only used in systems where the backhaul to the base station uses fixed size packets (such as with an ATM backhaul) and the conversion of the backhaul protocol is performed at the CPE rather than the base station. The two headers are identical except for the absence of a length field in the abbreviated MAC header.

The MAC header varies slightly for uplink and downlink.

The format of MAC downlink packets using the standard MAC header is shown in Figure 3.2-9. The format of MAC downlink packets using the abbreviated MAC header is shown in Figure 3.2-10.

<table>
<tr><td colspan="1">1</td><td colspan="2">PC</td><td colspan="1">E</td><td colspan="4">reserved</td></tr>
<tr><td colspan="4">reserved</td><td colspan="4">CID 15:12</td></tr>
<tr><td colspan="8">CID 11:4</td></tr>
<tr><td colspan="4">CID 3:0</td><td colspan="3">Frag</td><td colspan="1">PLP</td></tr>
<tr><td colspan="5">reserved</td><td colspan="3">Len 10:8</td></tr>
<tr><td colspan="8">Len 7:0</td></tr>
<tr><td colspan="8">Payload - Len Bytes</td></tr>
</table>

- Figure 3.2-9: Variable Length MAC Downlink Packet and Standard Header Format

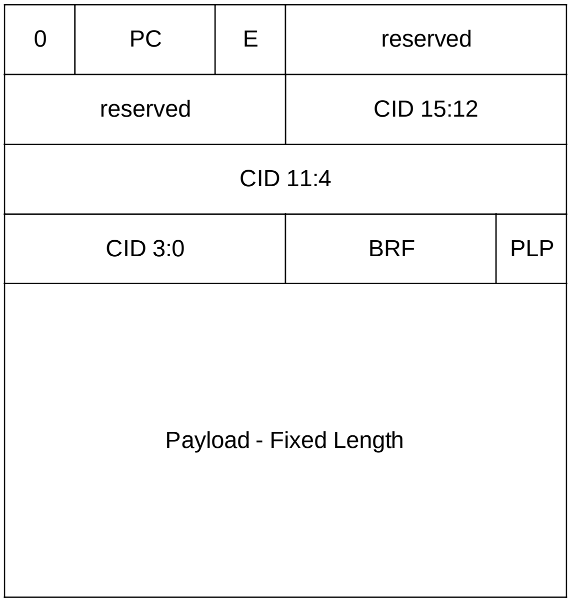

- Figure 3.2-10: Fixed Length MAC Downlink Packet and Abbreviated Header Format

The MAC header starts with the standard header flag. This flag is set to 1 in systems that allow variable length packets and is set to 0 in systems that use fixed length packets and the abbreviated MAC header. The two power control bits are used for fast, small adjustments in a CPEs power. Power is adjusted in relative rather than absolute amounts. The use of these bits is:

00 don’t change

01 increase power a little

11 decrease power a little

10 reserved.

Next, the encryption bit is set to 1 if the payload is encrypted and 0 if it is not. The MAC header is always sent unencrypted. The connection ID is a 16 bit destination identifier set up between the BS and CPE at the time of connection establishment. It is preceded by 8 bits reserved for future expansion of the connection ID field.

The next 3 bits control fragmentation. When the system is configured to use variable size packets (standard configuration), the MAC must perform fragmentation to efficiently use the air link bandwidth. In this case, the Fragmentation field has the following format:

010 Begin fragment of a fragmented message.

000 Continuation fragment of a fragmented message.

100 End fragment of a fragmented message.

110 Unfragmented message.

When the system is configured for fixed size packets, the MAC does not perform fragmentation. In this case these 3 bits are defined as reserved for backhaul fragmentation and are used to pass through backhaul specific fragmentation information.

The packet loss priority (PLP) bit is set to 1 for low priority packets that may be discarded first in congestion situations.

If the system uses the standard MAC header, the 11-bit length field, preceded by 5 reserved bits, indicates the number of bytes in the MAC packet payload

The MAC payload is a portion of a service type specific data element.

The format of the MAC uplink packets using the standard MAC header is shown in Figure 3.2-11. The format of the MAC uplink packets using the abbreviated MAC header is shown in Figure 3.2-12.

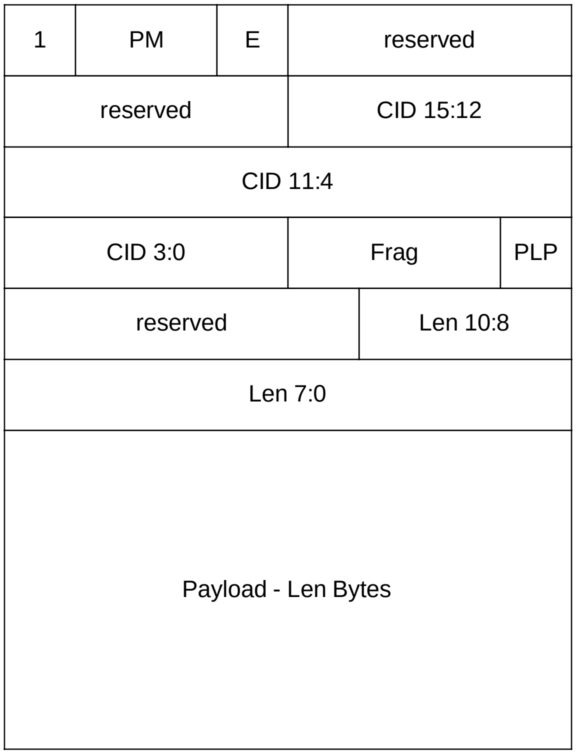

<table>
<tr><td>1</td><td>PM</td><td>E</td><td colspan="3">reserved</td></tr>
<tr><td colspan="3">reserved</td><td colspan="3">CID 15:12</td></tr>
<tr><td colspan="6">CID 11:4</td></tr>
<tr><td colspan="3">CID 3:0</td><td colspan="2">Frag</td><td>PLP</td></tr>
<tr><td colspan="4">reserved</td><td colspan="2">Len 10:8</td></tr>
<tr><td colspan="6">Len 7:0</td></tr>
<tr><td colspan="6">Payload - Len Bytes</td></tr>
</table>

• Figure 3.2-11: Variable Length MAC Uplink Packet and Standard Header Format

<table>
<tr><td>0</td><td colspan="2">PM</td><td>E</td><td colspan="4">reserved</td></tr>
<tr><td colspan="4">reserved</td><td colspan="4">CID 15:12</td></tr>
<tr><td colspan="8">CID 11:4</td></tr>
<tr><td colspan="4">CID 3:0</td><td colspan="3">BRF</td><td>PLP</td></tr>
<tr><td colspan="8">Payload - Fixed Length</td></tr>
</table>

- Figure 3.2-12: Fixed Length MAC Uplink Packet and Abbreviated Header Format

The uplink MAC header differs slightly from the downlink MAC header in that the 2-bit power control field is replaced by a 2-bit Poll Me field. The Poll Me bits are used to indicate that the CPE requests to be polled for bandwidth or connection requests. The PM bit 0 indicates that the request will be for a connection with QoS between **TBD** and 255. The PM bit 1 indicates that the request will be for a connection with QoS between 1 and **TBD**.

3.2.2.7.1 Continuing Grant Payload

The Continuing Grant (CG) payload from the MAC's point of view is simply some number of bytes, which are not allowed to be fragmented across multiple frames. To ensure quick response to a request for more bandwidth of a CG connection, the uplink bandwidth allocated to a CG connection that is not at its maximum rate is large enough to accommodate the connection's current rate plus a bandwidth request.

The CG payload structure is shown in Figure 3.2-13.

CG Payload - Unstructured from MAC pov

- Figure 3.2-13: CG Payload

3.2.2.7.2 DAMA Payloads (IP and Others)

The DAMA payload from the MAC's point of view is simply some number of bytes which are passed to the higher layers as a single unit. Fragmentation may be necessary if the packet will not fit in the available transmission slot, or if a portion of the scheduled transmission time is used to piggyback an additional request. There can be at most one packet per DAMA connection in a fragmentation and re-assembly state at any given time. However, any number of DAMA connections may have a packet in a fragmentation re-assembly state at the same time.

The DAMA payload format is shown in Figure 3.2-14.

Dama Payload - Unstructured from MAC pov

- Figure 3.2-14: DAMA Payload

3.2.2.7.3 Control Messages

Control messages are sent on the CPE’s basic connection. Multiple control messages to or from the same CPE may be packed into a single MAC packet. Fragmentation of MAC packets containing MAC Control messages is not allowed. MAC Control messages are specified in section 3.6.

Control messages have the form shown in Figure 3.2-15.

Message Type ID	Message Body (Message Type Specific)

- Figure 3.2-15: Control Message Structure

3.2.3 Mapping of MAC Entities to PHY Elements

The BS LL-MAA performs all allocation and mapping of the available bandwidth of a physical channel based on the priority and quality of services requirements of requests received from higher layers. Additionally, the availability of bandwidth is based on the modulation required to achieve acceptable BER between the BS and the individual CPEs. The BS MAC uses information from the PHY regarding signal quality to determine the modulation required for a particular CPE and, therefor, the bandwidth that is available. Once the BS LL-MAA has allocated uplink bandwidth to the CPEs, each CPE’s LL-MAA, in turn, allocates that bandwidth to the uplink requests it has outstanding.

The minimum physical unit the LL-MAA allocates is symbol based. The typical unit the LL-MAA allocates is the 52 byte payload of the 55 byte TC Data Unit (TDU). The FEC is performed on the TDU to create 600 bit PIs. Bandwidth needs that do not require FEC, such as the various transition gaps, are allocated in time units of 1 PS. Bandwidth needs that require FEC coding are allocated in TDUs, with each modulation, on the downlink, and each CPE’s transmission, on the uplink, padded to an integer multiple of TDUs to create an integer multiple of PIs. This padding is described in more detail in section 3.2.2.3.3. The number of PSs required to transmit a PI varies with modulation as mentioned previously. The MAC allows for a last shortened PI for efficiency.

Figure 3.2-16 shows how a stream of variable length MAC messages map to the 26 byte payloads of the TDUs, which map to PIs, and finally to PSs and symbols. In the figure it is assumed that there was no need to shorten a PI.

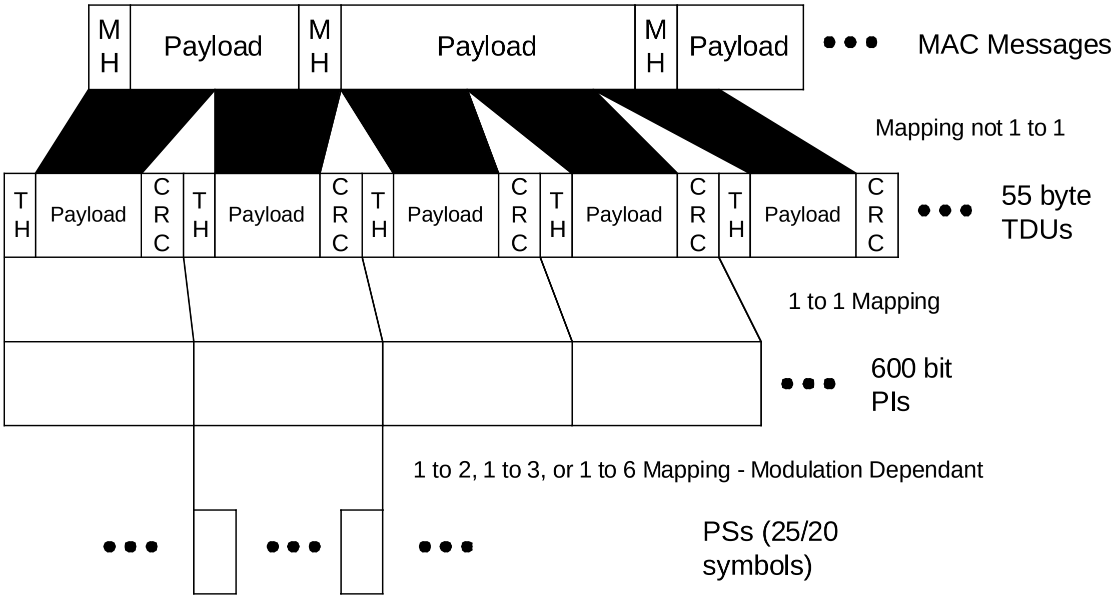

• Figure 3.2-16: Mapping of Variable Length Packets to PHY

3.2.3.1 Downlink Mapping of MAC to PHY

As was shown in Figure 3.2-3, the downlink subframe starts with a Frame Control Header containing a preamble of a fixed length, a PHY control section and a MAC control section. This Frame Control Header allows CPEs to synchronize with the downlink and determine the mapping of the uplink and the downlink.

Figure 3.2-17 shows the mapping of the body of the downlink subframe to the downlink needs of users. Within the subframe, TDUs are grouped by modulation. Within the modulation blocks, packets can be grouped by CPE, but do not need to be. All messages (other than in the frame header) for an individual CPE are transmitted with the same modulation. Each series of MAC packets at a particular modulation must be padded to be an integer multiple of a TDU to provide an integer multiple of a PI after coding. The padding uses the fill byte 0x55.

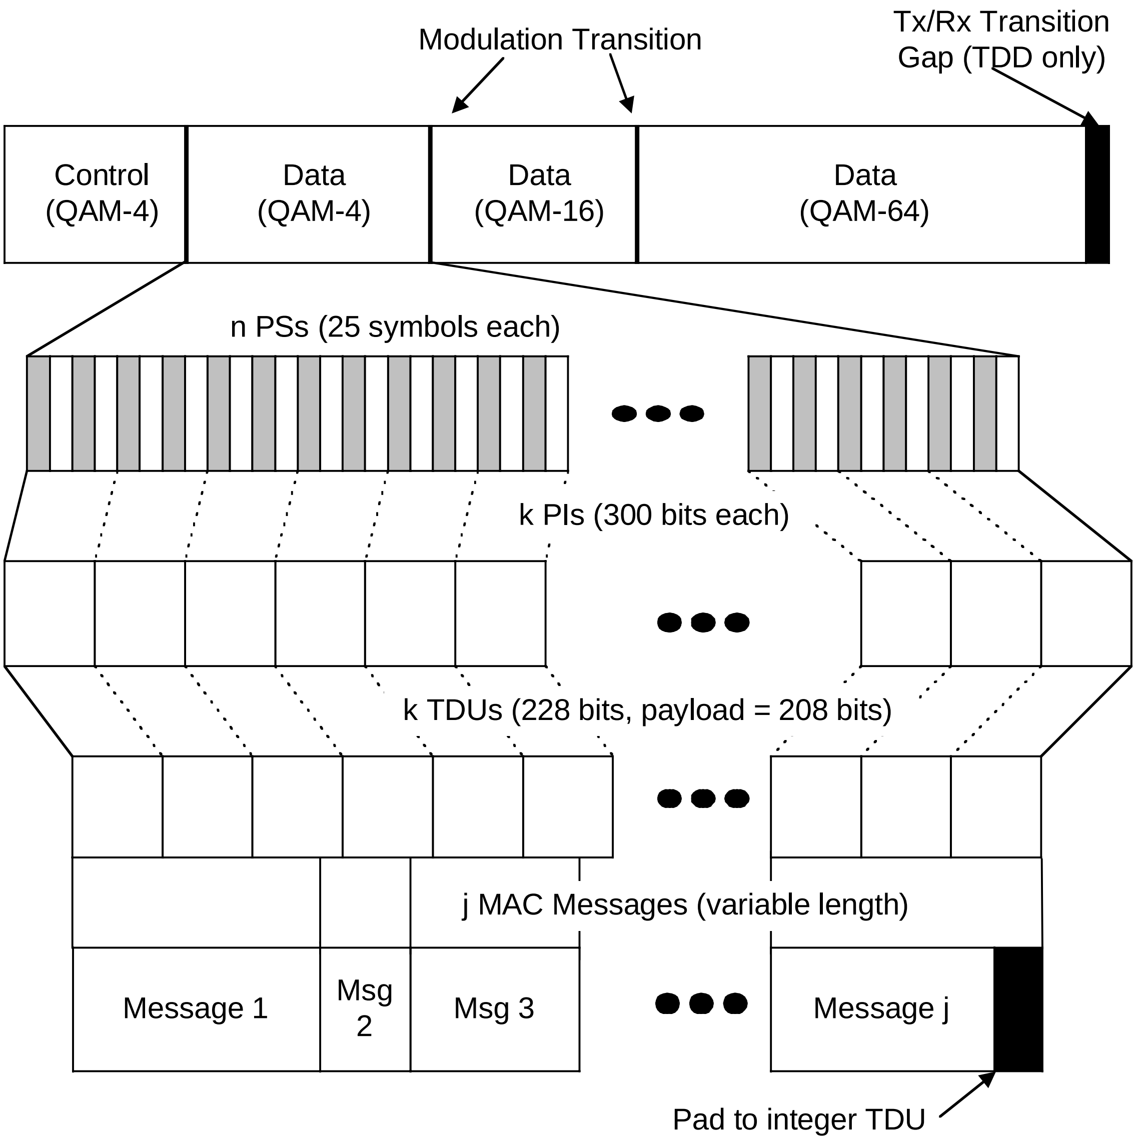

• Figure 3.2-17: Downlink Mapping of MAC Messages to PHY Elements

The MAC supports an additional mode advantegous for half duplex FDD. In this mode each CPE is mapped individually to the downlink resource. It is required that the such a transmission on the downlink would be prefixed with a short preamble on the PHY. An allocation map for downlink appears as part of the control information at the beginning of the frame.

3.2.3.2 Uplink Mapping of MAC to PHY

As shown in Figure 3.2-18, the uplink subframe starts with optional registration contention slots. Some slots of this type are allocated periodically to the PHY for use during station registration. Registration messages are proceeded by a 1 PS preamble. Registration and ranging messages must be sent alone. No other MAC control messages may be packed into the same MAC packet. Due to the short length of these MAC commands, the TDU for them is shortened. Next are slots that are allocated for responses to multicast and broadcast polls for bandwidth requirements. The bandwidth request messages, when sent in the bandwidth request contention period, must be proceeded by a 1 PS preamble. CPEs may pack additional bandwidth requests for other connections into the same MAC packet as part of the padding to a full TDU however shortening of a TDU is allowed as these MAC commands are short in length.

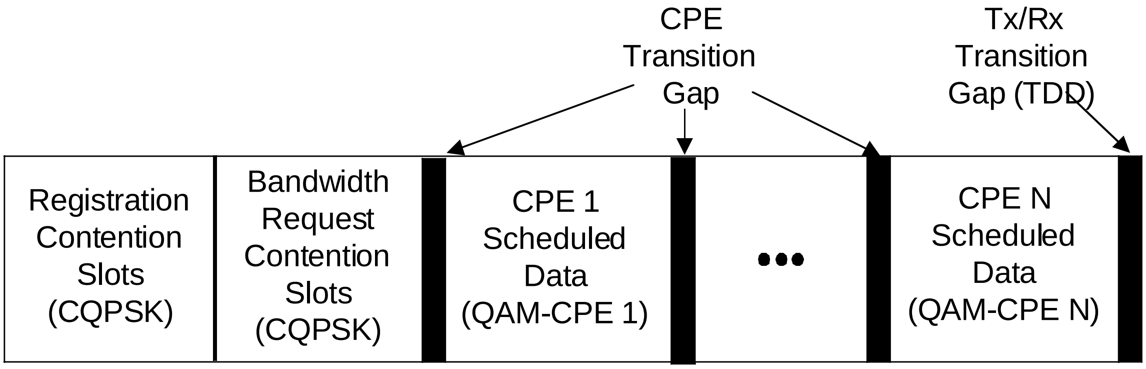
CPE Transition Gap
Tx/Rx Transition Gap (TDD)
Registration Contention Slots (CQPSK)
Bandwidth Request Contention Slots (CQPSK)
CPE 1 Scheduled Data (QAM-CPE 1)
•••
CPE N Scheduled Data (QAM-CPE N)

- Figure 3.2-18: Uplink Contention Access Slots

Figure 3.2-19 shows the mapping of the scheduled portion of the uplink subframe to the uplink needs of users. The CTG contains a 1 PS preamble to ensure synchronization with the new CPE. Within the subframe, TDUs are grouped by CPE. All messages (other than bandwidth requests transmitted in bandwidth request contention slots) from an individual CPE are transmitted with the same modulation. Each CPE's transmission must be padded to be an integer multiple of a TDU to provide an integer multiple of a PI after coding. The padding uses the fill byte 0x55.

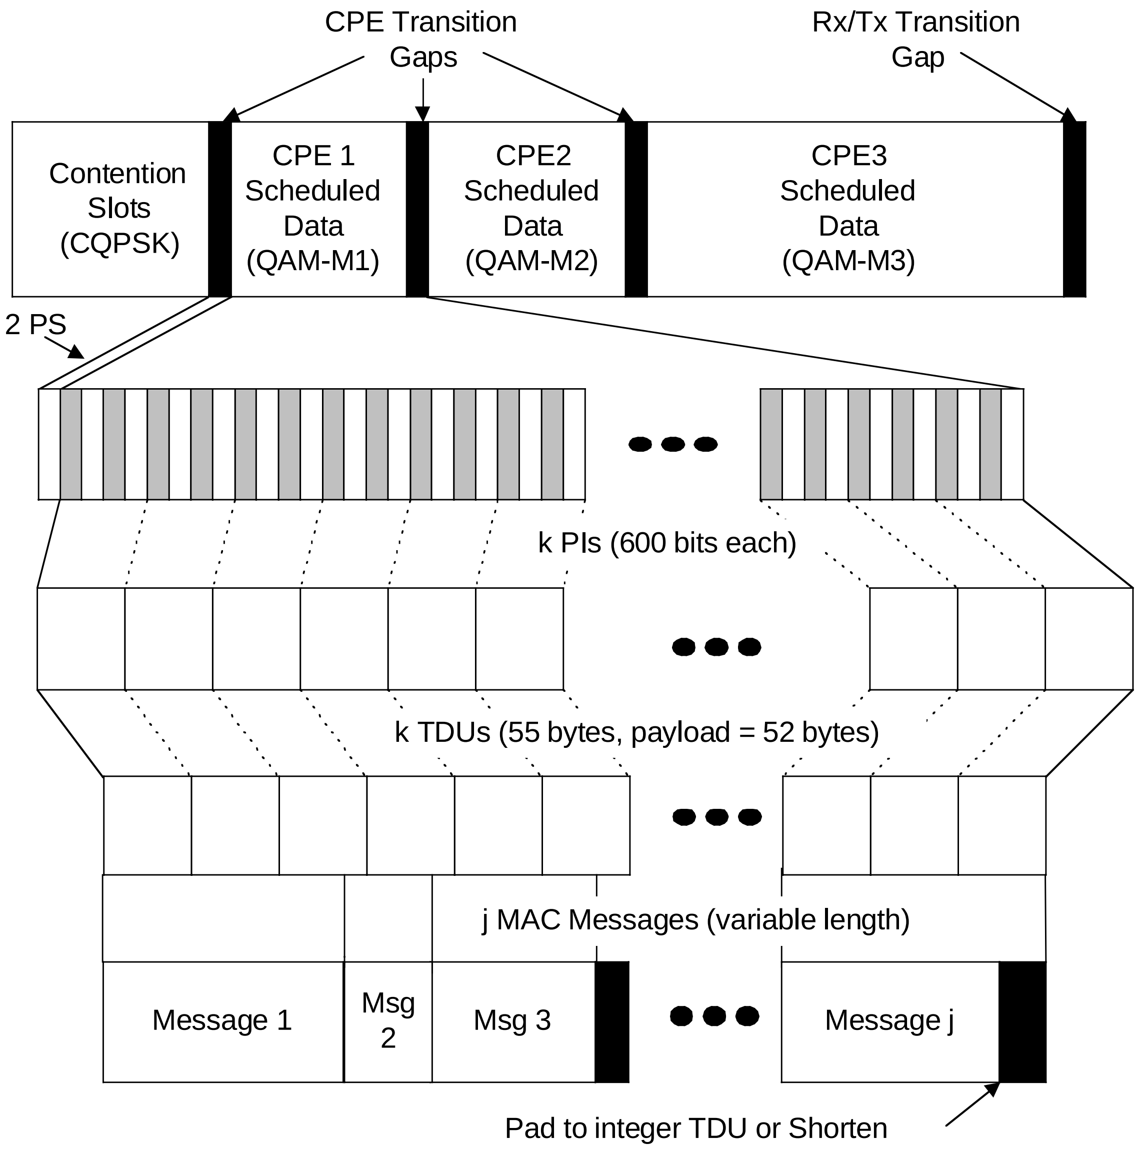
CPE Transition Gaps
Rx/Tx Transition Gap
Contention Slots (CQPSK)
CPE 1 Scheduled Data (QAM-M1)
CPE2 Scheduled Data (QAM-M2)
CPE3 Scheduled Data (QAM-M3)
2 PS
k PIs (600 bits each)
k TDUs (55 bytes, payload = 52 bytes)
j MAC Messages (variable length)
Message 1
Msg 2
Msg 3
Message j
Pad to integer TDU or Shorten

- Figure 3.2-19: Uplink Mapping of MAC Messages to PHY Elements

3.3 MAC Service Access Points

The architecture of the MAC is described in section 2.2. The protocol between the MAC and the higher layers expects a reliable link, such as TCP/IP. The protocol does not include Acks unless there is a need to return data.

The messages sent on service access points are tagged with a 1 byte code that identifies the service access point. These identifying codes are shown in Table 1.

- Table 1 Service Access Point Codes

Service Access Point	Service Access Point Identifier
Control SAP	0x01
CG SAP	0x02
DAMA SAP	0x03

3.3.1 Control Service Access Point

The Control SAP provides:

- The capability to exchange control information between the higher layers and the MAC.

There is a variety of information not related to a specific connection that must pass between the MAC and the higher layers. Connection specific information transfer is supported by the SAP for the relevant service type. The following messages are passed between the MAC and higher layers of the **BS** via the Control SAP. The message formats are implementation specific and are not subject to standardization.

1. The Register message is sent from the MAC to the higher layers when a CPE attempts registration.

2. The RegisterAck message is sent from the higher layers to the MAC in response to the Register message.

3. The RegisterComplete message is sent from the MAC to the higher layers at completion of registration.

4. The Deregistration is sent from higher layers to MAC when a CPE is leaving the system.

5. Keys are sent to the MAC from the higher layers.

6. The TDDSplit message is sent to the MAC from higher layers when the TDD split changes (TDD systems only).

7. The CellSizeChange Message is sent to the MAC from the higher layers when the cell size changes from the default.

8. Link Usage and Congestion Statistics are sent to the higher layers from the MAC.

The following primitives are passed between the MAC and higher layers of the **CPE** via the Control SAP. The message formats are implementation specific and are not subject to standardization.

1. The RegisterComplete is sent from CPE MAC to higher layers.

2. The Deregister message is sent from CPE higher layers to the MAC.

3. Keys are sent to the CPE MAC from the higher layers.

3.3.2 CG Service Access Point

The CG SAP provides the capability for the establishment and maintenance of CG connections, such as T1, fractional T1, E1, and other TDM connections. Compression of the CG link, by transmitting data only from active channels, can be used to reduce the air bandwidth requirements of TDM connections. This compression is performed by higher layers and is transparent to the MAC except for the concept of a current bandwidth and a maximum bandwidth for otherwise fixed rate services.

The capabilities provided by this SAP are:

- Connection establishment (always BS stimulated)
- Change in provisioning (always BS stimulated)
- Modification of current bandwidth by this side of link
- Modification of current bandwidth by the other side of the link
- Connection termination (always BS stimulated)
- Transfer of continuous grant data

The message formats are implementation specific and are not subject to standardization.

3.3.3 DAMA Service Access Point

The DAMA SAP provides the capability for the establishment and maintenance of ATM connections, packet connections, and the transfer of higher layer control messages via a CPE's higher layer control connection.

The capabilities provided to higher layers by the DAMA SAP are:

- Connection establishment
- Change in provisioning
- Connection termination
- DAMA data transfer

The message formats are implementation specific and are not subject to standardization.

3.3.4 QoS Level Definitions

The different QoS are described below.

- Table 2: QoS Descriptions

QoS	Class	Description & Fairness Algorithm
1	CG	Continuous Grant – all data available in a frame is sent that frame
2	MAC	Reserved for BS internal use.
3	MAC	Reserved for MAC Control Messages – Equally weighted, piggybacking Is preferred method of transport.
4-64	Real Time DAMA with ageing	Fair weighted queuing with weights derived real-time as a function of data pending. Aging may raise QoS of data.
65-127	Guaranteed Rate DAMA	Fair weighted queuing with weights derived real-time as a function of data pending.
128-199	Average Rate DAMA	Fair weighted queuing with weights statically derived from contracted average bandwidth
200	HLCM	Reserved for Higher Layer Control Connections – round robin with maximum burst size per frame.
201-255	Best Effort DAMA	Round Robin

3.4 Procedures for Media Access Control

The following timers are used in the implementation of MAC processes and protocols:

- Table 3: MAC Related Timers

Name	Description	Duration
PT1	The time allowed for a CPE to acquire downlink time and frequency synchronization on a physical channel before moving to the next.	**TBD**
PT2	The time waited by the CPE for hardware loops to lock. (Between first demodulation of the Frame Control Header and first transmit of registration request.)	**TBD**
MT1	The number of frames a CPE waits for a registration opportunity while registering before moving to the next physical channel.	**TBD**
MT2	The number of frames waited for response to a registration message.	**TBD**
MT3	The number of frames waited for response to control messages that do not have message specific timers.	**TBD**
MT4	The number of frames waited for response to setting the Poll Me bit before the CPE sets it again.	**TBD**
MT5	The number of frames waited for response to a bandwidth request sent via piggybacking or in response to an individual poll before re-sending the bandwidth request.	**TBD**
MT6	The number of frames waited for response to a bandwidth request sent in multicast or broadcast bandwidth request contention slots before starting the contention resolution process.	**TBD**
MT7	The number of frames a CPE waits for a registration opportunity while ranging before moving to the next physical channel.	**TBD**
MT8	The number of frames waited for a response to a ranging message.	**TBD**

3.4.1 Access and CPE Registration

3.4.1.1 General Process

The procedure by which a CPE joins the network is necessarily a complex one. A CPE must join the system on many levels, each frequently depending on the previous one, before it can be a full participant. These levels include:

Physical:	determining frequency channel
	Acquiring signal and framing
TC:	determination of modulation scheme
	Distance ranging
	Power level adjustment
MAC:	validation of CPE ID
	Provision of basic connection ID
Security:	authentication of user via public key
	Establishment of session key
Session:	establishment of additional connection IDs and QoS

When a CPE powers up, it searches for a frequency on which it can synchronize and locate the downlink subframe.

Having achieved downlink frame synchronization, the CPE waits for a registration opportunity. This may not occur in every frame. After the expiration of the MT1 timer, the CPE should try other frequencies. If no registration opportunity can be found at any frequency, the CPE should send back to the user an error indicating that no initial registration opportunity was found.

In the registration opportunity, the user CPE sends a registration request containing its 48 bit CPE ID and the PS in which it sent the message, relative to the start of the contention period.

The CPE then awaits a registration results message from the base station validating its CPE ID and providing a basic connection ID and a control connection ID. This message is sent, using CQPSK modulation, to a multicast connection ID reserved for unregistered CPEs, (reserved connection IDs are listed in section 3.2.1) with the CPE ID returned in a data field. The basic connection ID allows MAC to MAC communication between the BS and CPE. The control connection ID allows higher layer to higher layer communication between the BS and the CPE.

If the result code returned to the user is 0x03, this is an indication that the CPE ID supplied was invalid or that the CPE has somehow accessed the wrong BS. In this case the CPE reports an error to the user indicating a potential problem and continues the attempt to register on the next most likely channel. If the result code is 0x01, then the CPE is authorized and the CPE proceeds with ranging after applying the timing offset and power adjustment specified. If the result code is 0x02, the CPE is authorized but is requested to move to a different physical channel. The CPE applies the timing offset and power adjustment, but continues registration on the specified alternative channel.

If a registration results message is not received, but a registration collision message is received (available only with base stations capable of detecting collisions), the CPE assumes its message collided with another CPEs message. The CPE uses a slotted ALOHA contention resolution procedure to avoid continuing collisions.

If neither a registration results message nor a registration collision message is received from the BS, either an undetected collision occurred or the CPE transmitted with too low power. The CPE first tries progressively raising its power on subsequent attempts. If it reaches maximum power without receiving a message from the BS, the CPE uses a slotted ALOHA contention resolution procedure to avoid continuing collisions.

Figure 3.4-1 through Figure 3.4-3 show the overall process for CPE system access and registration. The details of timing acquisition are described in section 3.4.1.2. The details of ranging are described in section 3.4.1.3. The details of power leveling are described in section 3.4.1.4.

The message sequence for registration is shown in Figure 3.4-4.

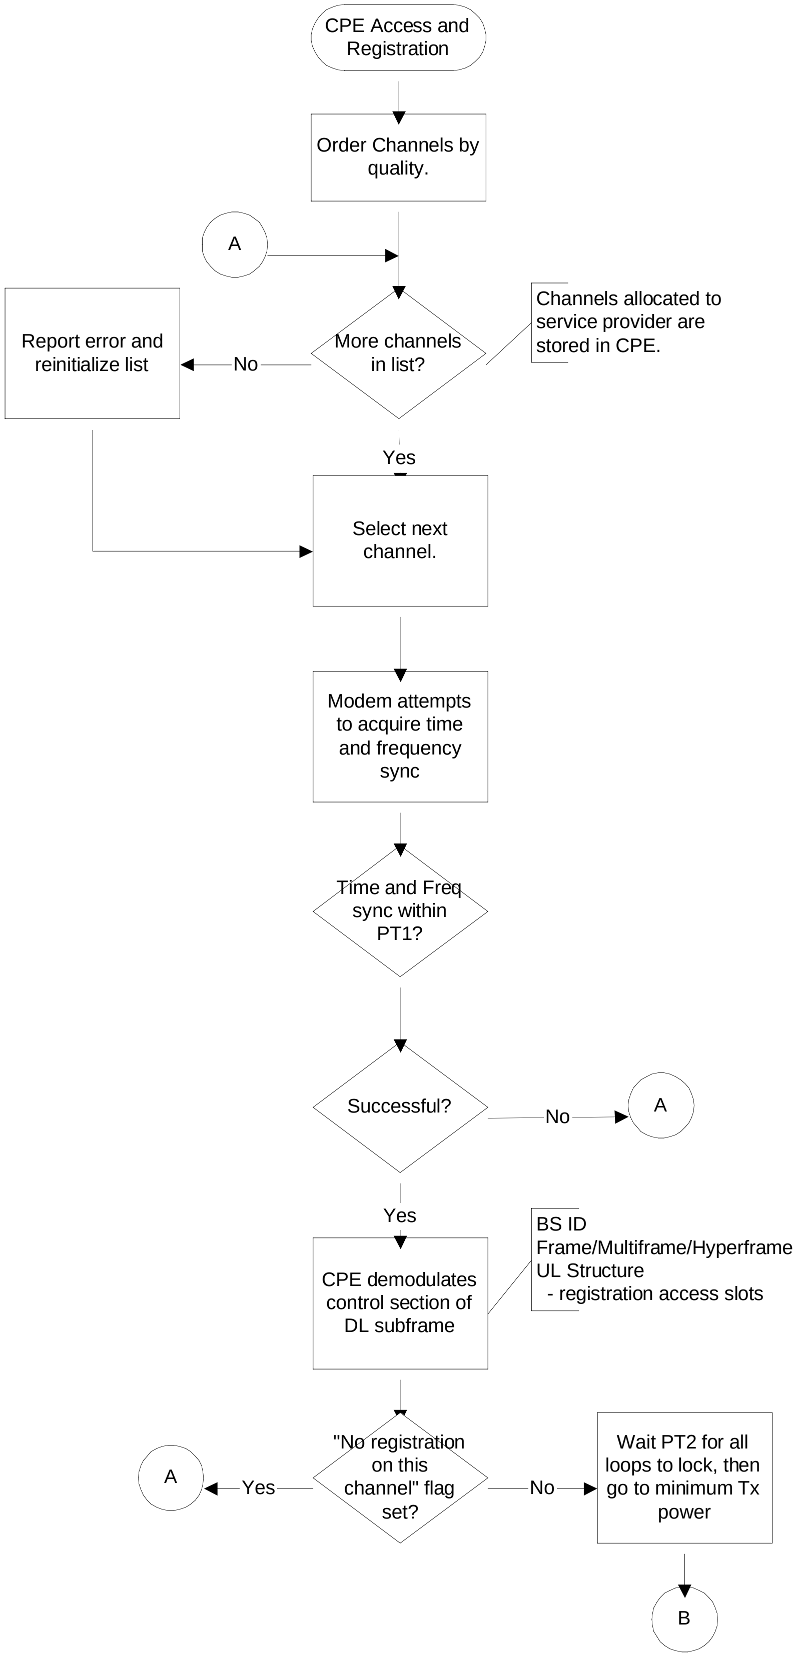
CPE Access and Registration
Order Channels by quality.
A
More channels in list?
No
Report error and reinitialize list
Channels allocated to service provider are stored in CPE.
Yes
Select next channel.
Modem attempts to acquire time and frequency sync
Time and Freq sync within PT1?
Successful?
No
A
Yes
CPE demodulates control section of DL subframe
BS ID
Frame/Multiframe/Hyperframe
UL Structure
- registration access slots
"No registration on this channel" flag set?
Yes
A
No
Wait PT2 for all loops to lock, then go to minimum Tx power
B

- Figure 3.4-1: CPE Registration

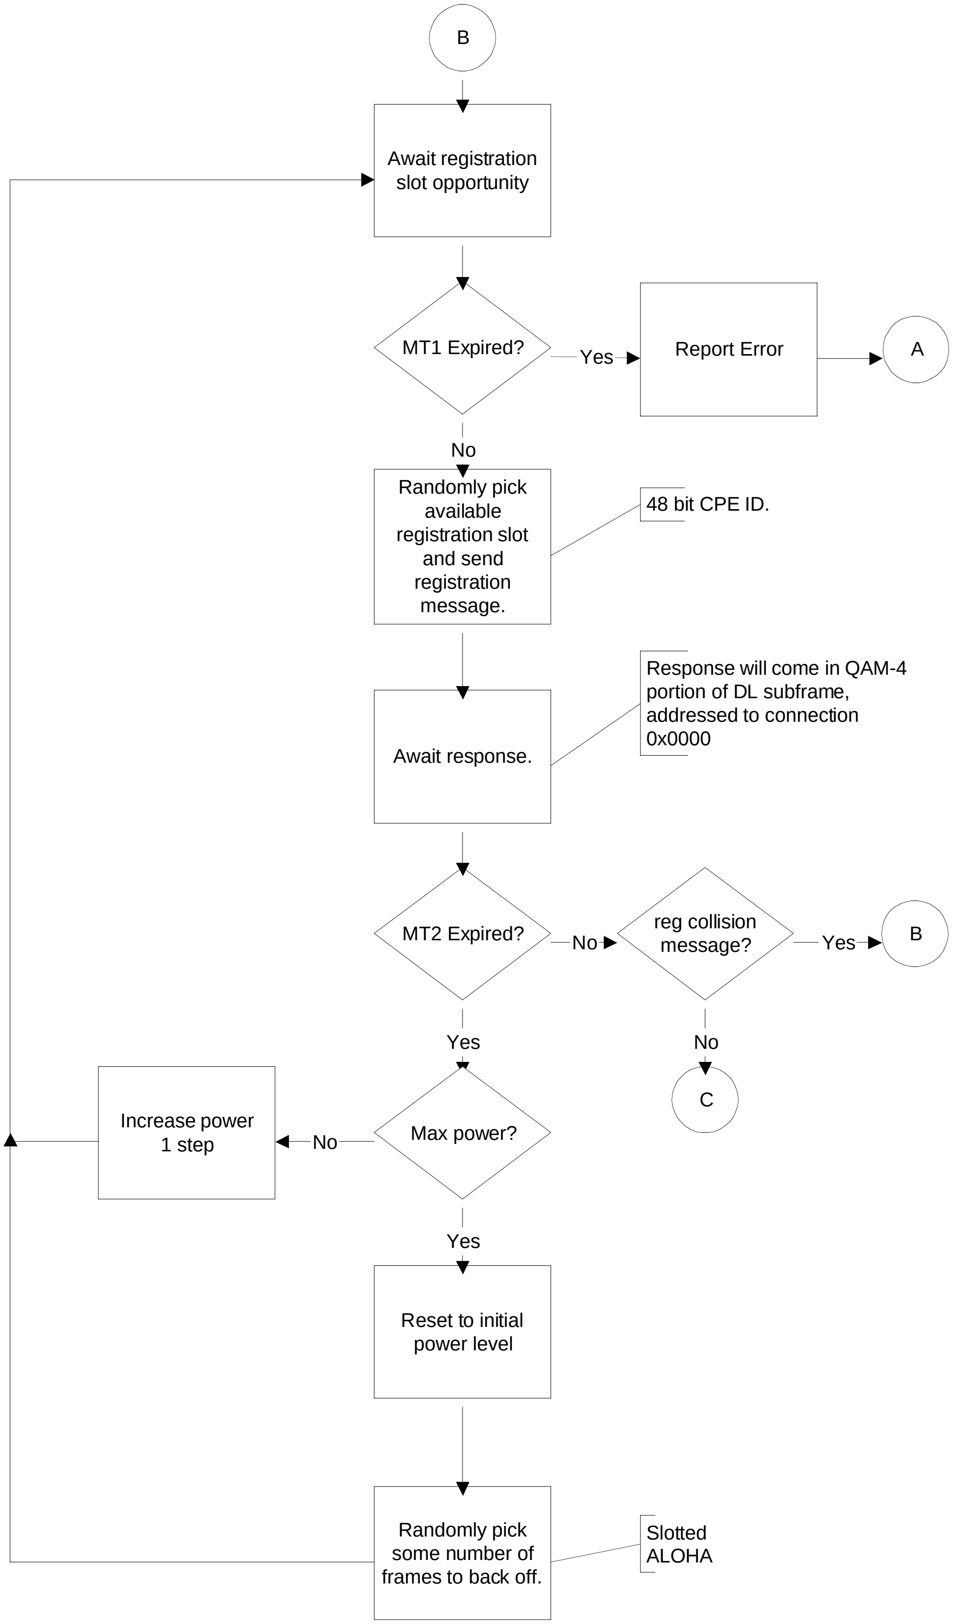
B
Await registration slot opportunity
MT1 Expired?
Yes
Report Error
A
No
Randomly pick available registration slot and send registration message.
48 bit CPE ID.
Await response.
Response will come in QAM-4 portion of DL subframe, addressed to connection 0x0000
MT2 Expired?
No
reg collision message?
Yes
B
Yes
No
C
Increase power 1 step
No
Max power?
Yes
Reset to initial power level
Randomly pick some number of frames to back off.
Slotted ALOHA

- Figure 3.4-2: CPE Registration, continued.

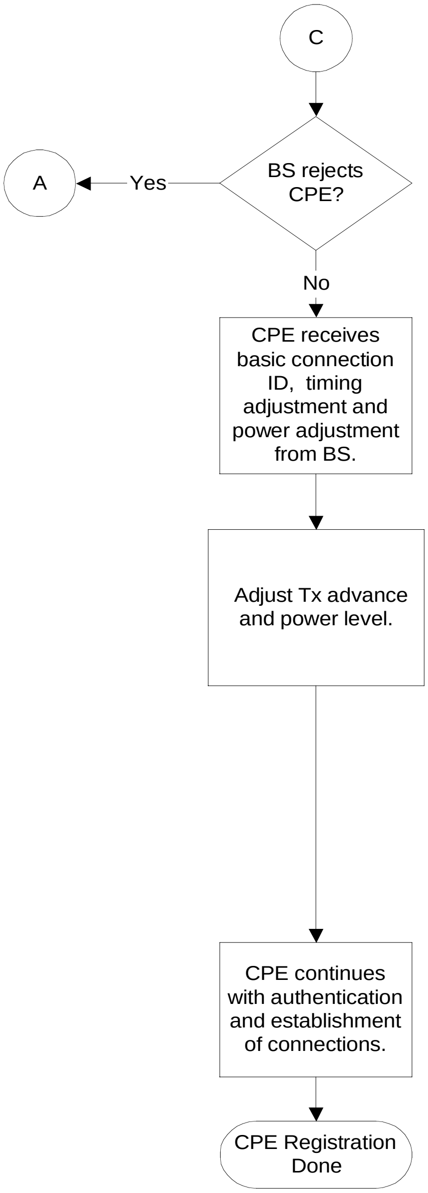
C
BS rejects CPE?
Yes
A
No
CPE receives basic connection ID, timing adjustment and power adjustment from BS.
Adjust Tx advance and power level.
CPE continues with authentication and establishment of connections.
CPE Registration Done

- Figure 3.4-3: CPE Registration, continued

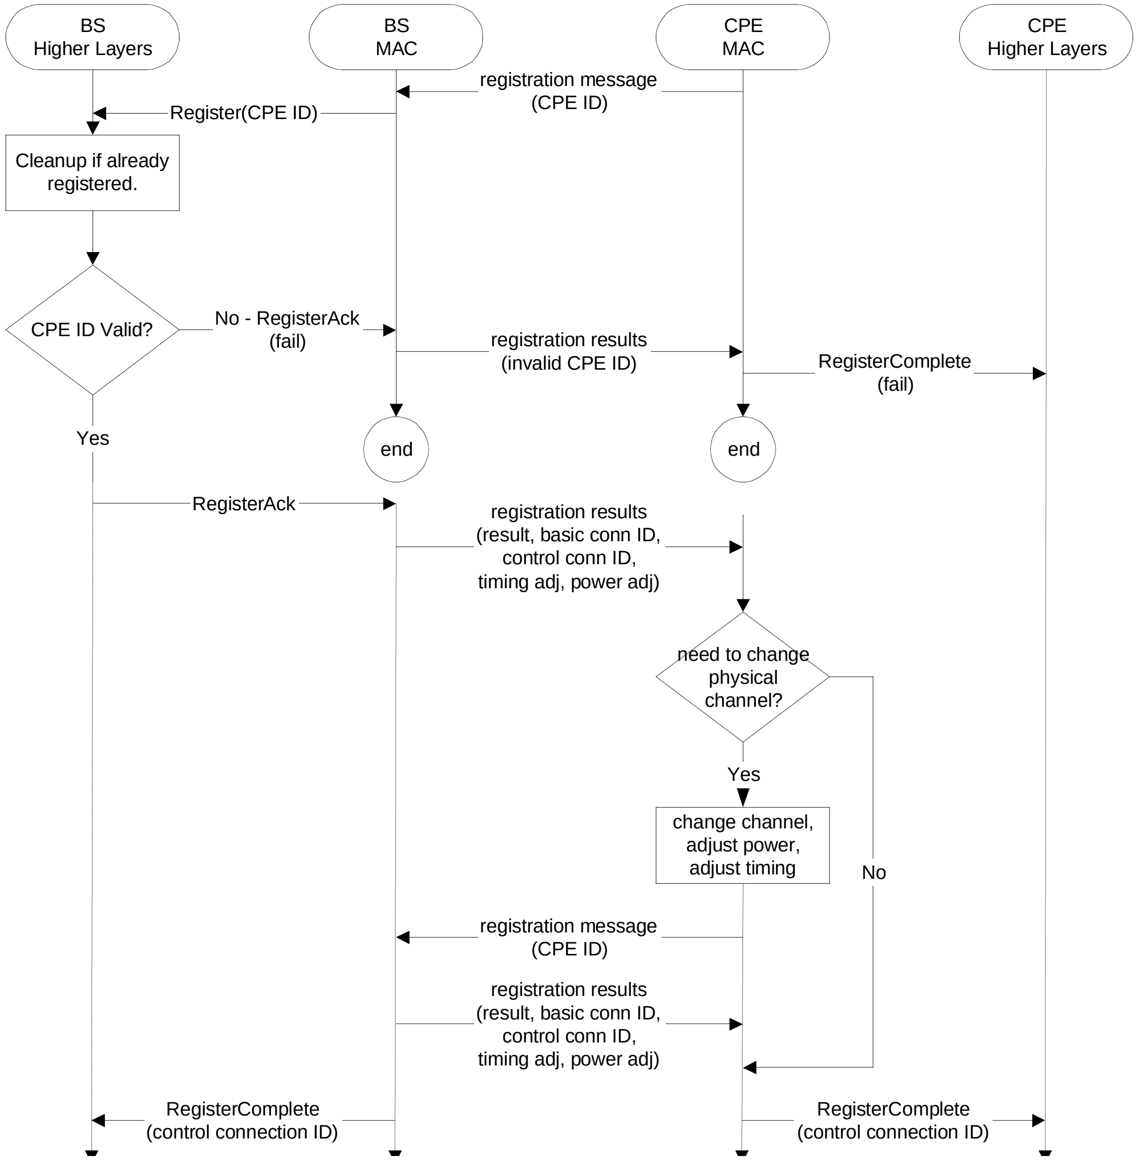

• Figure 3.4-4: CPE Registration High Level Message Sequence

3.4.1.2 Time and Frequency Acquisition

When a CPE first powers on, it needs to acquire synchronization with the BS downlink. This requires that the CPE command the modem to scan candidate frequencies from a pre-initialized list of the physical channels, belonging to the service provider, which the CPE has ranked in quality order.

The modem synchronizes, in time and frequency, to the preamble of the downlink subframe. How the modem accomplishes this is implementation dependant and is outside the scope of this document. The synchronization must be complete within the expiration of timer PT1 or the CPE moves to the next candidate frequency and tries again. Since time synchronization is on the downlink frame preamble, time synchronization automatically provides frame synchronization. The PHY Control portion of the downlink subframe contains the hyperframe, multiframe, and frame numbering information as well as BS identification information.

3.4.1.3 Ranging

Once the CPE has achieved downlink synchronization, it sets timer PT2 and awaits its expiration, allowing the hardware to stabilize. Then the CPE may transmit in the registration contention slots in the uplink subframe.

During the initialization process, the CPE must undergo ranging to determine its Tx timing advance before it is allowed to transmit in any part of the uplink subframe other than the registration contention slots. There are two parts to the ranging process. In the first part, in response to a registration opportunity, the CPE transmits a registration message described in section 3.6.3.1. Ranging is performed by the BS. It observes the arrival delay of the CPE transmission relative to the nominal start of the PS reported by the CPE in the registration message. This represents the two-way transmission time for the CPE's signal, since the CPE's assumption of when the frame starts is actually off by the initial propagation time of the downlink subframe. Round-trip delay information is sent back to the user via a registration results management message, addressed to the unregistered CPEs multicast group (connection ID 0x0000). This message is transmitted using QAM-4 modulation in the downlink subframe. The CPE applies the specified timing offset and power adjustment and proceeds.

In the second part of the ranging process, the BS may also periodically request a CPE to adjust its Tx advance using the ranging adjustment message.

3.4.1.4 Power Leveling

Because CPEs are at differing distances from the base station, their transmissions will be received at the BS at varying power levels depending on propagation losses due both to geographic and atmospheric-attenuation effects. As with the timing offset, the initial power adjustment is sent to the CPE in the registration results message. If the power adjustment is non-0, it is applied by the CPE.

The BS may also periodically request a CPE to increase or decrease its power level using the Power Adjustment message.

3.4.1.5 Registration Contention Resolution

If a CPE does not receive a registration results message within the expiration of timer MT2, or if it receives a registration collision message, it must assume that either a collision occurred or it transmitted the message with insufficient power. The first course of action taken by the CPE is to retry **TBD** times and then increase its transmit power, continuing until either it reaches maximum power, or until it receives a registration results message. In these cases, it uses a slotted ALOHA algorithm to wait a random number of registration opportunities before continuing registration.

3.4.1.6 CPE De-registration

When a CPE powers off for any reason, it must de-register. This allows for graceful discontinuing of service to the CPE. Additionally, CPEs may leave the system due to changes in provisioning. The message sequence for CPE de-registration is shown in Figure 3.4-5.

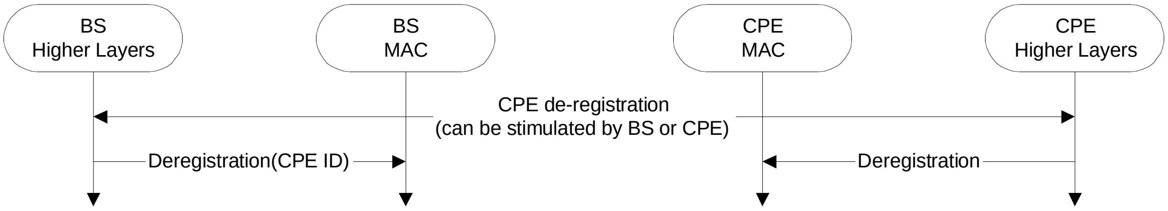

• Figure 3.4-5: CPE De-registration

3.4.2 Data Connections

Connections are established only in response to provisioning. Therefore, only the BS may initiate the establishment or termination of connections.

3.4.2.1 Connection Establishment

When a CPE registers with the BS, the BS MAC passes the CPE's ID to the higher layers. If there are connections provisioned for the CPE, the higher layers inform the BS HL-MAA of the total amount of guaranteed bandwidth (TDM, guaranteed rate ATM, etc.) provisioned to the CPE. This allows the HL-MAA to determine which physical channel to move the CPE to if necessary. Once the CPE is on the assigned physical channel, the BS HL-MAA allocates a DAMA connection to be used by the higher layers. The BS and CPE higher layers use this connection to perform such higher layer procedures as further authentication, key exchange, parameter download, etc.

Once the higher layers have concluded any higher layer registration protocols, the BS higher layers loop through the connections provisioned for the CPE and establish the connections. The BS HL-MAA generates an internal connection ID for each connection. The higher layers may exchange additional information regarding the connection. Once the connection establishment is complete, the higher layers my begin transferring user data on the connection.

Establishment of a connection is shown in Figure 3.4-6.

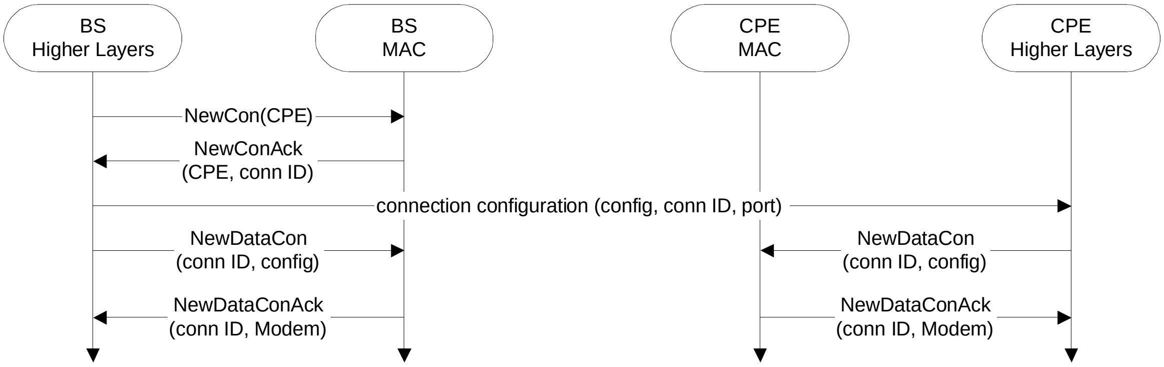

- Figure 3.4-6: Connection Establishment

A connection may also be added to the provisioning for a CPE. In this case, if the CPE is not yet registered, the BS higher layers save the information until the CPE registers. Otherwise, the connection is established identically to the establishment of the initially provisioned connections.

3.4.2.2 Steady State Connection

Connections may require ongoing maintenance. This maintenance is different for CG connections than for DAMA connection types.

CG connections are granted continuing bandwidth, meaning that a bandwidth request for a CG connection is a request for a constant amount of bandwidth on a periodic basis (usually every frame). Once a bandwidth request for a CG connection has been processed, there is no need for further bandwidth requests on that connection unless the bandwidth requirements of the connection change. If the bandwidth requirements of the CG connection change, either the CPE sends a bandwidth request to the BS, or the BS sends a notification to the CPE depending upon which side noticed the need for the change. For example, the particular CG connection may be a channelized T1. Figure 3.4-7 shows the sequence of events for a CPE detecting a DS0 going off hook. Figure 3.4-8 shows the sequence of events for the base station detecting a DS0 going off hook. Figure 3.4-9 shows the sequence of events for a CPE and the base station simultaneously detecting different DS0s going off hook.

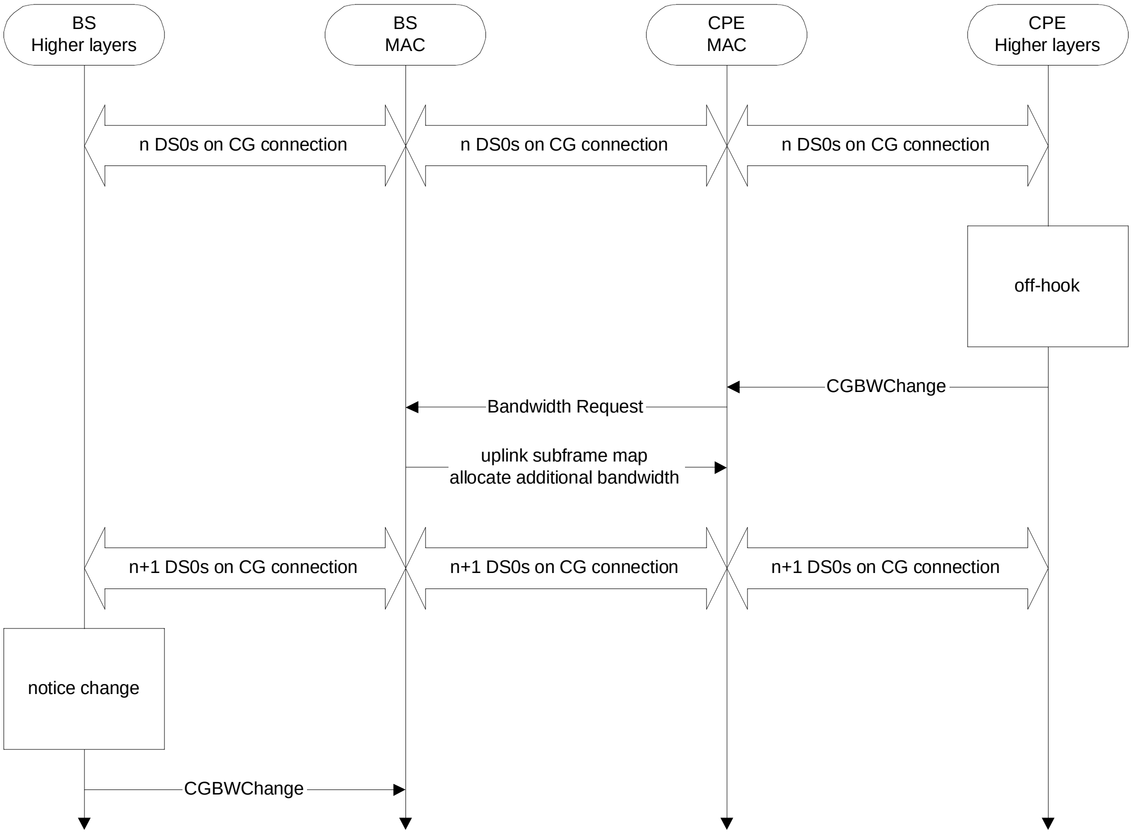

• Figure 3.4-7: CPE Initiated Continuous Grant Bandwidth Change

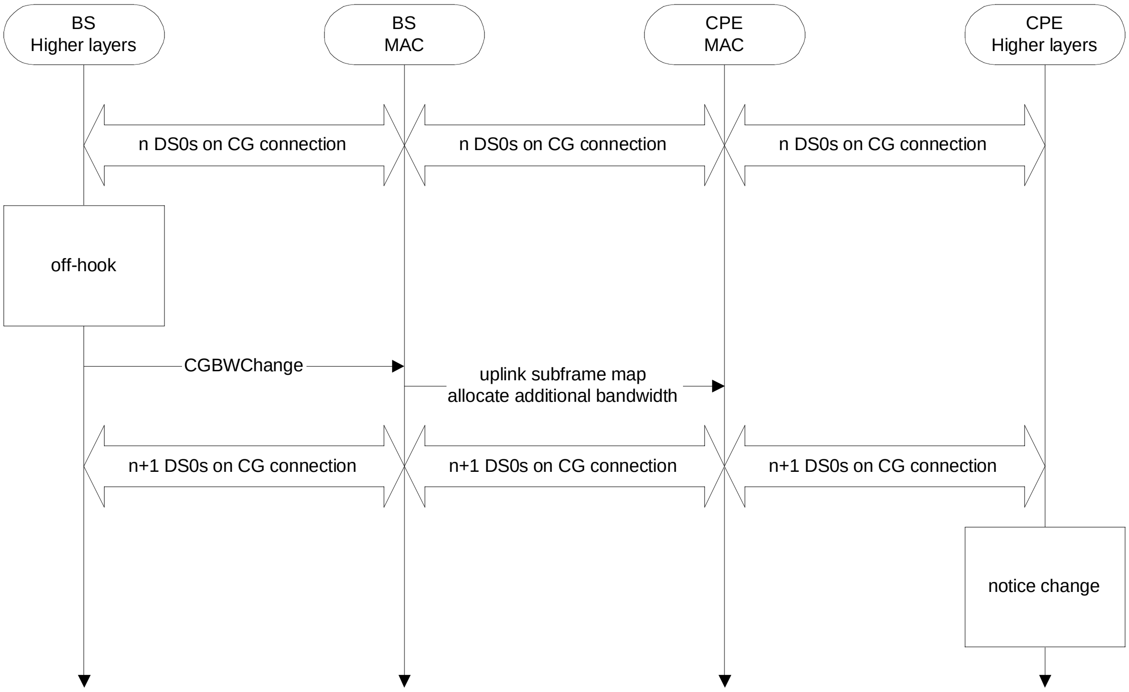

• Figure 3.4-8: Base Station Initiated Continuous Grant Bandwidth Change

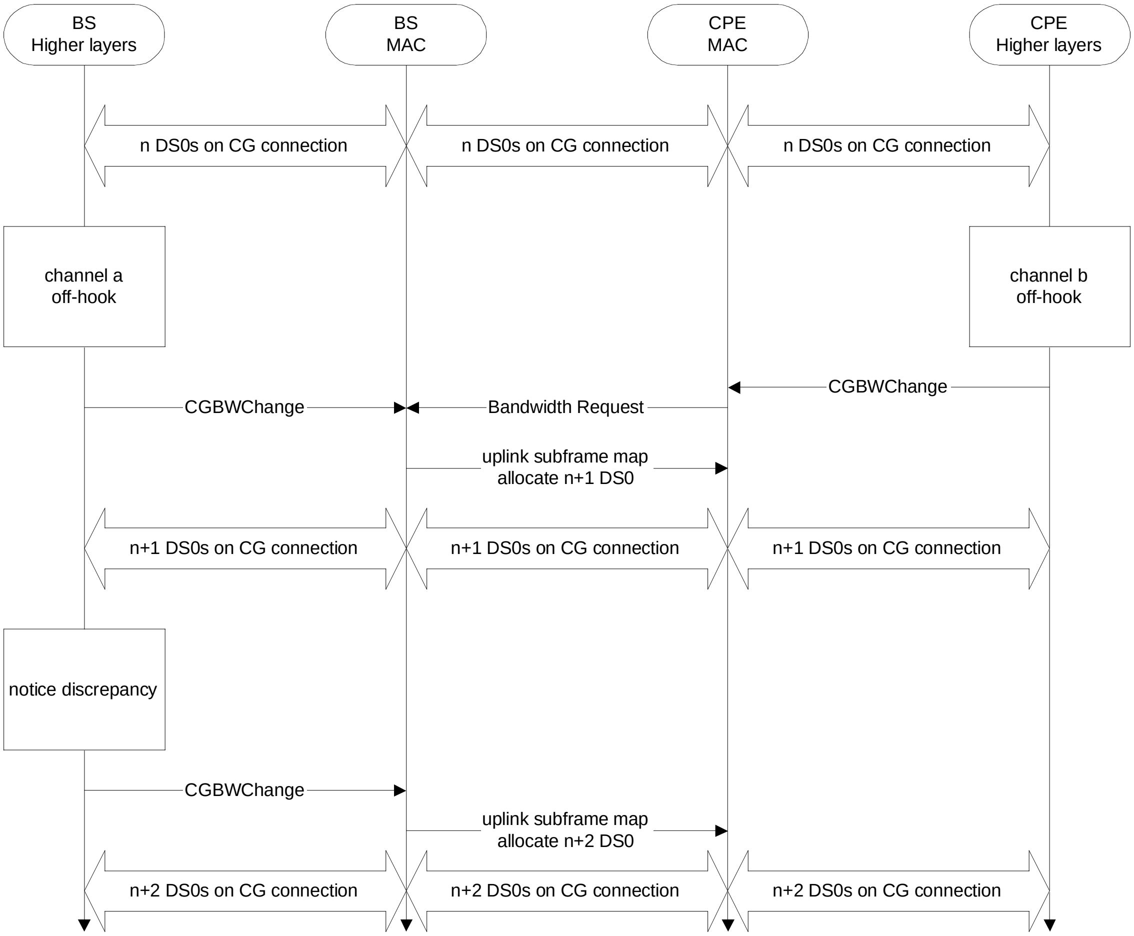

• Figure 3.4-9: Simultaneous Continuous Grant Bandwidth Change

DAMA connections are granted bandwidth on an as needed rather than periodic basis. On the downlink, the BS MAC simply uses the amount of data in its transmit queues in the bandwidth allocation algorithm. On the uplink, the CPE must request bandwidth either by piggybacking a bandwidth request or by waiting to be polled. The polling and bandwidth request process is described in section 3.4.3.

3.4.2.3 Connection Modification

Connections may be modified due to a change in provisioning. Connection modification is always initiated by the BS.

Modifications include:

- Termination
- Change in maximum bandwidth
- Change in guaranteed bandwidth

The information exchange sequence for provisioning change is shown in Figure 3.4-10.

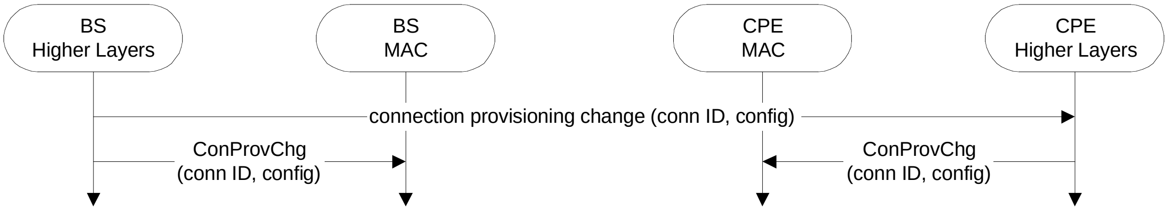

• Figure 3.4-10: Connection Provisioning Change

The information exchange sequence for connection termination is shown in Figure 3.4-11.

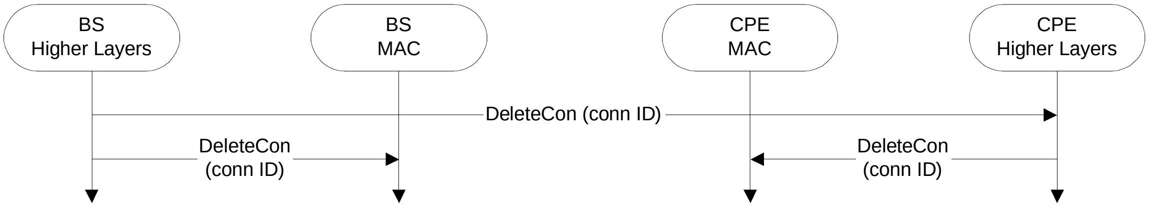

• Figure 3.4-11: Connection Termination

3.4.3 Polling/Bandwidth Requests

Note that at registration every CPE is assigned a dedicated connection ID for the purpose of sending and receiving control messages. Increasing (or decreasing) bandwidth requirements is necessary for all services except uncompressible constant bit rate CG services. The needs of uncompressible CG services do not change between connection establishment and termination. The requirements of compressible CG services, such as channelized T1, may increase or decrease depending on traffic. DAMA services are given resources on a demand assignment basis, as the need arises.

When a CPE needs to ask for bandwidth on a DAMA connection, it sends a message to the BS containing the immediate requirements of the DAMA connection. QoS for the connection was established at connection establishment and is looked-up by the BS.

There are numerous methods by which the CPE can get the bandwidth request message to the BS.

3.4.3.1 Polled

Polling is the process by which the BS allocates to the CPEs bandwidth specifically for the purpose of making bandwidth requests. These allocations may be to individual CPEs or to groups of CPEs. Allocations to groups of CPEs actually define bandwidth request contention slots (see section 3.4.3.1.2). The allocations are not in the form of an explicit message, but via an allocation (or increase) in the Uplink Map.

Note that polling is done on a CPE basis, bandwidth is requested on a connection ID basis, and bandwidth is allocated on a CPE basis.

3.4.3.1.1 Individual

When a CPE is polled individually, no explicit message is transmitted to poll the CPE. Rather, the CPE is allocated, in the Uplink Subframe Map, bandwidth sufficient to respond with a bandwidth request. If the CPE does not need bandwidth, it returns a request for 0 bytes (Note that 0 byte requests are only used in the individual polling case since explicit bandwidth for a reply has been allocated.). Active CPEs that do not set the Poll Me bit in some MAC packet header will not be polled individually. Only inactive CPEs and CPEs explicitly requesting to be polled will be polled individually. This saves bandwidth over polling all CPEs individually. Active CPEs respond to polling at their current uplink modulation, while inactive CPEs must respond at QAM-4 to ensure their transmission is robust enough to be detected by the BS.

The interpretation of bandwidth requests by the base station differs for CG connections and DAMA connections. For CG connections, the effect of a bandwidth request is to change the bandwidth allocated every frame. For DAMA connections, the effect is to reset the base station's perception of the data pending at the CPE for that connection.

The information exchange sequence for individual polling is shown in Figure 3.4-12.

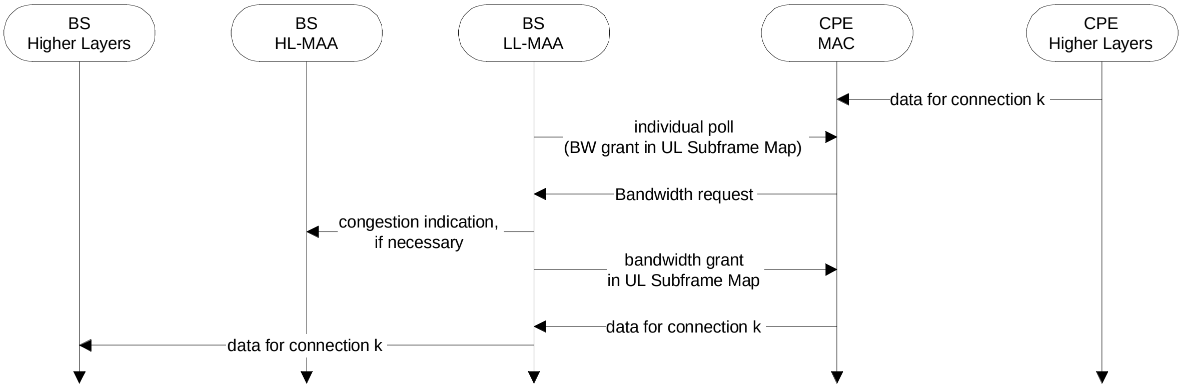

• Figure 3.4-12: Individual Polling

The individual polling process is shown in Figure 3.4-13.

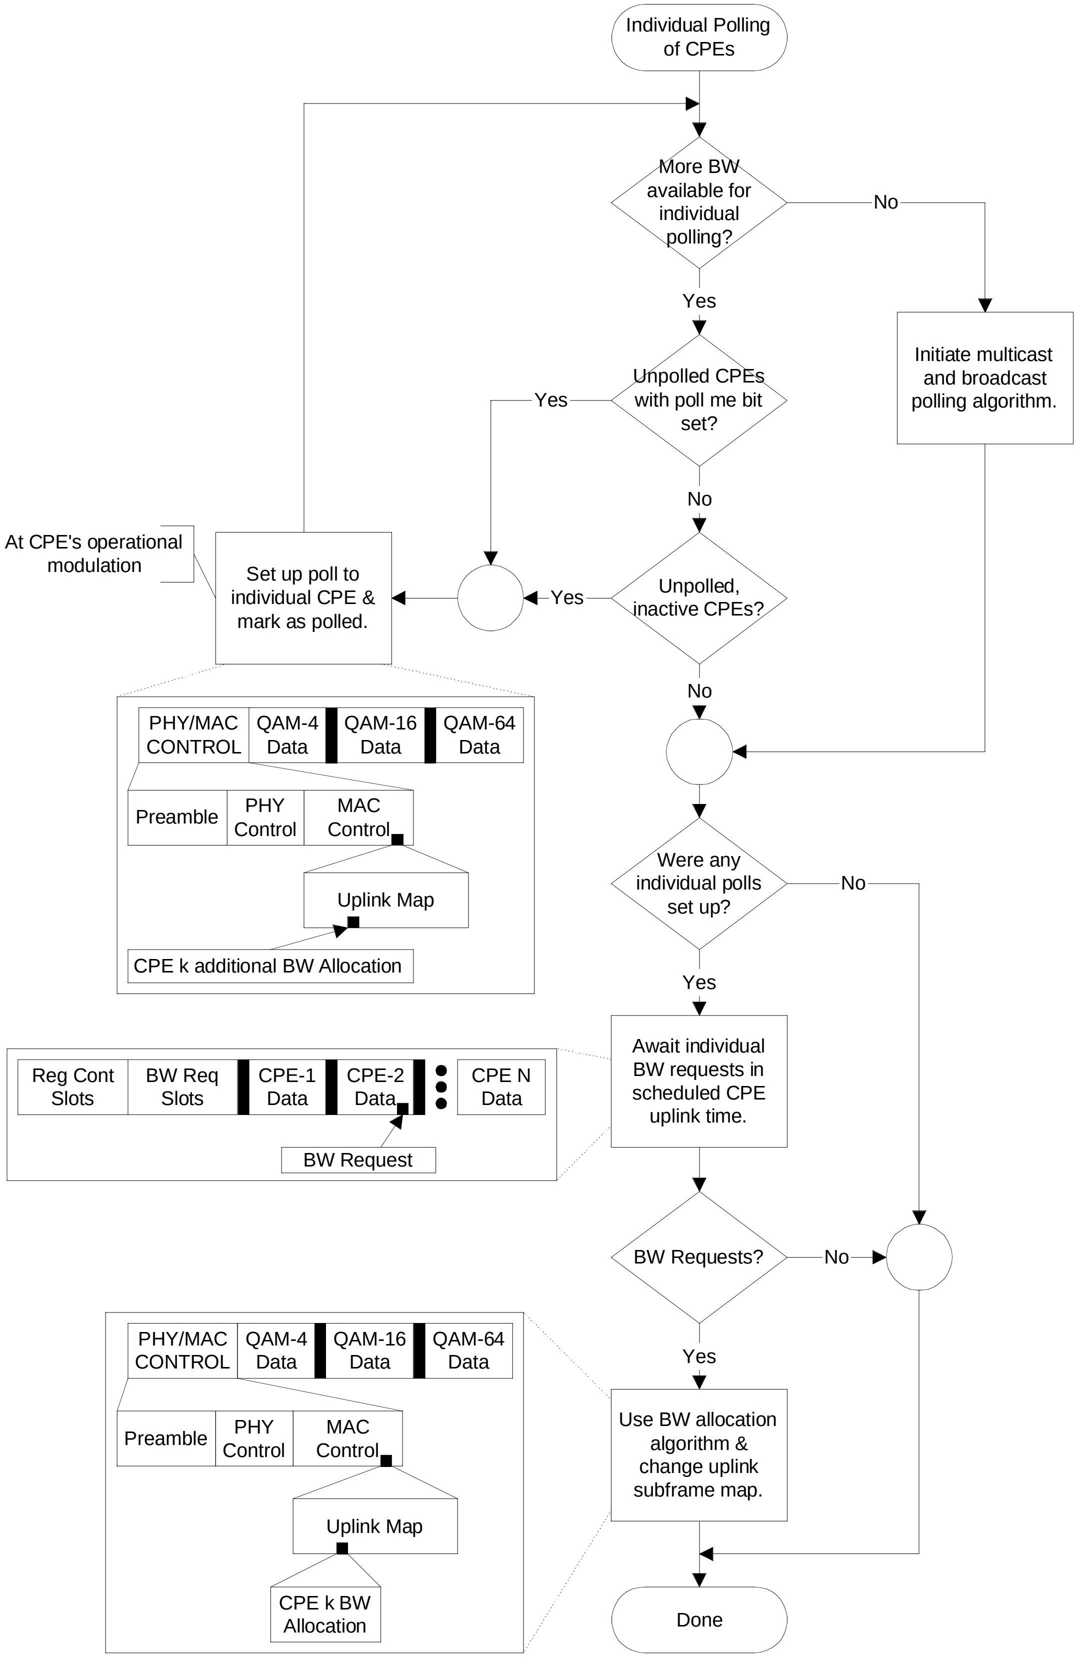
Individual Polling of CPEs
More BW available for individual polling?
No
Yes
Initiate multicast and broadcast polling algorithm.
Unpolled CPEs with poll me bit set?
Yes
No
Unpolled, inactive CPEs?
Yes
No
At CPE's operational modulation
Set up poll to individual CPE & mark as polled.
PHY/MAC CONTROL
QAM-4 Data
QAM-16 Data
QAM-64 Data
Preamble
PHY Control
MAC Control
Uplink Map
CPE k additional BW Allocation
Were any individual polls set up?
No
Yes
Await individual BW requests in scheduled CPE uplink time.
Reg Cont Slots
BW Req Slots
CPE-1 Data
CPE-2 Data
CPE N Data
BW Request
BW Requests?
No
Yes
Use BW allocation algorithm & change uplink subframe map.
PHY/MAC CONTROL
QAM-4 Data
QAM-16 Data
QAM-64 Data
Preamble
PHY Control
MAC Control
Uplink Map
CPE k BW Allocation
Done

- Figure 3.4-13: Individual Polling of CPEs

3.4.3.1.2 Multicast or Broadcast

If there are more CPEs that are inactive than there is bandwidth available for individual polling, some CPEs may be polled in multicast groups and a broadcast poll may be issued. Certain connection IDs are reserved for multicast groups and for broadcast messages, as described in section 3.2.1. As with individual polling, the poll is not an explicit message but bandwidth allocated in the Uplink Map. The difference is that rather than associating allocated bandwidth with a CPE's basic connection ID, the allocation is to a multicast or broadcast connection ID. This is shown in Figure 3.4-14.

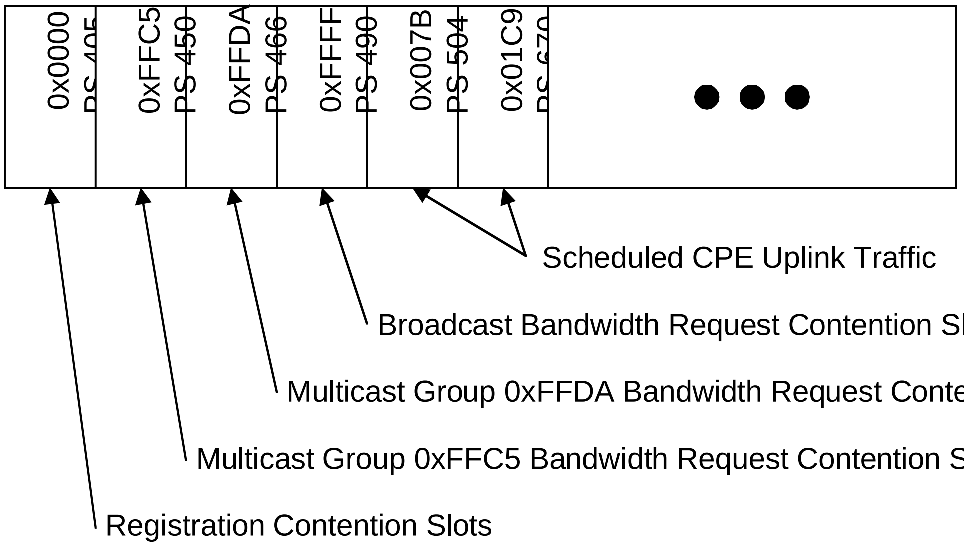
0x0000
PS 405
0xFFC5
PS 450
0xFFDA
PS 466
0xFFFF
PS 490
0x007B
PS 504
0x01C9
PS 670
●●●
Scheduled CPE Uplink Traffic
Broadcast Bandwidth Request Contention Slots
Multicast Group 0xFFDA Bandwidth Request Contention Slots
Multicast Group 0xFFC5 Bandwidth Request Contention Slots
Registration Contention Slots

- Figure 3.4-14: Uplink Map Structure

When the poll is directed at a multicast or broadcast connection ID, CPEs belonging to the polled group may request bandwidth using Bandwidth Request Contention Slots allocated in the uplink subframe. With multicast and broadcast polling, to reduce the likelihood of collision, only CPE's needing bandwidth reply. Zero-length bandwidth requests are not allowed in bandwidth request contention slots. CPEs always transmit using CQPSK modulation in the bandwidth request contention slots. The contention slots are sized to hold a 2 PS CTG and a bandwidth request message. The message requires a shortened TDU (as 1 PI = 15 PS is much greater in length than the actual data length required).

If an error such as an invalid connection ID occurs the BS sends an explicit error message to the CPE. If the BS does not respond with an error message or a bandwidth allocation within the expiration of timer MT5, the CPE assumes a collision has occurred and uses a slotted ALOHA scheme to back off and try at another contention opportunity. The multicast and broadcast polling process is shown in Figure 3.4-15.

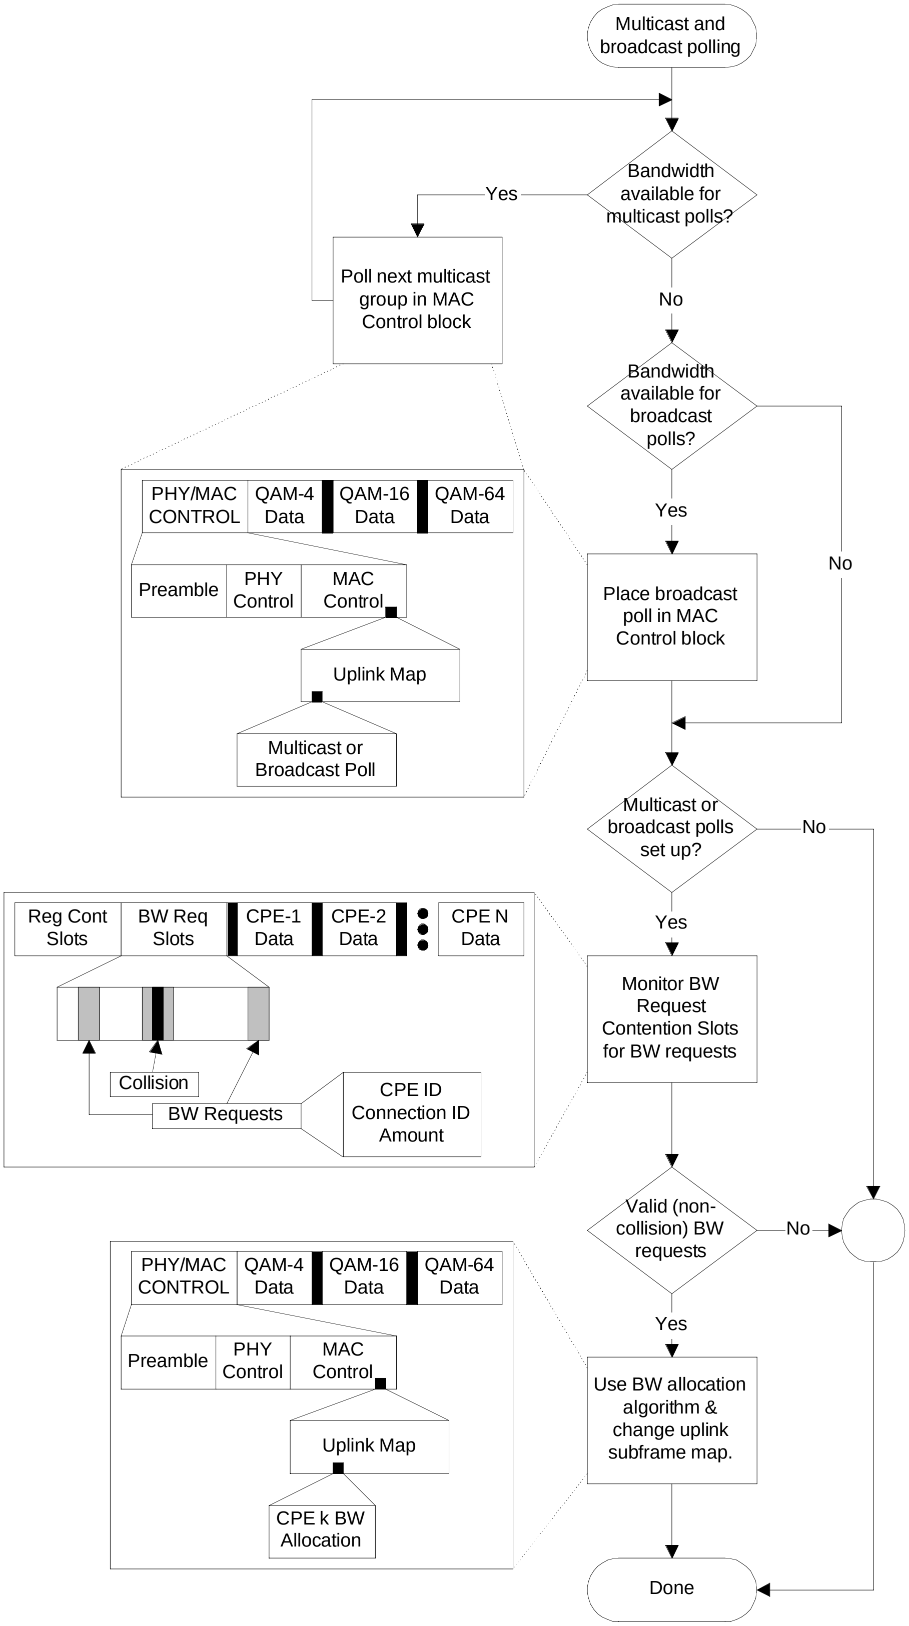
Multicast and broadcast polling
Bandwidth available for multicast polls?
Yes
No
Poll next multicast group in MAC Control block
Bandwidth available for broadcast polls?
Yes
No
Place broadcast poll in MAC Control block
PHY/MAC CONTROL
QAM-4 Data
QAM-16 Data
QAM-64 Data
Preamble
PHY Control
MAC Control
Uplink Map
Multicast or Broadcast Poll
Multicast or broadcast polls set up?
No
Yes
Monitor BW Request Contention Slots for BW requests
Reg Cont Slots
BW Req Slots
CPE-1 Data
CPE-2 Data
CPE N Data
Collision
BW Requests
CPE ID
Connection ID
Amount
Valid (non-collision) BW requests
No
Yes
Use BW allocation algorithm & change uplink subframe map.
PHY/MAC CONTROL
QAM-4 Data
QAM-16 Data
QAM-64 Data
Preamble
PHY Control
MAC Control
Uplink Map
CPE k BW Allocation
Done

- Figure 3.4-15: Multicast and Broadcast Polling of CPEs

3.4.3.1.2.1 Slotted Aloha Contention Resolution Process

Contention is necessary when there is insufficient time to poll all CPEs individually within a suitable interval. The BS is able to define contention periods both for multicast groups and also for all CPEs generally (broadcast).

After CPE scheduled data, control messages, and polling are allowed for; the base station allocates all unused time in the upstream part of the frame to contention, either for bandwidth requests or registration. User stations must transmit their requests at a random time within this interval to reduce the likelihood of collisions.

A CPE wishing to transmit in a request interval randomly selects a PS within the interval, and makes a request (in the associated starting PS). This randomization minimizes the probability of collisions.

A collision is presumed if there is no response from the base station to the request, within the expiration of timer MT5. If the BS does not respond within the expiration of timer MT5, then collision resolution is initiated.

The resolution process is as follows; assuming that the initial back off parameter is *I* and that the final back off parameter is *f*.

1. On the first collision, the CPE waits a random interval between 0 and 2^I registration opportunities and then tries again.

2. If another collision occurs, then the interval is doubled and the CPE tries again, repeating until the interval 2^f is reached.

If the CPE is still unsuccessful, an error must be reported to the system controller and the contention process aborted.

Note: this contention resolution mechanism may be slightly less efficient than others such as the ternary tree method, but collisions will be a relatively rare occurrence and the simplicity of the exponential backoff process will outweigh any loss in efficiency.

The collision resolution process is also used when there is a collision during the initial registration process, which is open to all CPEs. In this case there are separate ranging limits on the start and end of the backoff process (timer MT2 is used instead of timer MT5, for example), but otherwise the same principles apply. Additionally, in the registration case, the BS notifies the CPEs of collisions to avoid the CPEs erroneously increasing their transmit power.

3.4.3.2 Poll Me Bit

Currently active CPEs may set the poll me bit (bit PM in the MAC header) or the priority poll me bit (bit PPM in the MAC header) in a MAC packet to indicate to the BS that they need polled to negotiate a bandwidth change. To reduce the bandwidth requirements of individual polling, active CPEs will be individually polled only if one of these bits is set. Once the BS detects this request for polling, the process for individual polling is used to satisfy the request. The procedure by which a CPE stimulates the BS to poll it is shown in Figure 3.4-16. To minimize the risk of the BS missing the poll me bit, the CPE may set the bit in all MAC headers in the frame.

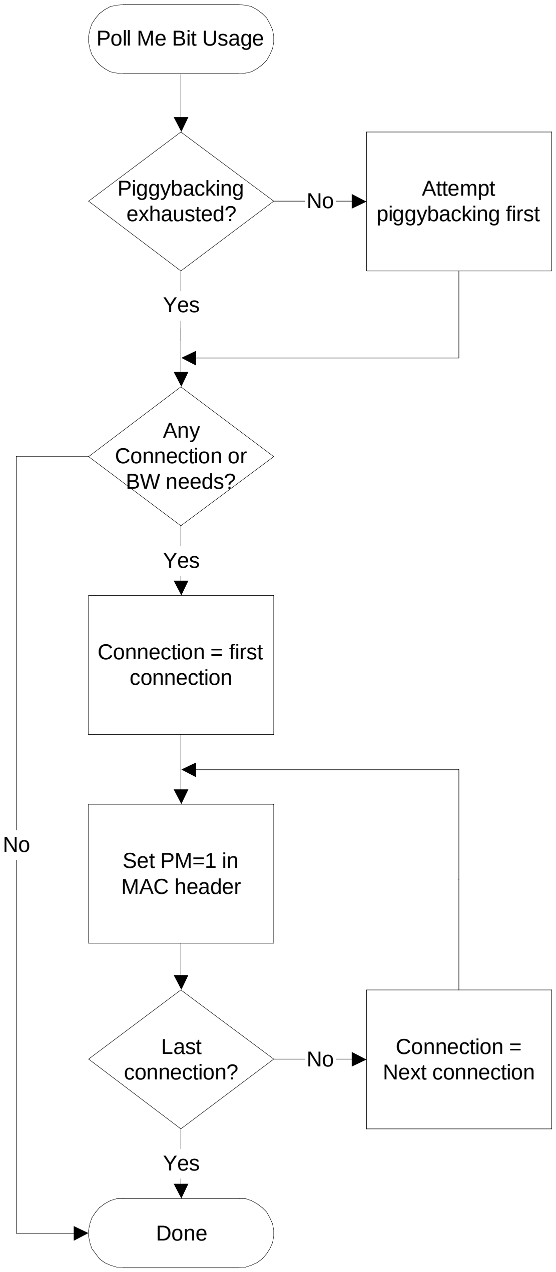
Poll Me Bit Usage
Piggybacking exhausted?
No
Attempt piggybacking first
Yes
Any Connection or BW needs?
Yes
No
Connection = first connection
Set PM=1 in MAC header
Last connection?
No
Connection = Next connection
Yes
Done

- Figure 3.4-16: Use of the Poll Me Bit to Stimulate Polling

The message sequence for requesting polling using a poll me bit is shown in Figure 3.4-17.

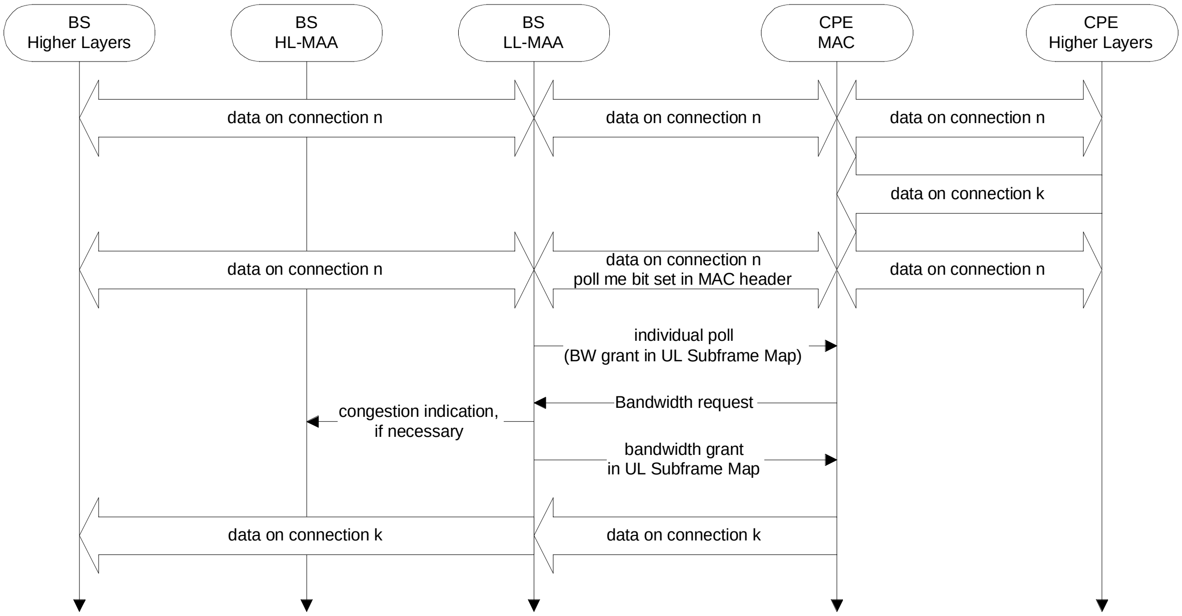

• Figure 3.4-17: Poll Me Bit Message Sequence

3.4.3.3 Piggybacking

To further reduce overhead bandwidth used by the bandwidth allocation process, currently active CPEs may piggyback a bandwidth request (or any other control message) on their current transmissions. They do this by using unused bandwidth in TDUs of existing allocations. The procedure for using excess bandwidth is shown in Figure 3.4-18.

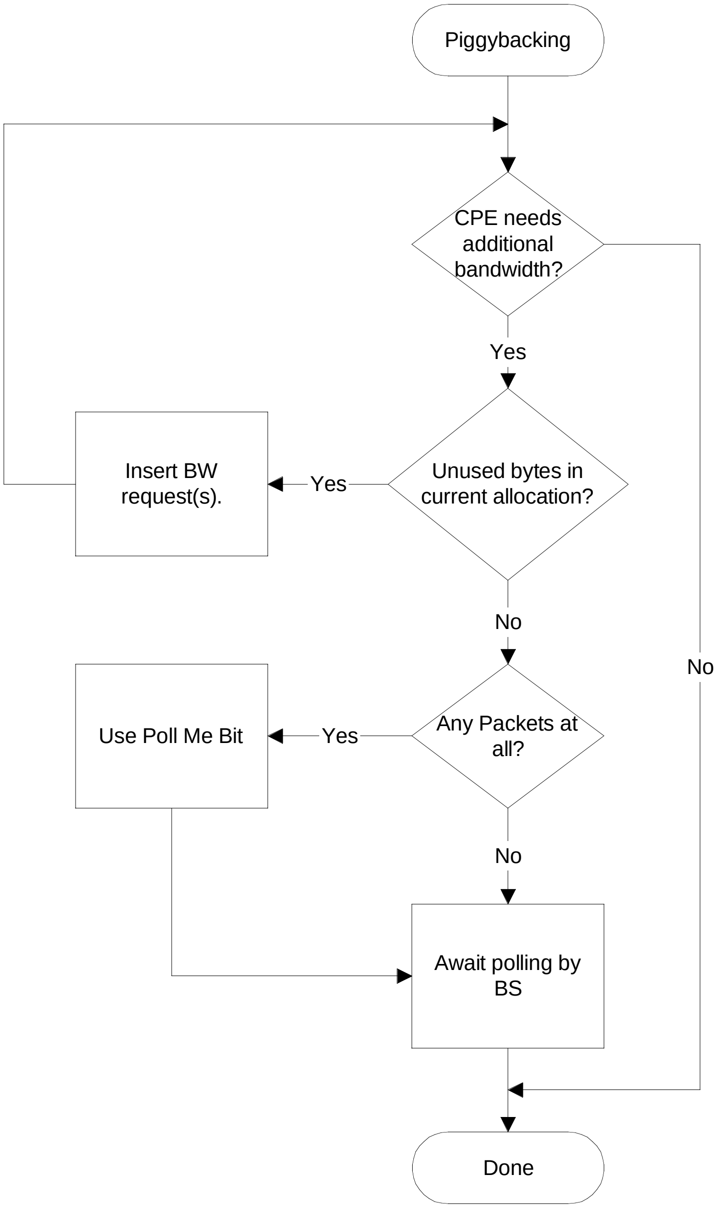

- Figure 3.4-18: Piggybacking Procedure

3.4.4 Bandwidth Allocation

The LL-MAA sublayer of the BS MAC is responsible for allocating the available bandwidth of a physical channel on the uplink and the downlink. Within the uplink and downlink subframes, the BS LL-MAA scheduler allocates the available bandwidth between the various services depending upon the priorities and rules imposed by their quality of service (QoS). Additionally, the HL-MAA sublayer of the BS MAC allocates across more than 1 physical channel.

3.4.4.1 Downlink Bandwidth Allocation – General

For each physical channel a set of queues, one for each QoS, hold the data ready to transmit to the CPEs on that physical channel. The higher layers of the BS protocol stack are responsible for the order in which data is place in the individual queues. The BS higher layers are free to implement any fairness algorithms or traffic shaping algorithms regarding the sharing of access between connections at the same QoS, without impacting the BS LL-MAA. Once the data is in the queues, it is the responsibility of the BS LL-MAA to allocate bandwidth based on the QoS. Separating these two algorithms allows the BS LL-MAA and the BS higher layers to concentrate on their own aspect of the bandwidth allocation problem independently.

3.4.4.2 Uplink Bandwidth Allocation – General

Uplink Bandwidth Allocation is very similar to downlink bandwidth allocation. The data queues, however, reside distributed across the individual CPEs.

The bandwidth allocated to a particular CPE, however, is sent in the form of a bandwidth allocation in the Uplink Map. The Uplink Map allocates a certain amount of bandwidth to a CPE, starting at a certain point in the frame. The CPE then allocates this bandwidth across its connections. This allows the CPE to use the bandwidth in a different manner than requested if it receives higher priority data while awaiting the allocation. Due to the dynamic nature of bandwidth allocation, the allocations are constantly changing. Because of this, a CPE may receive unsolicited modifications to the bandwidth granted on a frame by frame basis. If a CPE is allocated less bandwidth for a frame than is necessary to transmit all waiting data, the CPE must use the QoSs and their fairness algorithms to service its queues. The CPE may steal bandwidth from lower QoS connections to piggyback request for more bandwidth.

3.4.4.3 QoS Specific Fairness Algorithms

TBD

The fairness algorithms are vendor specific and may not be subject to standardization.

3.4.4.4 Bandwidth Allocation Algorithm

TBD

The bandwidth allocation algorithms are vendor specific and may not be subject to standardization.

3.4.5 Congestion

Due to the statistical multiplexing of connections on the physical channel, there is a possibility of congestion. Congestion may be handled in either of two ways. Based on contracted bandwidth, congestion may be predicted. If a physical channel is predicted to be congested due to the contracted rates of connections, CPEs may be handed off to another, less congested physical channel. Alternatively, if handover is not an option, or is predicted to not solve the congestion, marking of data may be used to determine which data may be dropped to relieve the congestion.

3.4.5.1 Handover to Another Physical Channel

TBD

3.4.5.2 Marking

For DAMA connections, data may be marked indicating whether it may be discarded when congestion occurs. During congestion, data marked for dropping, is dropped in reverse QoS order until the congestion is relieved.

3.4.6 CPE Modulation Change

When the BS requests from a CPE to transmit at a more robust modulation. To ensure receipt of the message by the CPE, the modulation change message is sent at the lower modulation. The CPE Acks at the new modulation. This sequence is shown in Figure 3.4-19.

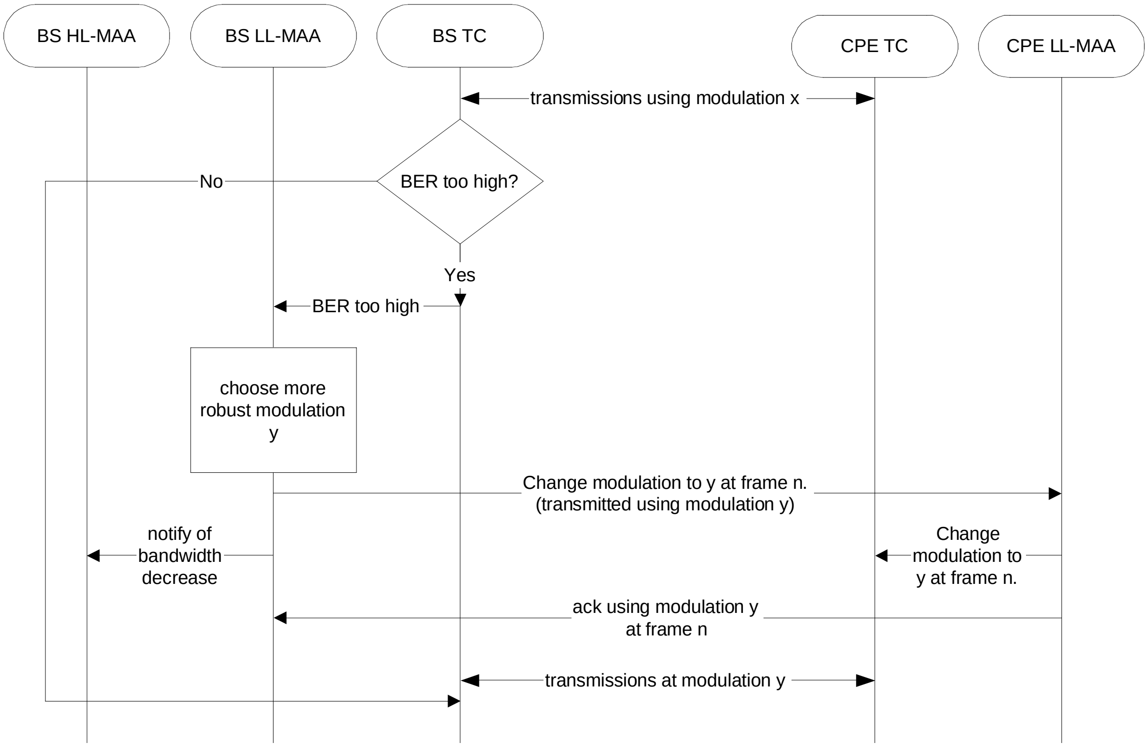
BS HL-MAA
BS LL-MAA
BS TC
CPE TC
CPE LL-MAA
transmissions using modulation x
No
BER too high?
Yes
BER too high
choose more robust modulation y
Change modulation to y at frame n. (transmitted using modulation y)
notify of bandwidth decrease
Change modulation to y at frame n.
ack using modulation y at frame n
transmissions at modulation y

• Figure 3.4-19: CPE Modulation Change – More Robust

When the BS requests from a CPE to transmit at a higher modulation rate then the modulation change message is sent at the old modulation rate of the CPE. In case the CPE is not capable of the new modulation, the CPE Acks at the old modulation, before the change is to take place. This sequence is shown in Figure 3.4-20.

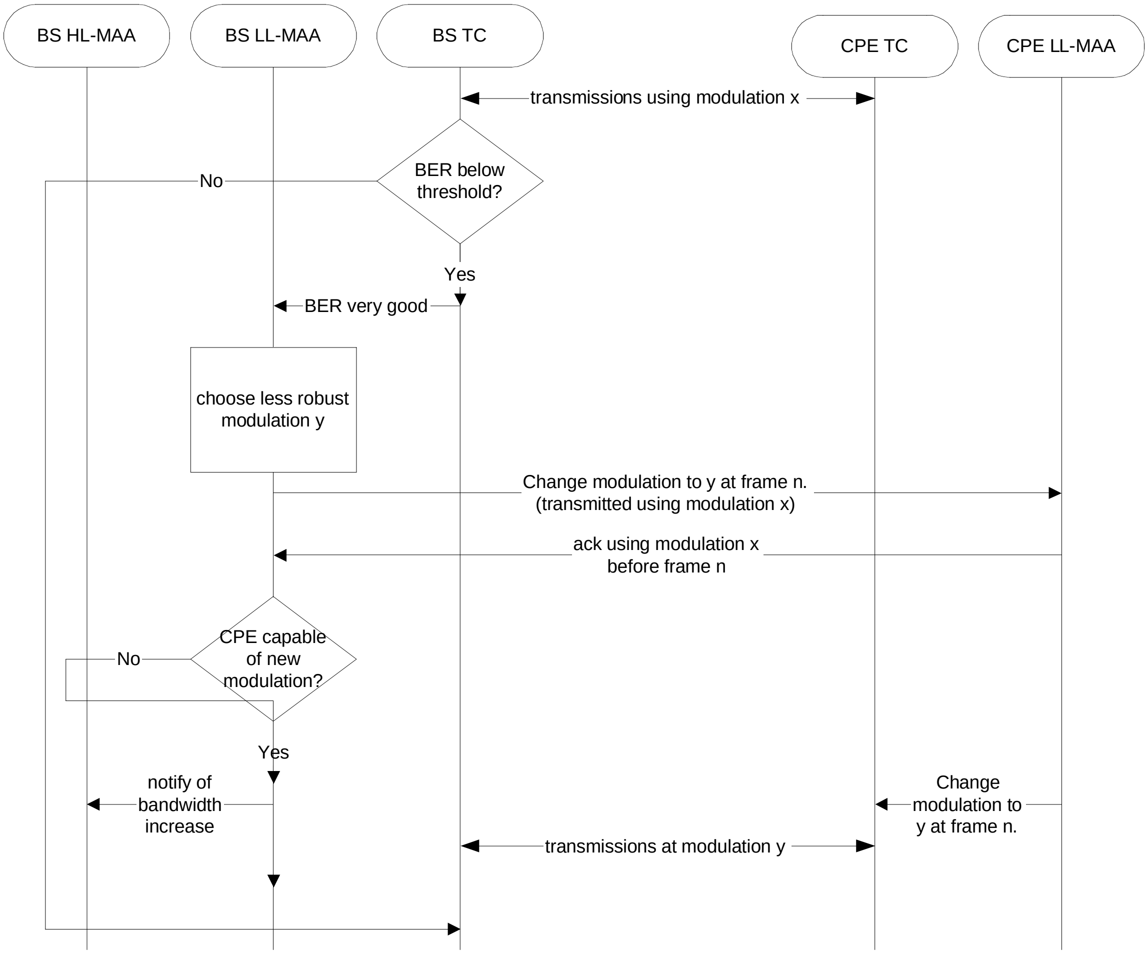
BS HL-MAA
BS LL-MAA
BS TC
CPE TC
CPE LL-MAA
transmissions using modulation x
BER below threshold?
No
Yes
BER very good
choose less robust modulation y
Change modulation to y at frame n. (transmitted using modulation x)
ack using modulation x before frame n
CPE capable of new modulation?
No
Yes
notify of bandwidth increase
Change modulation to y at frame n.
transmissions at modulation y

• Figure 3.4-20: CPE Modulation Change – Less Robust

3.4.7 Multicast Address Assignment

The BS may assign CPEs to multicast groups for two reasons:

- The CPE is involved in a multicast connection.
- The CPEs are being subdivided into groups for such purposes as polling.

In the first case this is indicated when the connection is set up. In the second case, the BS sends a multicast assignment message to the CPE and the CPE responds with a multicast assignment ack. The multicast assignment message can either add or remove a CPE from a multicast group.

Assignment to a multicast group is shown in Figure 3.4-21.

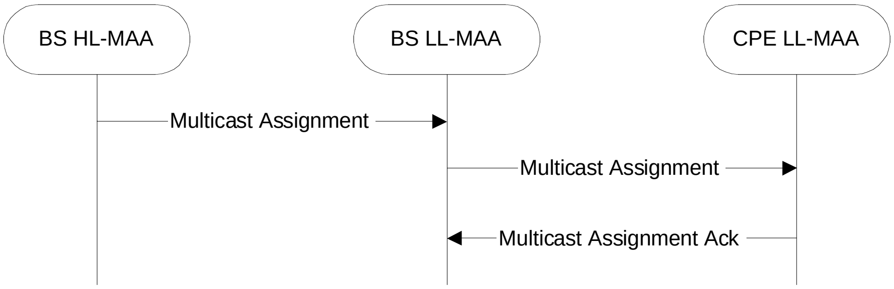

• Figure 3.4-21: Multicast Group Assignment

3.4.8 Privacy

The basic requirements of privacy are to prevent the scenario of a user decoding transmissions of other users and to make the system immune to random or accidental (unsophisticated) eavesdropping. It is understood that the system is not immune against an attack from a sophisticated hacker and the design requirement is purely to give a hard time to the hacker. It is intended that authentication and key exchange will be performed with a much greater degree of security than the privacy protection of the air link.

Functions to support privacy provided by the MAC falls into two categories:

- Capabilities provided by the MAC allowing higher layers to implement network provider policy specific privacy measures.
- Functions performed by the MAC that provide privacy.

The Convergence Subprocesses provide the capability, through the Higher Layer Control Message SAP, for higher layers in the BS and CPE to communicate. This interface may be used by the higher layers to implement network provider specific security measures such as:

- Additional authentication beyond checking a CPE's ID. A suggested authentication process is given in 5.2.1.
- Key exchange. A suggested key exchange process is given in 5.2.2.

The functions performed by the MAC to provide privacy are:

- Acceptance and storage of encryption keys from the higher layers.
- Sequencing through the encryption keys.
- Implementation of encryption.

The MAC accepts keys from the higher layers and stores them for use in the encryption process. The CPE MAC accepts a current key and a next key. The next key becomes the current key at a coordinated time. The BS MAC accepts a current key and a next key for each CPE. Encryption is performed on a per CPE basis, not a per connection basis. All connections to and from an individual CPE are encrypted with the same key. When keys are received from the higher layers, the BS higher layers give the BS MAC a time to start using the new key. When the time arrives, the BS MAC sends a message to the CPE MAC informing it to start using the next key as the current key. Changing from current key to the next key is independent for each CPE. The change of encryption keys is shown in Figure 3.4-22.

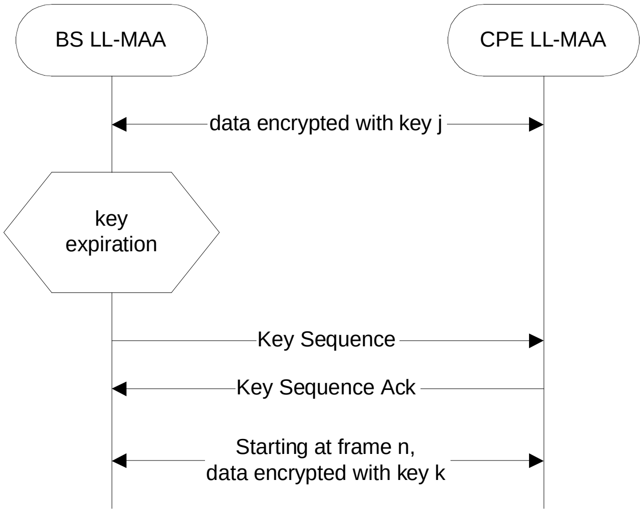

- Figure 3.4-22: Key Sequencing

Once encryption of transmissions to a CPE has begun, all user data transmissions directed to the CPE are encrypted. Encryption is used only on the payload of MAC packets sent to individual CPEs.

Since CPEs do not have a fixed transmission slot as they would in a cell phone system, GSM for instance, the actual implementation of the encryption process is structured to minimize loss of one message causing the CPE and BS encryption to be out of synch for future messages. For instance, if a CPE does not properly receive a PI containing the MAC header of a message directed to it, it does not know to run its encryption engine. Because of this, the encryption process does not retain history beyond the end of an individual MAC payload. The encryption engine restarts from a know point for each encrypted MAC payload, ensuring that loss of a packet does not force the loss of additional packets due to loss of encryption synchronization. For good encryption engines, such as DES, where the encryption is not merely an XOR of a stream of bits, this resetting of the encryption still provides a good measure of security, especially if the keys are changed on a regular basis.

To ensure that known data is never transmitted encrypted, a number of messages or message portions are not encrypted. MAC headers are not encrypted. MAC packets transmitted on broadcast and multicast connections are not encrypted. The Frame Control Header (PHY Control and MAC Control) is not encrypted. MAC basic connections are not encrypted. Preambles to transmissions are not encrypted.

3.5 TC-MAC Interface

The following primitives are sent between the MAC and the TC:

- Tx Timing Error
- Tx Time Advance
- Power
- Power Adjustment
- BER
- Modulation
- Encryption Keys
- Encryption Key Sequencing

3.5.1 Tx Timing Error and Timing Advance

In the base station, the timing error of each CPE is measured at registration and is monitored on a regular basis. The BS TC reports the timing error of the CPEs to the BS LL-MAA. During CPE registration, a new Tx timing advance is sent by the BS LL-MAA to the CPE LL-MAA in the Registration Results message. At other times, the Tx timing advance is sent by the BS LL-MAA to the CPE LL-MAA via the Tx Advance Change message. In this case, the CPE responds with a Tx Advance Ack message after passing the Tx timing advance to the CPE TC. The message sequence for the Tx Advance Change message is shown in Figure 3.5-1.

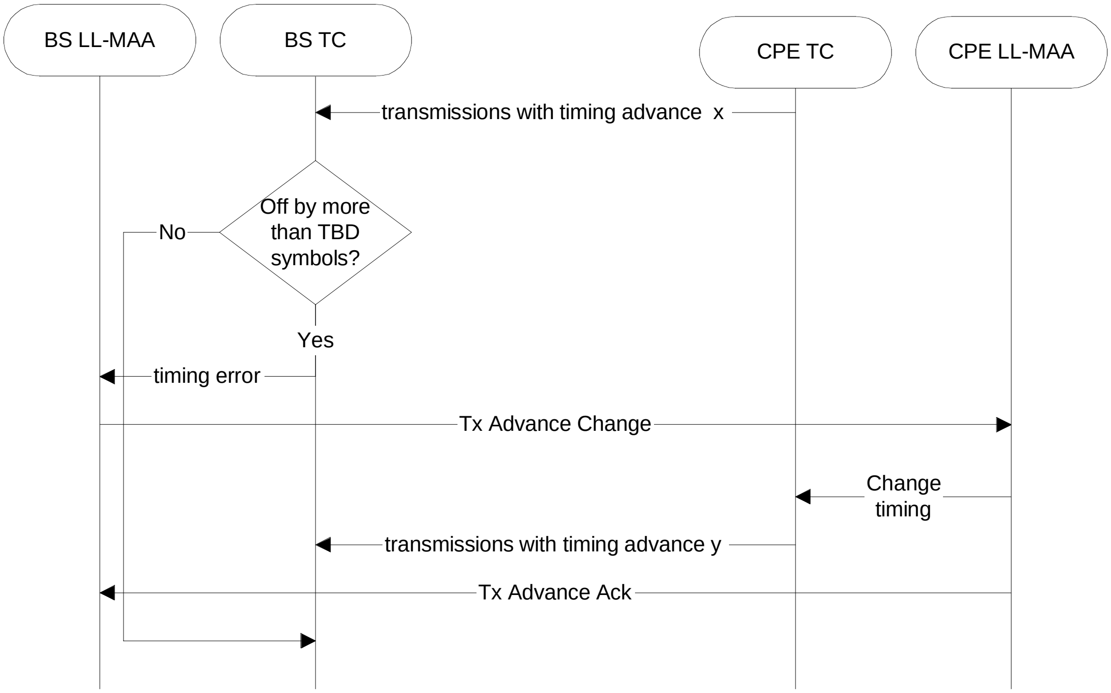
BS LL-MAA
BS TC
CPE TC
CPE LL-MAA
transmissions with timing advance x
Off by more than TBD symbols?
No
Yes
timing error
Tx Advance Change
Change timing
transmissions with timing advance y
Tx Advance Ack

• Figure 3.5-1: Transmit Timing Advance Change

3.5.2 Power and Power Offset

In the base station, the power of each CPE is monitored on a regular basis. The BS TC reports the power of a CPE to the BS LL-MAA. During CPE registration, a new power adjustment is sent by the BS LL-MAA to the CPE LL-MAA in the Registration Results message. At other times, if the power adjustment is minor, the power control (PC) bits in the MAC header are used to signal a CPE to adjust power. If the power change required is more substantial, the power adjustment is sent by the BS LL-MAA to the CPE LL-MAA via the Power Adjustment message. In this case, the CPE responds with a Power Adjustment Ack message after passing the power adjustment to the CPE TC.

When the PC bits are used, the BS may set them in every MAC header sent to that CPE that frame. This minimizes the possibility of loss of the adjustment.

Power adjustment is shown in Figure 3.5-2.

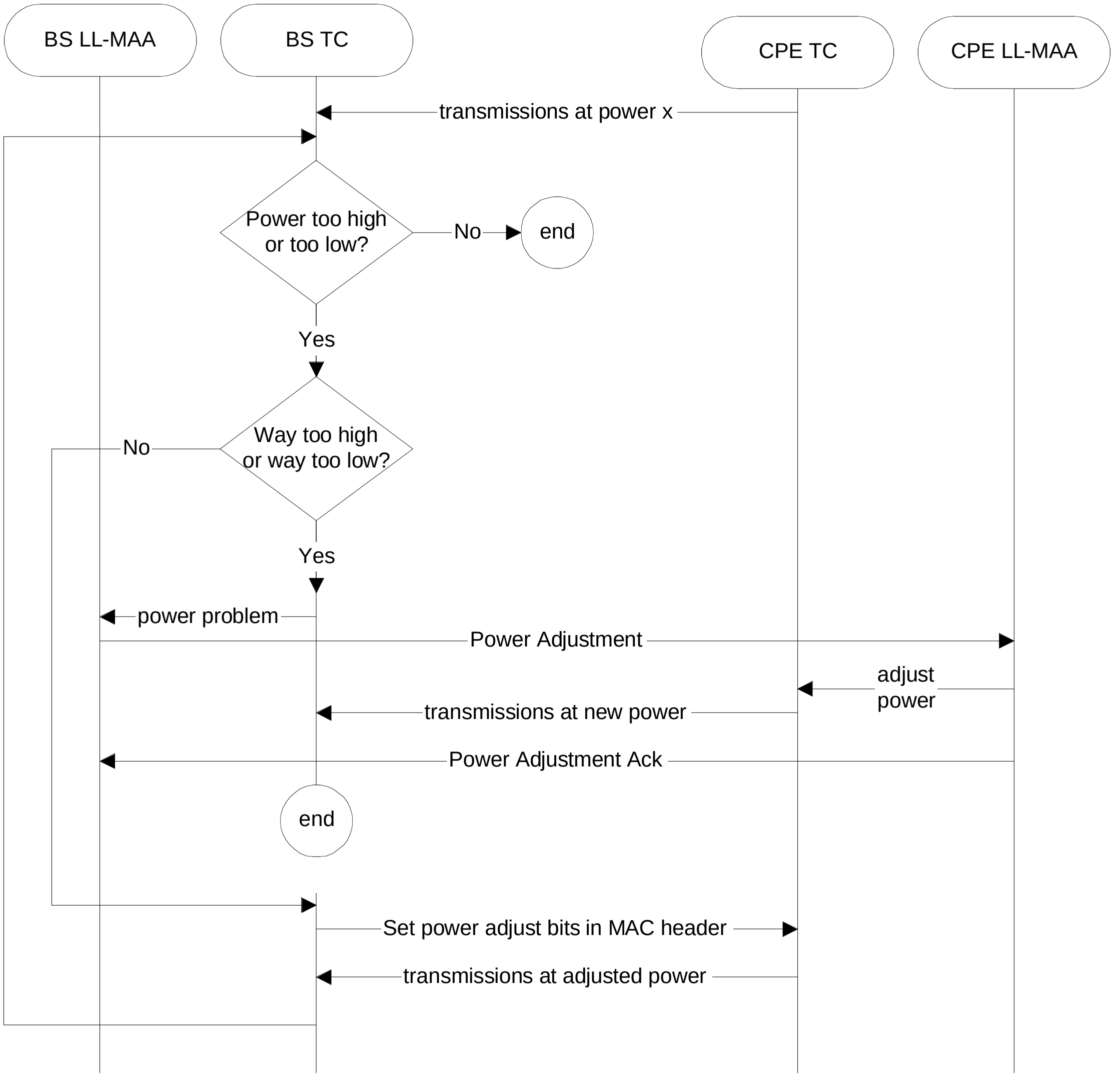
BS LL-MAA
BS TC
CPE TC
CPE LL-MAA
transmissions at power x
Power too high or too low?
No
end
Yes
Way too high or way too low?
No
Yes
power problem
Power Adjustment
adjust power
transmissions at new power
Power Adjustment Ack
end
Set power adjust bits in MAC header
transmissions at adjusted power

- Figure 3.5-2: Power Adjustment

3.5.3 BER and Modulation

In the base station, the BER of each CPE is monitored on a regular basis. The BS TC reports the BER and SNR of a CPE to the BS LL-MAA. The BS MAA sends a message to the CPE MAA requiring it to change modulation at a certain frame. The BS LL-MAA passes this change information and time to the BS TC. The CPE LL-MAA passes it to the CPE TC.

3.5.4 Encryption Keys and Encryption Key Sequencing

The LL-MAA in both the BS and the CPE passes encryption keys to the TC. The CPE has one set of encryption keys. The BS has one set per CPE. The LL-MAA in both the BS and the CPE tells the TC when to switch to the next key in the sequence.

3.6 MAC Messages

MAC and TC messages have the form described in section 3.2.2.7.3. Note that multiple MAC control messages to or from the same CPE may be packed into the same MAC packet.

Table 4 lists the MAC and TC messages.

- Table 4: MAC & TC Over the Air Messages

Message	**Direction**	**Purpose**	**Section**
PHY Control	BS -> CPE	Physical Layer Control	3.6.1
MAC Control	BS -> CPE	MAC Control	3.6.2
Registration	BS <- CPE	Registration	3.6.3.1
Registration Results	BS -> CPE	Registration	3.6.3.2
Re-register	BS -> CPE	Registration	3.6.3.3
Registration Collision	BS -> CPE	Registration	3.6.3.4
Change Modulation	BS -> CPE	Physical Layer Maintenance	3.6.4.1
Modulation Change Ack	BS <- CPE	Physical Layer Maintenance	3.6.4.2
Tx Advance Change	BS -> CPE	Physical Layer Maintenance	3.6.4.3
Tx Advance Ack	BS <- CPE	Physical Layer Maintenance	3.6.4.4
Power Adjustment	BS -> CPE	Physical Layer Maintenance	3.6.4.5
Power Adjustment Ack	BS <- CPE	Physical Layer Maintenance	3.6.4.6
Bandwidth Request	BS <- CPE	Connection Maintenance	3.6.5.1
Multicast Assignment	BS -> CPE	Connection Maintenance	3.6.5.2
Multicast Assignment Ack	BS <- CPE	Connection Maintenance	3.6.5.3
Key Sequence	BS -> CPE	Security	3.6.6.1
Key Sequence Ack	BS <- CPE	Security	3.6.6.2
Channel Change	BS -> CPE	Load Leveling	3.6.7.1
Channel Change Ack	BS <- CPE	Load Leveling	3.6.7.2

Table 5 gives the Message Type IDs for the various messages.

• Table 5: Message Type IDs

Message	ID
PHY Control	0x00
MAC Control	0x01
Registration	0x10
Registration Results	0x11
Re-register	0x14
Registration Collision	0x15
Change Modulation	0x21
Modulation Change Ack	0x22
Tx Advance Change	0x23
Tx Advance Ack	0x24
Power Adjustment	0x25
Power Adjustment Ack	0x26
Bandwidth Request	0xB0
Multicast Assignment	0xC0
Multicast Assignment Ack	0xC1
Key Sequence	0x40
Key Sequence Ack	0x41
Channel Change	0x50
Channel Change Ack	0x51

3.6.1 Physical Layer Control

The physical layer control information is always in the first message of the downlink subframe frame, following the 1 PS downlink subframe preamble. This message is sent using connection ID 0xFFFE. Table 6 shows the format of the physical layer control information.

• Table 6: PHY Control

Field	Size	Comments
Message Type ID	1 byte	Value = 0x00
Hyperframe Number	15 bits	Zero-based
Multiframe Number within Hyperframe	5 bits	Zero-based
Frame Number within Multiframe	4 bits	Zero-based
Maximum Tx Timing Advance	1 byte	Cell size expressed in PS
PHY Type	1 byte	0 = non-adaptive TDD 1 = TDD 2 = FDD
DL Start PS for QAM-16	2 bytes	0 = no QAM-16 this frame
DL Start PS for QAM-64	2 bytes	0 = no QAM-64 this frame
DL end PS	2 bytes	0 = no uplink this frame
Base Station ID	6 bytes	
Total	**18 bytes**	

3.6.2 MAC Control

The MAC Control message is broadcast to all CPEs every frame using connection ID 0xFFFE. It is concatenated with the PHY Control message, without it's own MAC header, as provide for in section 3.2.2.7.3. It provides the BS MAC version being used and an indication of whether this physical channel is open to CPE registration. It contains the Uplink Subframe Map for the next frame.

- Table 7: MAC Control

Field	**Size**	**Comments**
Message Type ID	1 byte	Value = 0x01
MAC version number	2 bytes	
Registration not allowed flag	1 bit	0 = registration on this channel OK 1 = no registration on this channel
Reserved	7 bits	
Reserved	6 bits	
Number of Uplink Subframe Map entries	10 bits	
Uplink Subframe Map	4 bytes per entry	14 bits = truncated CPE basic connection ID 11 bits = start PS for CPE 5 bits = start symbol within a PS for CPE 2 bits reserved
Total	**6+(4*entries) bytes**	

3.6.3 Registration

3.6.3.1 Registration Message

The registration message is sent by CPEs in the registration contention slots when performing registration. It is sent on connection ID 0x0000, which is reserved for registering CPEs.

- Table 8: Registration Message

Field	**Size**	**Comments**
Message Type ID	1 byte	Value = 0x10
CPE ID	6 bytes	48 bit CPE ID
Reserved	6 bits	
Physical Slot Sent	10 bits	PS the message was transmitted in, relative to the start of the frame.
Total	**9 bytes**	

3.6.3.2 Registration Results Message

The registration results message is sent in response to the Registration Message on connection ID 0x0000, which is reserved for registering CPEs. It is addressed to the CPE via its 48 bit CPE ID. If the result field is not a rejection (cause field = 0x03 or 0x04), the CPE must apply the power and timing adjustments proceed with ranging or move to the specified channel.

- Table 9: Registration Results Message

Field	**Size**	**Comments**
Message Type ID	1 byte	Value = 0x11
Result	1 byte	0x01 = continue ranging 0x02 = change channel and continue 0x03 = invalid CPE ID 0x04 = service not authorized
CPE ID	6 bytes	48 bit CPE ID
Basic Connection ID	2 bytes	Connection ID used between BS and CPE MACs
Control Connection ID	2 bytes	Connection ID used between BS and CPE higher layers
New Channel	2 bytes	Valid if Result = 0x02
Tx Timing Advance	2 bytes	In _ symbol units
Power Adjustment	1 byte	Signed, relative power adjustment
Minimum Non-terminal Fragment Size	1 byte	In bytes
Non-terminal Fragment Step Size	1 byte	In bytes
Total	**19 bytes**	

3.6.3.3 Re-register Message

The Re-register message is sent to a CPE that needs to re-register. It is sent on the CPE's basic connection ID.

- Table 10: Re-register Message

Field	**Size**	**Comments**
Message Type ID	1 byte	Value = 0x14
Cause	1 byte	Values **TBD**
Total	**2 bytes**	

3.6.3.4 Registration Collision

The Registration Collision message is sent by the Base Station when it detects a collision in the registration contention slots. It is sent on the registration connection ID 0x0000. This message is **optional** and is only sent by base stations capable of detecting collisions.

- Table 11: Registration Collision

Field	**Size**	**Comments**
Message Type ID	1 byte	Value = 0x15
Frame in which collision was detected	3 bytes	15 bits – hyperframe modulo 128 5 bits – multiframe within hyperframe 4 bits – frame within multiframe
Total	**4 bytes**	

3.6.4 Physical Layer Maintenance

3.6.4.1 Change Modulation Message

The Change Modulation message is sent by the BS to the CPE on the CPE's basic connection ID. If it is a command to go to a more robust (lower bits per symbol) modulation, the message is sent QAM-4 to increase the likelihood of reception by the CPE. Otherwise, it is sent at the current (before the change) operational modulation for the CPE. This message must be transmitted by the BS at least 10 frames before the modulation change is to take place.

- Table 12: Change Modulation Message

Field	**Size**	**Comments**
Message Type ID	1 byte	Value = 0x21
Modulation	1 byte	4 = QAM-4 16 = QAM-16 64 = QAM-64
Frame at which to change modulation	3 bytes	15 bits – hyperframe modulo 128 5 bits – multiframe within hyperframe 4 bits – frame within multiframe
Total	**5 bytes**	

3.6.4.2 Modulation Change Ack

The CPE responds to the Change Modulation message with a Modulation Change Ack. If the Change Modulation message indicated a change to a less robust modulation, the Modulation Change Ack may be sent before the frame indicated in the Change Modulation Message using the CPE's old modulation. Otherwise, it is sent the frame indicated at the new, more robust, modulation. Errors are always indicated before the frame indicated, using the CPE's current modulation.

- Table 13: Modulation Change Ack

Field	**Size**	**Comments**
Message Type ID	1 byte	0x22
Ack Code	1 byte	0x00 = OK 0x01 = error 0x02 = CPE not capable of indicated modulation
Total	**2 bytes**	

3.6.4.3 Tx Advance Change

The Tx Advance Change message is sent by the BS to the CPE to adjust the CPE's timing advance. It is sent on the CPE's basic connection ID.

- Table 14 : Tx Advance Change Message

Field	Size	Comments
Message Type ID	1 byte	Value = 0x23
Tx Timing Advance Adjustment	1 byte	Signed, in _ symbol units
Total	**2 bytes**	

3.6.4.4 Tx Advance Ack

The Tx Advance Ack message is sent by a CPE on its basic connection ID in response to the Tx Advance Change message.

- Table 15: Tx Advance Ack

Field	Size	Comments
Message Type ID	1 byte	Value = 0x24
Ack Code	1 byte	0x00 = OK 0x01 = cumulative Tx Advance would be negative 0x02 = cumulative Tx Advance would be > max Tx advance in PHY Control message
Total	**2 bytes**	

3.6.4.5 Power Adjustment

The Power Adjustment message is sent by the BS to the CPE to adjust the CPE’s power level. It is sent on the CPE’s basic connection ID.

- Table 16: Power Adjustment Message

Field	Size	Comments
Message Type ID	1 byte	Value = 0x25
Power Adjustment	1 byte	Signed, relative value
Total	**2 bytes**	

3.6.4.6 Power Adjustment Ack

The Power Adjustment Ack is sent by the CPE in response to the Power Adjustment message. It is sent on the CPE’s basic connection ID.

- Table 17: Power Adjustment Ack

Field	Size	Comments
Message Type ID	1 byte	Value = 0x26
Ack Code	1 byte	0x00 = OK 0x01 = at max power 0x02 = at min power
Total	**2 bytes**	

3.6.5 Connection Maintenance

3.6.5.1 Bandwidth Request

The Bandwidth Request message is sent by the CPE to the BS to request bandwidth in which to send data for a specific connection. The message is sent on the basic connection ID of the CPE.

- Table 18: Bandwidth Request

Field	**Size**	**Comments**
Message Type ID	1 byte	Value = 0xB0
Connection ID	2 bytes	
Amount Requested	2 bytes	In bytes per frame for CG connections and total bytes pending for DAMA connections
Total	**5 bytes**	

3.6.5.2 Multicast Assignment

The Multicast Assignment message is sent to a CPE to include it in a multicast polling group. Multicast data connections are set up using the Connection Establishment message. This message is normally sent on a CPE's basic connection ID. It may also be sent to a group of CPE's on a previously set up multicast connection ID.

- Table 19: Multicast Assignment Message

Field	**Size**	**Comments**
Message Type ID	1 byte	Value = 0xC0
Join/Leave	1 byte	0x01 = join multicast group 0x00 = leave multicast group
Multicast Connection ID	2 bytes	
Total	**4 bytes**	

3.6.5.3 Multicast Assignment Ack

The Multicast Assignment Ack is sent by the CPE in response to the Multicast Assignment message. It is sent on the basic connection of the CPE.

- Table 20: Multicast Assignment Ack

Field	**Size**	**Comments**
Message Type ID	1 byte	Value = 0xC1
Join/Leave	1 byte	0x01 = join multicast group 0x00 = leave multicast group
Multicast Connection ID	2 bytes	
Ack Code	1 byte	0x00 = OK **TBD** = error
Total	**5 bytes**	

3.6.6 Security

3.6.6.1 Key Sequence

The key Sequence Message is sent by the BS to the CPE notifying it of an upcoming change to the next key in the CPE's key sequence. This message is also used to inform the CPE at which frame to start encryption after the CPE registers. This message is sent on the CPE's basic connection ID. This message must be transmitted by the BS at least 10 frames before the key change is to take place.

- Table 21: Key Sequence Message

Field	**Size**	**Comments**
Message Type ID	1 byte	Value = 0x40
Key sequence number	1 byte	Range [0,7], 0xFF = stop encrypting
Frame in which to change keys	3 bytes	15 bits – hyperframe modulo 128 5 bits – multiframe within hyperframe 4 bits – frame within multiframe
Total	**5 bytes**	

3.6.6.2 Key Sequence Ack

The Key Sequence Ack is sent by the CPE in response to the Key Sequence message. It is sent on the CPE's basic connection ID.

- Table 22: Key Sequence Ack

Field	**Size**	**Comments**
Message Type ID	1 byte	Value = 0x41
Key sequence number	1 byte	Range [0,7]
Ack Code	1 byte	0x00 = OK **TBD** = error
Total	**3 bytes**	

3.6.7 Load Leveling

3.6.7.1 Channel Change

The Channel Change Message is sent by the BS to the CPE to direct it to change physical channels at a given frame. The new channel will be time synchronized with the old channel and will use the same reference frequency, so there is no re-synchronization necessary. TDM data will not be lost or delayed in the change over. Other data may be delayed at the CPE or lost at the base station. The message is sent on the CPE's basic connection ID. This message must be transmitted by the BS at least 10 frames before the channel change is to take place.

- Table 23: Channel Change Message

Field	**Size**	**Comments**
Message Type ID	1 byte	Value = 0x50
Allocation that frame	3 bytes	10 bits – PS at which CPE UL starts 4 bits - reserved 10 bits – number of PS allocated to CPE
New Channel	2 bytes	
Frame at which to change	3 bytes	15 bits – hyperframe modulo 128 5 bits – multiframe within hyperframe 4 bits – frame within multiframe
Total	**11 bytes**	

3.6.7.2 Channel Change Ack

The Channel Change Ack is sent by the CPE in response to the Channel Change Message. It is sent on the CPE's basic connection ID.

- Table 24: Channel Change Ack

Field	**Size**	**Comments**
Message Type ID	1 byte	Value = 0x51
Ack Code	1 byte	0x00 = OK **TBD** = error
Total	**2 bytes**	

Chapter

4

4 Convergence Sublayers

The following sections describe the convergence sublayers.

4.1 TDM Compression

Efficient support of TDM traffic is maintained by transmitting only the active DS0 channels, in cases where the trunk is channelized. If the trunk is not channelized, for example if it is used for data traffic, then no compression is possible.

Interpretation of the signaling is the responsibility of the user station. (This is desirable because there are many variants of the signaling protocols in use. Rather than require the base station to be conversant with all of them, each user station is responsible for monitoring the particular variant in use by the attached equipment on the user's site.) The user station modifies its bandwidth requests as it detects channels going off-hook or on-hook. Grants of transmission opportunities are on a continuing basis for TDM: they are valid until changed, with no need for new requests, though the grants still appear in each allocation map.

For example - a T1, in a one-millisecond packetization interval up to 192 bytes of data are carried plus 3 bytes of Time Slot Map. The Time Slot Map indicates whether a channel is active or suppressed.

Formats for TDM packets are given in Figure 4.1-1.

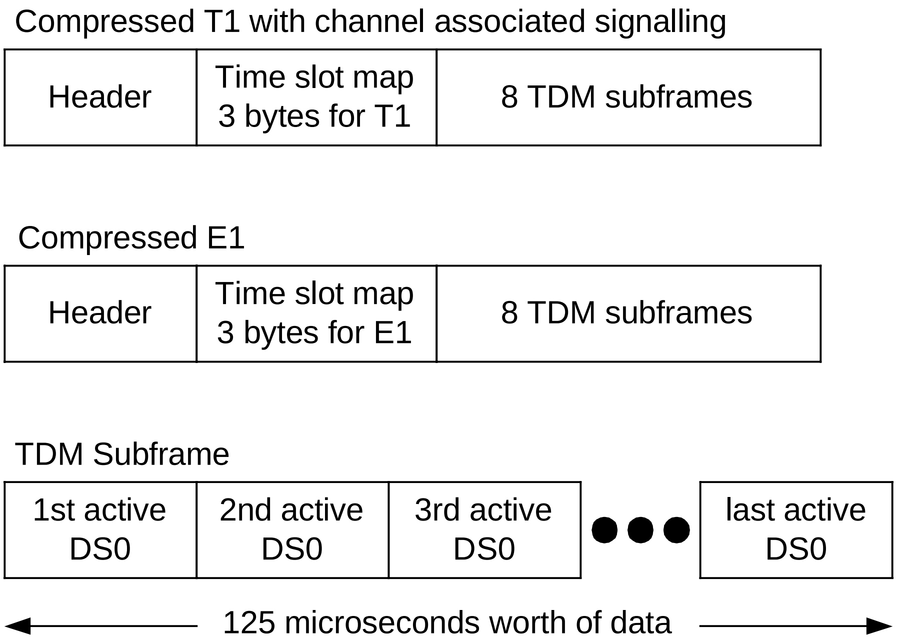
Compressed T1 with channel associated signalling
Header
Time slot map 3 bytes for T1
8 TDM subframes
Compressed E1
Header
Time slot map 3 bytes for E1
8 TDM subframes
TDM Subframe
1st active DS0
2nd active DS0
3rd active DS0
●●●
last active DS0
125 microseconds worth of data

- Figure 4.1-1: TDM Compression

Both ends of the air link monitor the state of the Robbed Bits Signalling (RBTS). Changes in the RBTS state are sent using and out of band control channel.

Since the BWA frame structure imposes a 1-millisecond granularity on the traffic from any individual user station, the base station must be prepared to provide a de-jitter buffer that will permit receiving a millisecond's worth of TDM data from each station and feed it out into the backbone network with conventional timing (normally one byte per DS0 at 125-microsecond intervals).

4.2 ATM Compression

ATM compression is accomplished by sending the first header in a series of identical headers, along with a run-length count. If ATM header compression is performed it shall be done within frame boundaries (e.g. 1 millisecond). The most-significant bit of the 8-bit count field is used to indicate that the last cell of this compression group is the last cell of the packet, i.e., that the AAL 5 last-cell bit is set. Hence there is no need to send this header as a separate compression group even though it differs (in this bit) from the previous header. Header compression is shown schematically in Figure 4.2-1.

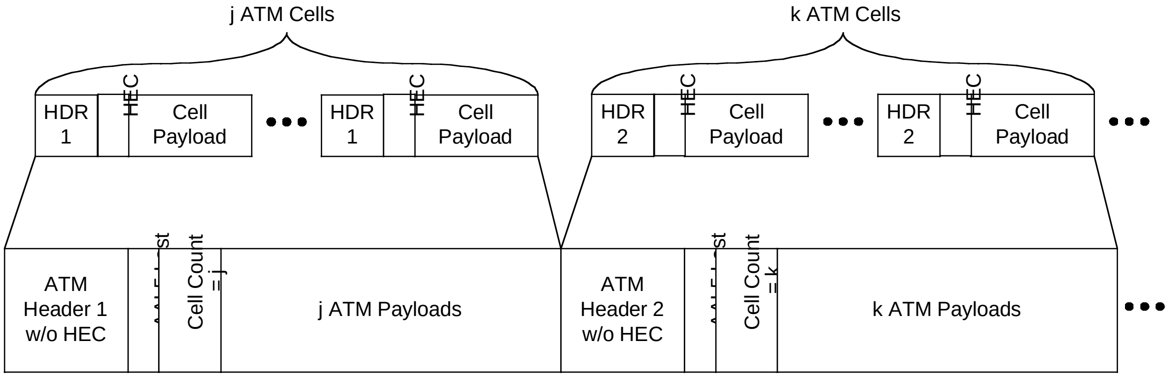

- Figure 4.2-1: Compression of Identical ATM headers

Chapter

5

5 Policies

The following sections describe policies within the BS and CPE that, while not requirements of the MAC, ensure that the MAC is used with its intended efficiencies.

5.1 Polling Groups

Polling groups are established to minimize the amount of overhead involved in polling the users, while still providing fast access. Users may be considered members of one of 4 groups:

Active	with requests outstanding or made within the last two frames
Recently active	activity with the last few seconds
Pausing	no activity with the last few seconds
Inactive	no activity with the last few minutes

The boundaries between the different groups may be defined by the service operator as part of the system configuration. Likewise, the frequencies of polling can be set as well. The following polling frequencies are the defaults:

Active	no polling since CPE can piggyback or use poll-me bit
Recently active	individual poll every **TBD** milliseconds
Pausing	individual poll every **TBD** milliseconds
Inactive	don’t poll individually, have user make contention request

5.2 Security

5.2.1 Authentication

For the authentication process the following is assumed:

- Each CPE has an ID which has a public part PBID and a private part PVID.
- PBID may be distributed freely over a non secured link (i.e., at registration); The PBID is the 48 bit CPE ID.
- PVID remains private and is known only to the CPE and the NMS. It is never transmitted over the air. PVID size is determined mainly by privacy strength requirements (the longer the better) and memory limitations (as there is some limit to memory size)

- Each CPE has an additional seed number SU that is initialized to some known value at power on reset (or registration).

- SU may be updated by a procedure independent of the authentication process (See section 5.2.2).

- SU is known also to the NMS (per CPE).

- It is assumed that both the NMS and any CPE are capable of performing a predetermined one-way hash function, X=H(PBID,PVID,SU,R) where H() is the hash function and X and R are integer numbers

The authentication process is simple. In response to an authentication cycle (i.e., registration) the NMS generates a random number R and uses it to calculate X using PBID, PVID, SU of the specific CPE and the hashing function. R is transmitted to the user. The user calculates X' based on the received R and its own parameters PBID, PVID and SU. It transmits X' back to the NMS. The NMS compares X' to X. If they are the same, the authentication cycle successfully completes.

Parameter size issues:

Even if the hash function seems to be irreversible it is possible to guess the parameters if their size (bit length) is small. For example assume that all parameters are known besides PVID. There are only 2^n combinations to go through where n is the PVID size. A small n may make the authentication algorithm vulnerable.

5.2.2 Key Exchange

The key exchange algorithm is a two part process. It involves the exchange of the seed SU and the generation of a new key, using the new SU.

5.2.2.1 SU Update

If the SU parameter can be changed upon request, the authentication and privacy strength of the system is increased. The tricky part is to update the SU parameter “secretly”.

Assume that there is another fixed and secret parameter K (a key), unique for each CPE - known only to the CPE and the NMS. When the system requires a SU update the following will occur:

- NMS generates a random number R. Using a one-way hash function G, a new SU (noted as NSU) is generated by NSU=G(R,SU,K). All CPEs know the one-way hash function as well (the algorithm).

- The NMS orders the CPE to go into SU-update mode, and transmits to the CPE the random number R as a parameter.

- The CPE calculates a new SU based on the received R and its current SU and K.

- The CPE performs a reverse authentication test: It generates a random number P and uses the authentication one-way hash function X=H(PBID,PVID,NSU,P) to challenge the NMS in a similar way that the authentication phase is performed (but in a reverse manner – it is the NMS which is challenged).

- If the NMS response back X' to the CPE is equal to X then an SU update confirmation message is sent to the NMS and the process successfully terminates.

5.2.2.2 Key Generation

In the case of either an authentication event or SU update event the encryption key is generated using a one-way hash function F(PBID,PVID,SU,R) where R is the random number used at the challenge process. Note that F is different than H and G and should not generate the challenge sequence which is assumed to be known to the public.

5.2.3 Remark on One-way Hash Functions

Although these functions are the core of the authentication and privacy scheme this does not mean that they have to be standard. To the best of my knowledge there is no direct "NSA restriction" on such functions although they must be disclosed as part of the approval process to the NSA.

The bottom line is that the functions should be chosen according to design implementation issues (i.e., hardware vs. software based). There are various types of one-way hash functions – some may be candidates or we can invent our own.

Chapter
6

6 Functional Requirements and Criteria Table

(1) Meets System Requirements:
Meets all system requirements; Supports IP/ATM and STM; Support of different QoS constraints

(2) Mean Access Delays & Variance:
Different allocation requests mechanisms gurantee and control latency sensitive services (collision based, piggy backing etc.). This proposal can support easily a bounded delay service as it has access to various QoS parameters.

(3) Payload & BW efficiency
This proposal is optimized for BWA with minimum overheads. Address translation minimizes pointer sizes. This proposal does not need to support anything besides wireless PMP and therefore has no additional header burden.

(4) Simplicity of implementation
The implementation of this MAC can be done with the current available level of processing power. As it is partitioned to a lower and higher level functionality, the lower level functionality which tends to be a real time process can be implemented easily in an ASIC.

(5) Scalibility
The MAC was designed to deal with channels carrying more than 100 Mbps of traffic per upstream channel hence can easily handle any realistic traffic scenario.

(6) Service Support Flexibility
The MAC supports various types of services as IP, ATM and STM. It is optimized for the business environment. It can support any new service optimally as it is independent of its higher layer choice (not optimized specifically only for Residential IP or only for ATM)

(7) Robustness
The MAC is very robust and it can easily recover from burst error scenarios. For example if a user data is hit on one of its packets within a frame, the MAC can re-synchornize on the next PDU immediately without losing the whole burst

(8) Security
The MAC can interact with any basic public/private key encryption system

(9) Maturity
This MAC is new as it is optimized for BWA and does rely on other existing technologies. The MAC is currently being deployed for BWA field trials using millimeter wave frequencies.

(10) Sign-on Process
As most of the other processes within the system, the sign-on process is fully automatic. The user parameters are tracked through operation and changes are made on the fly (i.e., modulation change, RF carrier change).

(11) Adequacy of mangment functions
See (10).

(12) Convergence with existing technologies
See (6).

(13) Ability to work with PHY variations

The MAC can work with any duplex scheme. And can be tailored to work with any proper PHY which has minimal latency impact (i.e., minimum or no interleaving) as any MAC would require for BWA. It supports real time subscriber base adaptive modulation.

This proposal either fulfills mandatory requirements or does not preclude items which were pointed out as part of mandatory requirements.