IEEE P802.22
Wireless RANs

A PHY/MAC Proposal for IEEE 802.22 WRAN Systems
Part 2: The Cognitive MAC

Date: **2006-02-23**

Author(s):

Name	Company	Address	Phone	email
John Benko	France Telecom (FT)	USA		John.Benko@francetelecom.com
Yoon Chae Cheong	SAIT	Korea	+82-31-280-9501	Yc.cheong@samsung.com
Carlos Cordeiro	Philips	USA	+1 914 945-6091	Carlos.Cordeiro@philips.com
Wen Gao	Thomson Inc.	USA	+1-609-987-7308	wen.gao@thomson.net
Chang-Joo Kim	ETRI	Korea	+82-42-860-1230	cjkim@etri.re.kr
Hak-Sun Kim	Samsung Electro-Mechanics	Korea	+82-31-210-3500	hszic.kim@samsung.com
Stephen Kuffner	Motorola	USA	+1-847-538-4158	stephen.kuffner@motorola.com
Joy Laskar	Georgia Institute of Technology	USA	+1-404-894-5268	joy.laskar@ece.gatech.edu
Ying-Chang Liang	Institute for Infocomm Research (I2R)	Singapore	+65-68748225	ycliang@i2r.a-star.edu.sg

Abstract

We propose a MAC layer to be used as the basis for the future IEEE 802.22 WRAN standard operating in the TV bands. The proposed MAC layer is in some respects inspired by the IEEE 802.16 standard, but it provides major extensions, improvements and also simplifications in order to meet the 802.22 functional requirements.

Co-Author(s):				
Name	**Company**	**Address**	**Phone**	**email**
Myung-Sun Song	ETRI	Korea	+82-42-860-5046	mssong@etri.re.kr
Soon-Ik Jeon	ETRI	Korea	+82-42-860-5947	sijeon@etri.re.kr
Gwang-Zeen Ko	ETRI	Korea	+82-42-860-4862	gogogo@etri.re.kr
Sung-Hyun Hwang	ETRI	Korea	+82-42-860-1133	shwang@etri.re.kr
Bub-Joo Kang	ETRI	Korea	+82-42-860-5446	kbj64370@etri.re.kr
Chung Gu Kang	ETRI	Korea	+82-2-3290-3236	ccgkang@korea.ac.kr
KyungHi Chang	ETRI	Korea	+82-32-860-8422	khchang@inha.ac.kr
Yun Hee Kim	ETRI	Korea	+82-31-201-3793	yheekim@khu.ac.kr
Moon Ho Lee	ETRI	Korea	+82-63-270-2463	moonho@chonbuk.ac.kr
HyungRae Park	ETRI	Korea	+82-2-300-0143	hrpark@mail.hangkong.ac.kr
Martial Bellec	France Telecom	France	+33 2 99 12 48 06	Martial.Bellec@francetelecom.com
Denis Callonnec	France Telecom	France	+33-4-76-764412	Denis.Callonnec@francetelecom.com
Luis Escobar	France Telecom	France	+33-2-45-294622	Luis.Escobar@francetelecom.com
Francois Marx	France Telecom	France	+33-4-76-764109	Francois.Marx@francetelecom.com
Patrick Pirat	France Telecom	France	+33-2-99-124806	Ppirat.ext@francetelecom.com
Kyutae Lim	Georgia Institute of Technology	USA	+1-404-385-6008	ktlim@ece.gatech.edu
Youngsik Hur	Georgia Institute of Technology	USA	+1-404-385-6008	yshur @ece.gatech.edu
Chee Wei Ang	I2R	Singapore	+65-68748225	angcw@i2r.a-star.edu.sg
Anh Tuan Hoang	I2R	Singapore	+65-68748225	athoang@i2r.a-star.edu.sg
Peng-Yong Kong	I2R	Singapore	+65-68748225	kongpy@i2r.a-star.edu.sg
Haiguang Wang	I2R	Singapore	+65-68748225	
Yufei Blankenship	Motorola	USA		Yufei.Blankenship@motorola.com
Brian Classon	Motorola	USA		Brian.Classon@motorola.com
Fred Vook	Motorola	USA		Fred.Vook@motorola.com
Jeff Zhuang	Motorola	USA		Jeff.Zhuang@motorola.com
Kevin Baum	Motorola	USA		Kevin.Baum@motorola.com
Tim Thomas	Motorola	USA		Tim.Thomas@motorola.com
David Grandblaise	Motorola	France	+33 1 69 35 25 82	David.Grandblaise@motorola.com
Dagnachew Birru	Philips	USA	+1-914-945-6401	Dagnachew.Birru@philips.com
Kiran Challapali	Philips	USA	+1-914 945-6356	Kiran.challapali@philips.com
Vasanth Gaddam	Philips	USA	+1-914-945-6424	Vasanth.Gaddam@philips.com
Monisha Ghosh	Philips	USA	+1-914-945-6415	Monisha.Ghosh@philips.com
Duckdong Hwang	SAIT	Korea	+82-31-280-9513	duckdong.hwang@samsung.com
Chung Jaehak	SAIT	Korea	+82-32-860-8421	jchung@inha.ac.kr
Kim Jaemyeong	SAIT	Korea	+82-32-860-8420	jaekim@inha.ac.kr
Ashish Pandharipande	SAIT	Korea	+82-010-6335-7784	pashish@ieee.org
Yoo Sangjo	SAIT	Korea	+82-32-860-8304	sjyoo@inha.ac.kr
Jeong Suk Lee	Samsung Electro-Mechanics	Korea	+82-31-210-3217	js0305.lee@samsung.com
Chang Ho Lee	Samsung Electro-Mechanics	Korea	+82-31-210-3217	changholee@samsung.com
Wangmyong Woo	Samsung Electro-Mechanics	Korea	+82-31-210-3217	wmwoo@samsung.com
David Mazzarese	Samsung Electronics Co. Ltd.	Korea	+82 10 3279 5210	d.mazzarese@samsung.com
Baowei Ji	Samsung Telecom America	USA	+1-972-761-7167	Baowei.ji@samsung.com
Max Muterspaugh	Thomson Inc.	USA	+1-317-587-3711	Max.muterspaugh@thomson.net
Hang Liu	Thomson Inc.	USA	+1-609-987-7335	hang.liu@thomson.net
Paul Knutson	Thomson Inc.	USA	+1-609-987-7314	paul.knutson@thomson.net

Josh Koslov	Thomson Inc.	USA	+1-609-987-7337	josh.koslov@thomson.net

Contents

List of Figures

List of Tables

1. Introduction

This document describes the Cognitive MAC (CMAC) layer proposed to be used as the IEEE 802.22 WRAN [2] PMP medium access control standard. As it shall be presented, not only CMAC meets all the functional requirements set forth by the IEEE 802.22 WG [3], but it is built upon a design cornerstone which can be described in a single word: *coexistence*.

With this major coexistence design goal in mind, CMAC provides a rich set of tools for coexistence and protection of incumbents services and introduces a novel coexistence beacon protocol (CBP) that allows 802.22 BSs with overlapping coverage areas to coordinate and efficiently share the scarce radio spectrum[1]. Additionally, CMAC includes channel management and measurement functions which provide further flexibility and efficiency in spectrum management.

In some respects, CMAC has been inspired by the IEEE 802.16 MAC standard for FBWA (Fixed Broadband Wireless Access) [5], and so it is a connection oriented MAC providing significant flexibility in terms of QoS support. On the other hand, many limitations and complexities of the 802.16 MAC have been identified when subjected to the 802.22 functional requirements, which have led us to introduce major changes including simplifications and enhancements.

To make an effective use of the radio spectrum, CMAC regulates downstream medium access by TDM (Time Division Multiplex), while the upstream is managed by using a DAMA (Demand Assigned Multiple Access) TDMA system. In CMAC, the BS manages all the activities within its 802.22 cell[2] and associated CPEs are always under the control of the BS.

Throughout the rest of this document, we provide a detailed description of CMAC, including its frame formats, protocols, algorithms, and coexistence mechanisms. As it is shown, this proposal either fulfils the mandatory requirements or does not preclude items which have been included as part of the mandatory requirements.

1.1 Overview

In an 802.22 cell, multiple CPEs are managed by a single BS which controls the medium access. The downstream direction (from BS to CPEs) is regulated by TDM and typically broadcast, while CPEs shall listen only to those messages addressed to them. The upstream direction (from CPEs to BS) is shared by CPEs on a demand basis, according to a DAMA TDMA scheme. 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 BS after receipt of a request from the user.

CMAC supports unicast (addressed to a single CPE), multicast (addressed to a group of CPEs) and broadcast (addressed to all CPEs in a cell) services. In particular for measurement activities, multicast management type of connections are very suitable as they allow vendor-specific clustering algorithms to be implemented and the measurement load to be shared.

CMAC implements a combination of access schemes that efficiently controls contention between users while at the same time meeting the delay and bandwidth requirements of each user application. This is accomplished through four different types of upstream scheduling mechanisms that are implemented using unsolicited bandwidth grants, polling, and contention procedures. The use of polling simplifies the access operation and guarantees that applications receive service on a deterministic basis if it is required. For example, real-time

[1] See Section 24 for definition.
[2] See Section 24 for definition.

applications like voice and video require service on a more uniform basis and sometimes on a very tightly controlled schedule. On the other hand, data applications are typically more delay tolerant and hence contention may be used to avoid the individual polling these CPEs. Contention is also suitable for saving resources, as it is possible to avoid polling CPEs that have been inactive for a long period of time.

CMAC is connection-oriented, and as such connections are a key component which require active maintenance and thus can be dynamically created, deleted, and changed as the need arises (maintenance requirements may vary depending upon the type of service). A connection defines both the mapping between peer convergence processes that utilize the MAC and a service flow (one connection per service flow). 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.

Another concept that is central to CMAC is that of a service flow on a connection. A service flow defines the QoS parameters for the PDUs that are exchanged on the connection, and provide a mechanism for upstream and downstream QoS management. In particular, they are integral to the bandwidth allocation process as a CPE requests upstream bandwidth on a per connection basis (implicitly identifying the service flow). The BS, in turn, grants bandwidth to a CPE as an aggregate of grants in response to per connection requests from the CPE.

1.2 Reference Architecture

In general, the major goals in the definition of a suitable reference architecture for 802.22 WRANs based on cognitive radios are with respect to flexibility and, at the same time, efficiency. With this in mind, here we propose the use of the reference architecture model depicted in Figure 1. As we can see, here we suggest that the MAC (i.e., CMAC in this proposal) can either natively support IP or else CSs (Convergence Sublayers) may be included in case more than one network layer technology needs to be supported.

The unique and distinctive characteristic of the proposed architecture is that it is scalable and so its capacity can be expanded over time, as the need arises. Hence, our proposed architecture is comprised of one or more PHY/MAC air interface module and a new entity called Spectrum Manager (SM). In our proposal, CMAC is designed to effectively deal with either single or multiple channels simultaneously, provided these multiple channels are *contiguous* in frequency domain (hereby called channel bonding). However, since we expect the TV bands to be fragmented and the occupation of a channel to vary with time, it is of paramount importance to design an air interface that can also take advantage of channels that are non-contiguously distributed over the entire TV band, and hence provide for increased capacity (hereby called channel aggregation). All these important characteristics are supported by our proposed architecture shown Figure 1.

The SM has a key role in the overall architecture as it allows the system to take advantage of non-contiguous channels while keeping the simplicity of CMAC (and also of the PHY) and allowing the system to scale (and also evolve) over time. In other words, the SM allows for an effective channel aggregation mechanism to be implemented. It is the entity (which could reside in the management layer of the 802 reference model) responsible for maintaining an updated global view of the target RF (i.e., TV bands in case of 802.22) spectrum (done through measurements as described later) and assigning the identified free channels for use by the various MAC/PHY modules (similar to a resource allocator). To allow a greater level of flexibility, the SM assigns channels (possibly disjoint unless directional antennas are used) to the MAC/PHY modules based on several criteria such as number of terminals associated to each these modules, traffic requirements, ranging (e.g., lower frequency channels could be assigned to the module dealing with farther away terminals), and so on. Figure 2 gives an example of possible channel assignments to a set of arbitrary modules.

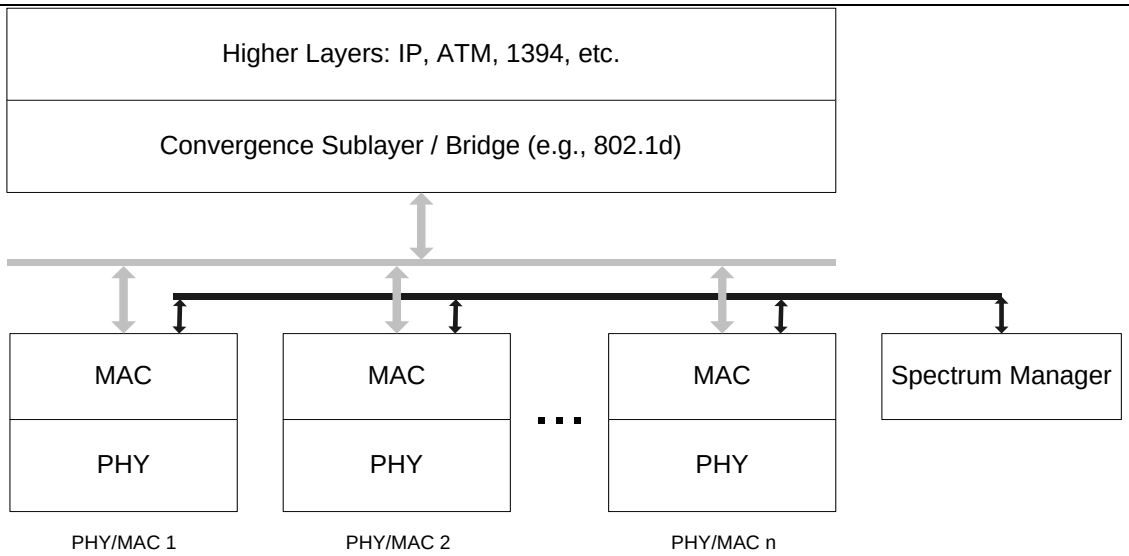

Figure 1 – The proposed reference architecture

The SM has also other capabilities such as taking requests from the various modules. For example, if an interference situation arises (e.g., with incumbents or other 802.22 cells) during normal operation in the channel, this is detected by a particular MAC which shall then be capable of taking appropriate actions to resolve the issue such as switching channels. In order to do this, the MAC may inquire the SM about the most suitable channel (or set of channels) to switch to (e.g., based on several criteria including the number of CPEs with which it is dealing with, the average CPE range, traffic type), and uses the informed response from the SM to perform the switching operation.

In the specific case of 802.22, we propose an instantiation of this proposed architecture. Since BSs can be more complex while CPEs should possess vary low complexity, we propose 802.22 BSs to support the architecture of Figure 1 by incorporating the SM and the possibility of progressively adding PHY/MAC air interface modules as demands increases. In other words, the capacity at the BS is augmented by increasing the number of PHY/MAC modules. However, from the CPE side only a single MAC/PHY module could be supported with no need to implement a SM, as CPEs are fully under control by the BS. With this arrangement, it would be possible to design the system with greater flexibility (and hence scalability), capacity, and efficiency, while keeping CPEs with low complexity.

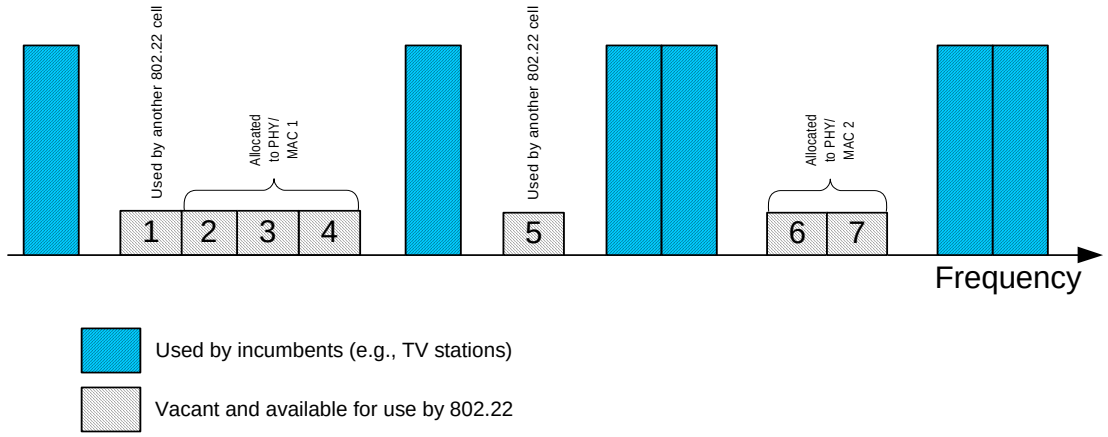

Figure 2 –Illustrative diagram of spectrum allocations. Channels 1 and 5 are in use by overlapping 802.22 cells, while channels 2-4 are allocated to PHY/MAC 1 (i.e., channel bonding is used and thus it is achieved 3 times as much bandwidth as a single channel) and channels 6-7 are assigned to PHY/MAC 2. Also, proper frequency separation is enforced in order to protect incumbent services.

From a practical implementation point of view, the SM could be implemented in many ways such as a programmable logic device giving high system flexibility. Algorithms could be developed within the SM that could make an efficient use of the radio spectrum as per various criteria (as outlined above), while the overall architecture would still provide a MAC and PHY with complexity comparable to existing wireless systems.

1.3 Basic Terms and Definitions

In this section we define some basic terminology that is needed to understand the hierarchical structure of the proposed MAC protocol. Further definitions and acronyms can be found in Sections 24 and 25.

- **Superframe** (see Section 3): Defined by the transmission from the BS of a preamble and the SCH. It is comprised of a number of **Frames**;
- **Frame** (see Section 4): Comprised of one DS and one US **Subframe**, where BS and CPEs communicate with each other;
- **Subframe** (see Section 4): Formed by a number of **Bursts**;
- **Burst** (see Section 4): Defined by a two dimensional segment of MAC logical channel (frequency) and MAC slot (time). It may comprise multiple **MAC PDUs** that are encoded with the same physical modulation and coding;
- **MAC PDU** (see Sections 4 and 6): The smallest unit of transmission/reception by the MAC. It is comprised of the MAC header, the payload, and CRC.

2. Addressing and Connections

Each 802.22 station shall have a 48-bit universal MAC address, as defined in IEEE Std 802-2001. This address uniquely defines the station from within the set of all possible vendors and equipment types. It is used during the initial ranging process to establish the appropriate connections for a CPE, as well as for coexistence purposes. It is also used as part of the authentication process by which the BS and CPE each verify the identity of the other.

Connections are identified by a 16-bit CID, thus allowing a total of 64K connections within each downstream and upstream channel. At CPE initialization, two pairs of management connections (upstream and downstream) shall be established between the CPE and the BS, while a third pair of management connections may be optionally generated[3]. The three pairs of connections reflect the fact that there are inherently three different levels of QoS for management traffic between a CPE and the BS. The basic connection is used by the BS MAC and CPE MAC to exchange short, time-urgent MAC management messages, whereas the primary management connection is used by the BS MAC and CPE MAC to exchange longer, more delay-tolerant MAC management messages (Table 24 specifies which MAC management messages are transferred on which of these two connections). Finally, the secondary management connection is used by the BS and CPE to transfer more delay tolerant, standards-based (e.g., DHCP, TFTP, and SNMP) messages which are carried in IP datagrams. Use of the secondary management connection is required only for managed CPE, and the messages carried on these types of connections may be packed and/or fragmented.

Requests for transmission are based on these CIDs, since the allowable bandwidth may differ for different connections, even within the same service type. For example, a CPE 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.

[3] The CIDs for these management connections shall be assigned in the RNG-RSP and REG-RSP messages, which provide a total three CID values. The same CID value is assigned to both members (upstream and downstream) of each connection pair.

CIDs may be used in many different ways. For example, multiple higher-layer sessions may operate over the same CID as in the case of various users within a company communicating with TCP/IP to different destinations. Here, since all users operate within the same overall service parameters (they are sharing the same CID), all of their traffic is pooled for request/grant purposes. Despite of that, 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. Thus, CIDs can be considered a connection identifier even for nominally connectionless traffic like IP, since it serves as a pointer to destination and context information.

3. General Superframe Structure

The superframe structure employed in CMAC is depicted in Figure 3, where it can be seen that it is comprised of three main parts:

- A PHY preamble (not discussed here)
- A Superframe Control Header (SCH) – see Section 5.1.
- A number of frames – see Section 4.

At the beginning of every superframe, the transmitter (i.e., the BS) sends special preamble and SCH (with a known modulation/coding) through each and every channel it is currently using for communication with its associated CPEs. Any device tuned to any of these channels and who synchronizes and receives the SCH, is able to obtain all the information it needs in order to establish communication with the transmitter (in this case, the BS). During the lifetime of a superframe, multiple MAC frames are transmitted which may span multiple channels and hence can provide better system capacity, range and data rate. During each MAC frame the BS has the responsibility to manage the upstream and downstream direction, which may include ordinary data communication, measurement activities, coexistence procedures, and so on.

The minimum superframe length (excluding the superframe preamble and SCH) shall be the maximum frame size (see Table 27). This is needed to guarantee that overlapping 802.22 BSs can efficiently coexist and share resources.

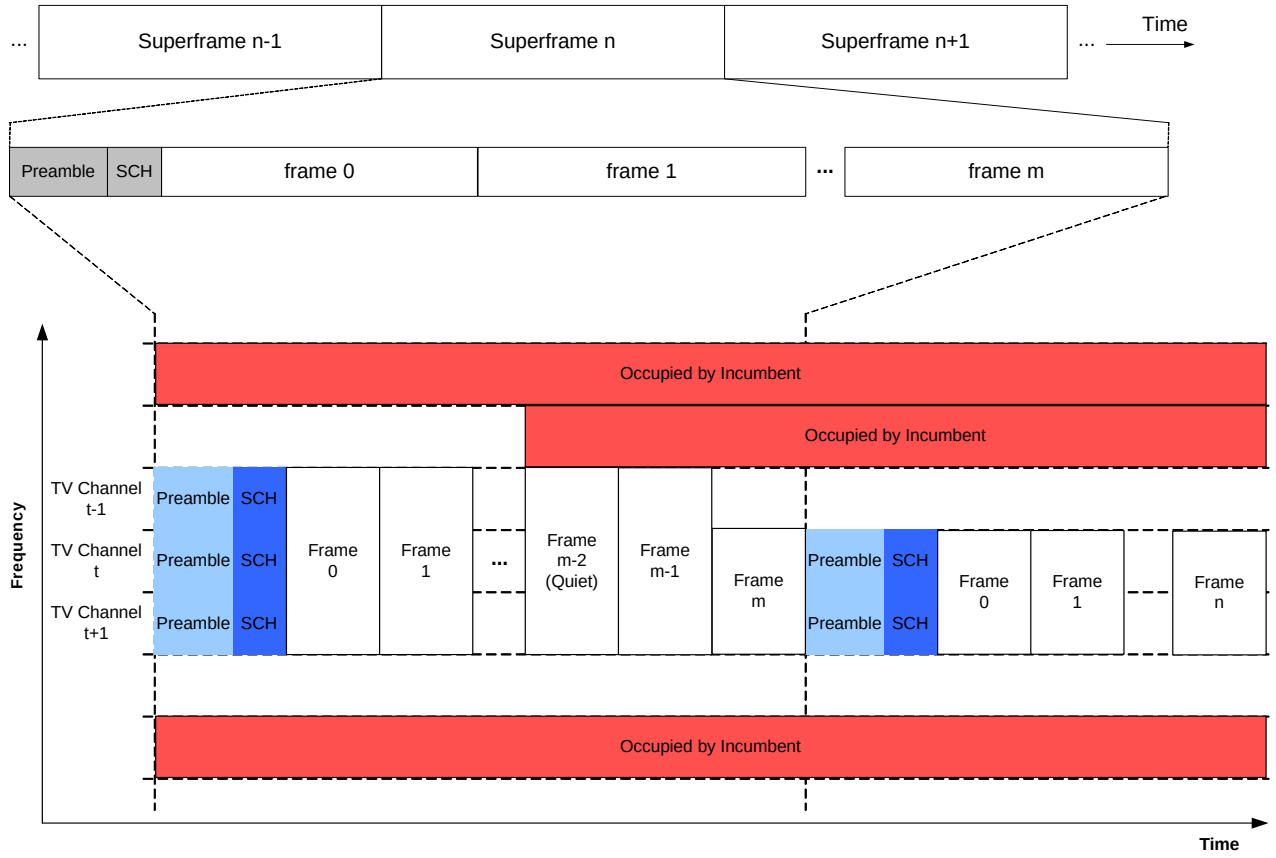

Figure 3 – General superframe structure

4. General Frame Structure

The frame structure is an important component to an efficient MAC protocol. In order to define it, we first have to decide on a particular multiplexing scheme as it may have an impact on the frame structure organization. After careful consideration of pros and cons, revision of documents 22-05-0020-00-0000_TDD_FDD_Tradeoffs.doc and 22-05-0022-02-0000_TDD_vs_FDD.doc, and given the envisioned application areas of IEEE 802.22 in rural and remote areas, we believe that supporting both FDD and TDD would incur unnecessary complexity. In addition, given that our current proposal incorporates a novel and distributed coexistence mechanism that overcomes the well-known limitation of TDD when there are multiple overlapping BSs, and that TDD has the advantage of not requiring frequency planning (which is of paramount importance for unlicensed operation in TV bands), we have decided in our proposal to support the use of TDD as the only duplexing mode of choice[4]. Besides, the protocol architecture adopted here (see 1.2) already brings with it some of the aspects of FDD as it allows flexible assignment of frequencies to various PHY/MAC air interface modules. The top-down frame structure employed in CMAC is illustrated in Figure 4, while Figure 5 depicts another view of the frame structure which is comprised of a series of time slots that can be allocated for either downstream of upstream traffic. It is critical to note at this point that burst depicted in Figure 4 can be transmitted across several logical subchannels as defined by the PHY. This is shown in Figure 28, which depicts how a frame can be actually transmitted (in time and frequency) by the PHY layer.

[4] Another added advantage of TDD is that it makes it considerably easier to extend a PMP MAC protocol (e.g., CMAC) to support mesh architecture. Although mesh support is currently out of scope in 802.22, it should be considered in selecting a suitable duplexing mode.

As we can see from Figure 4 and Figure 5, a frame is comprised of two parts: a predominantly downstream (DS) subframe and an upstream (US) subframe. The boundary between these two segments is adaptive, and so the control of the downstream and upstream capacity can be easily done. The downstream subframe consists of only one downstream PHY PDU with possible contention intervals for coexistence purposes. An upstream subframe consists of contention intervals scheduled for initialization (e.g., initial ranging), bandwidth request, UCS (Urgent Coexistence Situation) notification, and possibly coexistence purposes and one or multiple upstream PHY PDUs, each transmitted from different CPEs. Each frame is formed by an integral number of fixed size (MAC) slots, which are, in turn, an integral number of modulation symbols (currently, 1 MAC slot = 1 modulation symbol). As we shall see later, the definition of slots facilitates the implementation of various other schemes such as bandwidth allocation management and coexistence.

A downstream PHY PDU starts with a preamble (not discussed here), which is used for PHY synchronization. The preamble is followed by a FCH burst, which, as described in 5.2, specifies the burst profile and length of one or several downstream bursts immediately following the FCH. A DS-MAP message (see 8.2), if transmitted in the current frame as controlled by the Lost DS-MAP Interval parameter specified in Table 226, shall be the first MAC PDU in the burst following the FCH. An US-MAP message (see 8.4) shall immediately follow either the DS-MAP message (if one is transmitted) or the FCH. If UCD and DCD messages (see 8.3 and 8.1, respectively) are transmitted in the frame, they shall immediately follow the DS-MAP and US-MAP messages. Although burst #1 contains broadcast MAC control messages, it is not necessary to use the most robust well-known modulation/coding. A more efficient modulation/coding may be used if it is supported and applicable to all the CPEs of a BS.

In the upstream direction, if a CPE does not have any data to be transmitted in its US allocation, it shall transmit an US PHY burst containing a general MAC header (see 6.1.1) with its basic CID, together with a bandwidth request subheader (see 805512728.1332372.56.1.3.1) with BR = 0. This would allow the BS to reclaim this CPE's allocation and use the resource for some other purpose. Preceding upstream CPE PHY bursts, the BS may schedule up to three contention windows (see 14): the Initialization window is used for ranging, the BW window is for CPEs to request upstream bandwidth allocation from the BS, while the UCS Notification window is used by CPEs to report an urgent coexistence situation with incumbents. In particular, the UCS Notification window (see 21.1) can only be used by CPEs which do not currently have upstream allocations with the BS and need to report to the BS that it has detected an incumbent in one of the channels in use by the BS. See section 21.1.4 for details on the incumbent notification mechanism.

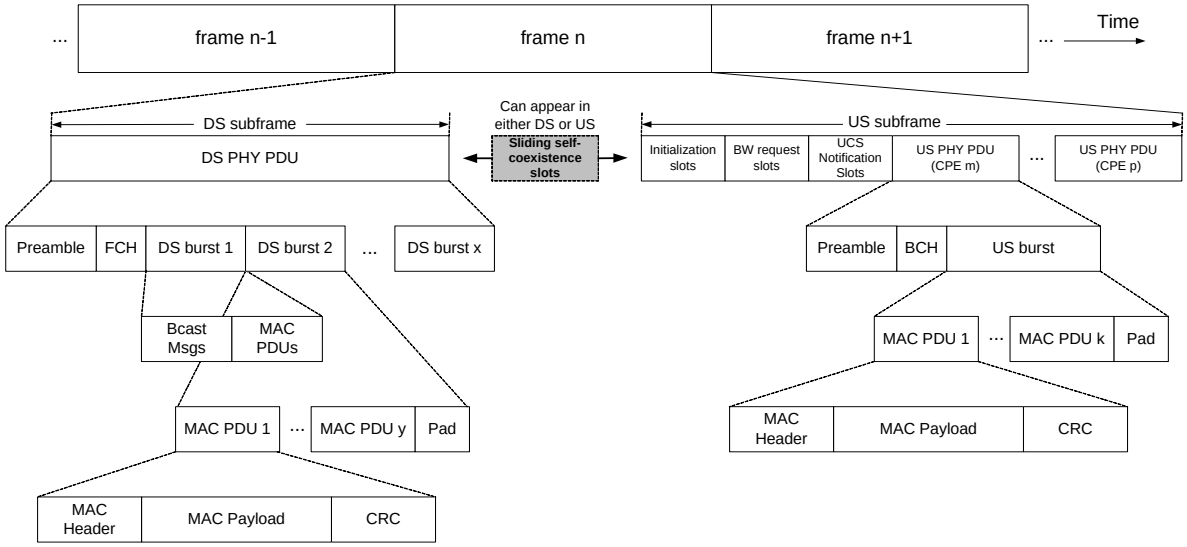

Figure 4 – Frame structure

Another important aspect to consider in the frame structure design is good coexistence with other overlapping 802.22 cells. It is common sense that coexistence is a key issue for the performance of any wireless technology which intends to operate in unlicensed bands (this is one of the reasons why IEEE has established the 802.19 Coexistence TAG, and the IEEE 802.16 WG established the 802.16h TG), as it ultimately impacts the widespread adoption of the technology. This is particularly true in the case of 802.22 where multiple BSs may be operating with overlapping coverage areas. Furthermore, since these BSs may belong to different operators, coordination or frequency planning cannot be assumed. In view of these aspects, 802.22 is faced with major challenge that may severely impact its success: coexistence with itself or, in other words, *self-coexistence*. Moreover, given the experiences in other WGs where their approach is to consider coexistence only after the standard has been approved (i.e., coexistence as an afterthought), here it is advocated that for coexistence to be really effective it has to be included as a key design goal of the air-interface, and not as an afterthought as it is often the case.

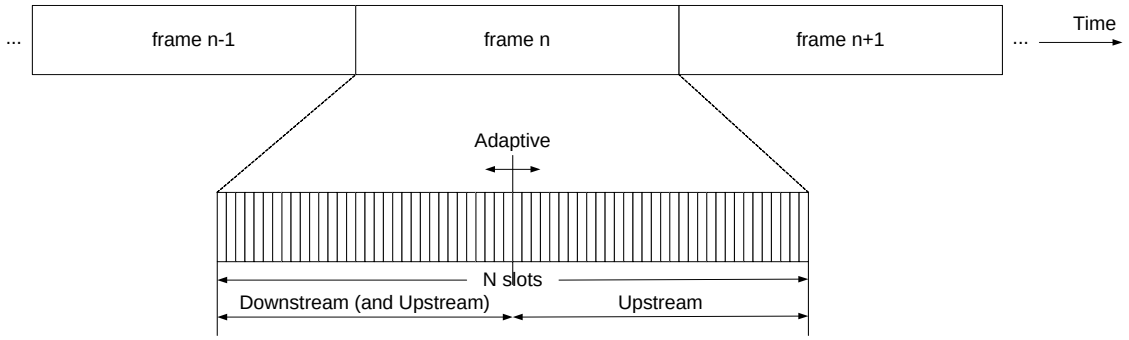

Figure 5 – The slotted structure of the CMAC frame. The boundary between upstream and downstream is adaptive.

To deal with scenarios where multiple 802.22 BSs possess overlapping coverage areas, we introduce the idea of Sliding Self-coexistence Slots (SSS) and coexistence beacons (see Section 6.1.2), all incorporated in CBP. The SSS (depicted in Figure 4) can appear in either the downstream or upstream part of a frame (although it is anticipated that it will mostly be scheduled in the upstream subframe to accommodate simpler receiver designs). Through CMAC, the BS is capable of scheduling SSS at any time during a frame with the goal of improving coexistence with nearby cells. These beacons are transmitted by selected CPEs and carry important information about the transmitter's ongoing service flows with the BS and also about the 802.22 cell as a whole. With this information, CPEs belonging to other collocated BSs are capable of implementing mechanisms that allow better inter-cell coexistence such as "interference-free" scheduling (see 21.2).

Whenever a CPE is neither receiving nor sending data to its BS, it is capable and ready to receive coexistence beacon packets possibly transmitted by nearby CPEs belonging to other BSs. More importantly, a frame synchronization mechanism is defined so that multiple collocated 802.22 cells can efficiently communicate with each other. These and other schemes make the proposed coexistence mechanisms to be highly effective. In addition, the BS also has the possibility to use the SSS and define what should be done during this time window. The SSS may have two purposes as indicated by the BS: listening for beacons (from nearby CPEs associated to other BSs) or transmitting beacons. Therefore, a high degree flexibility in terms of coexistence is supported which allows for improved system performance.

5. Control Headers

As can be seen in Figure 3 and Figure 4, there are a total of three control headers that provide essential information about the superframe (SCH), frame (FCH) and CPE burst (BCH). In the following subsections, we describe the content and goal each of these control headers in detail.

5.1 Superframe Control Header

The SCH specification is shown in Table 1. As we can see, it provides essential information about the 802.22 cell and brings with it many benefits including the support for channel bonding, a certain control over the time a device takes to join the network, better coexistence with incumbents and FCC Part 74 systems employing beacon signals, and so on.

The ST field serves the purpose of better coexistence amongst future wireless systems operating in the same band. It defines a way for systems to identify themselves and implement mechanisms for better coexistence[5]. The CT serves to identify the purpose for the transmission of the SCH. In CMAC, transmission of a SCH indicates two possible types of content may follow: a superframe (see 3) or a coexistence beacon (see 6.1.2). Therefore, the CT field is used to distinguish the type of content following the SCH. As we shall see later on, this distinction is needed to support CBP which is employed to implement better coexistence and sharing of the radio spectrum with other 802.22 systems. The use of the FS, FDC, Tx ID, CN and NC fields are straightforward and can be found in Table 1. Since a SCH may contain further IEs, the Length field is used to specify the total length of the SCH.

Table 1 – Superframe control header format

Syntax	Size	Notes
Superframe_Control_Header_Format() {		1 OFDM symbol long and transmitted with well-known modulation/coding (e.g., QPSK rate ½)
ST = 0	7 bits	System Type Indicates the type of the system using this band. See Table 2.
CT	1 bits	Content Type Indicates the type of the content following the transmission of the SCH. • Superframe = 0 • CBP Beacon = 1
FS	8 bits	Frames per Superframe Indicates the number of frames within a superframe. Frames within a superframe have a fixed size (Table 27) which preferably does not change across superframes.
FDC	8 bits	Frame Duration Code See Table 27
TTQP	16 bits	Time To Quiet Period The amount of time it will take for the next scheduled quiet period. This serves to synchronize the quiet period of overlapping BSs. As shown in Table 3, the TTQP is divided into two subfields: Time Scale and Time. The Time Scale subfield defines the scale for the Time subfield as shown in Table 148. The Time subfield consists of a 15 bit unsigned integer number.
DQP	16 bits	Duration of Quiet Period The estimated duration of the next scheduled quiet period. This is specified similar to the TTQP field.
PP	1 bit	Preamble Present Indicates whether the preamble of the following

[5] For this to happen, a common language (e.g., with a known preamble and modulation/coding) would have to be defined across all wireless systems operating in the same band. However, since we cannot predict all future wireless systems operating in the TV bands, what we propose here is for 802.22 systems and Part 74 devices operating in the TV bands (e.g., wireless microphones) to define a common signalling mechanism utilizing the same structure as specified in Table 1, where, for example, the ST field would be the first field and used to discern among the various system types.

		frame is present or not. For example, a frame preamble may not be needed when the cell is operating in a single physical (i.e., TV) channel
Tx ID	48 bits	Address that uniquely identifies the transmitter of the SCH (CPE or BS)
CN	8 bits	Channel Number Indicates the starting physical (i.e., TV) channel number used by the BS.
NC	5 bits	Number of Channels In case channel bonding is used, this field indicates the number of additional consecutive physical (i.e., TV) channels used by the BS. In the basic mode, NC = 3 (i.e., three TV channels).
Reserved	1 bit	
GIF	1 bit	Guard Interval Factor Specifies the GIF used by the PHY in the frame transmissions of this superframe. Pre-determined values are: 0 = O-QAM 4 = Default mode used for superframe transmission
Length	8 bits	The length of the information following the SCH
IEs	*Variable*	Optional Information Elements which would be transmitted after the SCH. They are: • MAC version (7.2) • Current transmit power (7.3) • Part 74 acknowledgement (7.7) • Location configuration information (8.21.3.1.4)
HCS	8 bits	Header Check Sequence See Table 6
}		

Table 2 – System type

Value	System
0	802.22 WRAN
1	Part 74 (see 21.1.7)
2	802.11 WLAN
3	802.15 WPAN
4	802.16 WMAN
5-127	*Reserved*

Table 3 – TTQP field

Bits: 1	15
Time Scale	Time

5.2 Frame Control Header

The format of the FCH is shown in Table 4. Since FCH decoding is critical, the FCH shall be transmitted using a robust modulation/coding. The FCH contains the length of the four (DS-MAP, US-MAP, DCD, UCD) critical downstream bursts that may immediately follow the FCH (Length == 0 indicates the absence of a given burst). Location and profile of other bursts are specified in the DS-MAP (see 8.2) and US-MAP (see 8.4) messages. Lastly, a HCS field occupies the last byte of the FCH.

Table 4 – Frame control header format

Syntax	Size	Notes
Frame_Control_Header_Format() {		Transmitted with well-known modulation/coding (e.g., QPSK rate ½)
DS-MAP Length	8 bits	Size in bits. If there are any unused IEs, the first unused IE shall have all fields encoded as zeros.
US-MAP Length	8 bits	Size in bits. If there are any unused IEs, the first unused IE shall have all fields encoded as zeros.
DCD Length	8 bits	Size in bits. If there are any unused IEs, the first unused IE shall have all fields encoded as zeros.
UCD Length	8 bits	Size in bits. If there are any unused IEs, the first unused IE shall have all fields encoded as zeros.
Repetition Indication	1 bit	
Short Training Sequence Present	1 bit	Indicates, for the next frame preamble, if a short training sequence is present. For the first frame of the superframe, the short training sequence shall always be present.
Reserved	6 bits	
HCS	8 bits	Header Check Sequence See Table 6
}		

5.3 Burst Control Header

The BCH format is shown in Table 5. The BCH is transmitted by CPEs before any upstream burst. The main purpose of including the BCH in the overall upstream subframe is to reliably and uniquely identify the CPE who has transmitted a given PHY burst, and its associated BS. This is an important requirement for effective coexistence and protection of incumbent services, as it allows the system to detect and pinpoint other 802.22 stations which may be mistakenly using channels currently occupied by incumbents.

Table 5 – Burst control header format

Syntax	Size	Notes
Burst_Control_Header_Format() {		
CPE ID	48 bits	Address that uniquely identifies the CPE
BS ID	48 bits	Address that uniquely identifies the BS to which this CPE is associated with
HCS	8 bits	Header Check Sequence See Table 6
}		

6. MAC PDU Formats

The proposed MAC PDUs is illustrated in Figure 6 (Figure 4 depicts how the MAC PDU fits in the overall frame structure when used for intra-cell communication). Each PDU begins with a fixed-length generic MAC header, which shall be followed by a Payload. The Payload shall consist of zero or more subheaders, zero or more IEs,

and zero or more MAC SDUs and/or fragments thereof. The payload information may vary in length, so that a MAC PDU may represent a variable number of bytes. This allows the MAC to tunnel various higher-layer traffic types without knowledge of the formats or bit patterns of those messages. Finally, a MAC PDU shall carry a CRC (see 9.5).

Figure 6 – MAC PDU format

6.1 MAC Headers

A total of two MAC headers are defined in CMAC: the general MAC header used for intra-cell communication and the beacon MAC header used for inter-cell communication. A MAC PDU employing the beacon MAC header shall always be preceded by the SCH, which is not the case for a MAC PDU using the general MAC header. As stated before, the general MAC header can be seen as an intra-cell MAC header, as it is utilized solely for communication between BS and CPEs within a cell. On the other hand, the beacon MAC header is used by the CBP protocol with the intention of inter-cell communication and to foster good coexistence amongst overlapping 802.22 BS, and as such can be seen as an inter-cell MAC header.

In addition to these headers, there is also the possibility of including subheaders and special payloads associated to the general MAC header. In the following subsections we present the MAC headers defined by CMAC, together with the subheaders and special payloads.

6.1.1 General

The format of the general MAC header is shown in Table 6. This header begins each MAC PDU containing either higher layer traffic or MAC management data.

Since CMAC is a connection-oriented MAC, an important component of the general MAC header is the CID which serves to identify an existing service flow between the BS and CPE. Two other critical fields included in the header for the purpose of coexistence are the UCS and the CN. These fields are used by CPEs to quickly signal the BS of a newly detected urgent coexistence situation with incumbents. For example, these can be utilized to cope with the case where the BS is engaged in communication with a CPE and incumbents are detected by this CPE in a channel currently in use. Here, this urgent situation can be immediately conveyed to the BS by having the CPE set both the UCS bit to 1 and the CN field to the corresponding channel number now occupied by an incumbent service. The general MAC header also includes other fields (e.g., security related), and these can be found in Table 6.

Table 6 – Generic MAC header format

Syntax	Size	Notes
General_MAC_Header_Format() {		
EC	1 bits	Encryption control 0 = payload is not encrypted 1 = payload is encrypted
Type	6 bits	Indicates the subheaders and special payload types present in the message payload See Table 7
Reserved	3 bits	*Reserved*
EKS	2 bits	Encryption key sequence The index of the traffic encryption key (TEK) and initialization

		vector used to encrypt the payload. This field is only meaningful if the EC field is set to 1
UCS	1 bit	Urgent Coexistence Situation Used by the CPE to indicate the BS about an urgent coexistence situation with incumbents in the channel(s) currently being used by the BS. 0 = no incumbent (default) 1 = incumbent detected
CN	8 bits	Channel Number It indicates exactly which of the channel(s) is under an urgent coexistence situation with incumbents (UCS == 1) or is vacant (UCS ==0).
Length	11 bits	The length in bytes of the MAC PDU including the MAC header and the CRC
CID	16 bits	Connection identifier
HCS	8 bits	Header check sequence The transmitter shall calculate the HCS value for the first six bytes of the header, and insert the result into the HCS field (the last byte of the MAC header). It shall be the remainder of the division (Modulo 2) by the generator polynomial $g(D = D^8 + D^2 + D + 1$ of the polynomial D^8 multiplied by the content of the header excluding the HCS field. (Example: [EC Type]=0x80, BR=0xAAAA, CID=0x0F0F; HCS should then be set to 0xD5).
}		

Table 7 – Encoding of the Type field

Type bit	Value
#5 most significant bit (MSB)	Bandwidth Request subheader Indicates whether this is a bandwidth request frame, and hence contains a special payload related to bandwidth allocation (see Table 10) 1 = present; 0 = absent
#4	ARQ feedback payload 1 = present; 0 = absent
#3	Extended type Indicates whether the present Packing or Fragmentation Subheaders is Extended 1 = Extended 0 = not Extended. Applicable to connections where ARQ is not enabled
#2	Fragmentation subheader 1 = present; 0 = absent
#1	Packing subheader 1 = present; 0 = absent
#0	Downstream: FAST-FEEDBACK Allocation subheader Upstream: Grant Management subheader 1 = present; 0 = absent

6.1.2 Beacon

In the context of CMAC, CPE beacons are inter-cell packets used by CBP only (see 21.2.1) and which are transmitted with the goal of improving coexistence amongst overlapping 802.22 cells (i.e., self-coexistence). These beacons are transmitted under the control of the BS during an active self-coexistence window and share the same beacon MAC header as described in Table 8. Since their goal is to improve self-coexistence, these beacons are sometimes referred to as *coexistence beacons*.

As discussed in 21.2.1, the beacon MAC header is utilized in the transmission of CBP packets. Overall, the CBP packets provide information about the CPE's current cell of attachment as well as the CPE's traffic flows with its

BS (see 21.2.1). Specifically, conveying the information about the traffic flows of a CPE with its BS is the responsibility of the beacon MAC header and subsequent IEs carried in the payload.

To describe its traffic flows with the BS, the beacons transmitted by CPEs shall carry beacon IEs (see Table 9) in their payload that provide necessary and sufficient information about the CPE's traffic reservations with the BS (CPEs with no traffic reservations with the BS need not transmit beacons). Stations (either CPEs or BSs) belonging to other 802.22 BSs and who receive a coexistence beacon, can then improve coexistence amongst BSs through a variety of mechanisms such as interference-free scheduling. These beacon IEs shall be the only type of information present in the payload of a beacon PDU, that is, no other information other than beacon IE shall be present in the payload.

Table 8 – Beacon MAC header format

Syntax	Size	Notes
Beacon_MAC_Header_Format() {		
Frame Number	8 bits	The frame number in which this message is transmitted. See definition in Table 26
Transmission Offset	8 bits	Indicates the offset (in units of slots) relative to the start of the first slot of the PHY PDU (including preamble) where the current frame (i.e., the beacon itself) is transmitted. The time instants indicated by the Transmission Offset values are the transmission times of the first slot of the beacon including preamble (if present).
BS ID	48 bits	Address that uniquely identifies the BS to which this CPE is associated with
Channel Number for Backup	8 bits	See definition in Table 26
Number of Channels for Backup	8 bits	See definition in Table 26
Starting DS Allocation Channel	8 bits	Starting logical channel from the perspective of the overall DS allocation by the BS.
Ending DS Allocation Channel	8 bits	Ending logical channel from the perspective of the overall DS allocation by the BS.
Ending DS Allocation Slot	7 bits	Ending time slot from the perspective of the overall DS allocation by the BS.
Starting US Allocation Channel	8 bits	Starting logical channel from the perspective of the overall US allocation by the BS.
Ending US Allocation Channel	8 bits	Ending logical channel from the perspective of the overall US allocation by the BS.
Starting US Allocation Slot	7 bits	Starting time slot from the perspective of the overall US allocation by the BS.
Length	8 bits	See definition in Table 6
HCS	8 bits	See definition in Table 6
}		

Table 9 – Beacon IE

Syntax	Size	Notes
CPE_Beacon_IE_Format() {		
Element ID	8 bits	
Length	8 bits	
Direction	1 bit	Indicates whether this reservation is for Upstream direction (set to 0) or Downstream direction (set to 1)
Reserved	4 bits	*Reserved*
Frame Offset	8 bits	Indicates the offset (in units of slots) of this CPE's reservation with the BS (whether DS or US) relative to the start of the first slot of the PHY PDU (including preamble) where the frame is transmitted. The time instants indicated by the Frame Offset values are the transmission times of the first slot of the CPE reservation including

		preamble (if present).
Duration	8 bits	Indicates the duration (in units of slot duration) of this CPE's reservation with the BS (whether DS or US)
CoS	3 bits	Indicates the priority of the reservation
Channel Number	8 bits	The initial logical channel number of this reservation
Number of Channels	8 bits	The number of logical channels that this reservation spans
}		

6.1.3 MAC Subheaders and Special Payloads

Six types of subheaders may be present. The per-PDU subheaders (i.e., Bandwidth Request, Fragmentation, FASTFEEDBACK_Allocation, and Grant Management) may be inserted in MAC PDUs immediately following the Generic MAC header. If both the Fragmentation subheader and Grant Management subheader are indicated, the Grant Management subheader shall come first. If both the Grant Management subheader and Bandwidth Request subheader are indicated, the Grant Management subheader shall come first. The FAST-FEEDBACK Allocation subheader shall always appear as the last per-PDU subheader.

The only per-SDU subheader is the Packing subheader. It may be inserted before each MAC SDU if so indicated by the Type field. The Packing and Fragmentation subheaders are mutually exclusive and shall not both be present within the same MAC PDU.

When present, per-PDU subheaders shall always precede the first per-SDU subheader.

6.1.3.1 Bandwidth Request Subheader

Table 10 – Bandwidth Request subheader format

Syntax	Size	Notes
BWR_Subheader_Format() {		
Type	3 bits	Indicates the type of the bandwidth request header 000 = incremental 001 = aggregate
BR	21 bits	Bandwidth request The number of bytes of upstream bandwidth requested by the CPE. The bandwidth request is for the CID. The request shall not include any PHY overhead.
Length	8 bits	Length of possible bandwidth allocation constraints IEs (see Section 8.23.2.1) relative to this request. If present, these IEs shall immediately follow all present subheaders or special payloads.
}		

6.1.3.2 Fragmentation Subheader

Table 11 – Fragmentation subheader format

Syntax	Size	Notes
Fragmentation_Subheader_Format() {		
FC	2 bits	Indicates the fragmentation state of the payload: 00 = no fragmentation 01 = last fragment 10 = first fragment 11 = continuing (middle) fragment
if (ARQ-enabled Connection)		
BSN	11 bits	Sequence number of the first block in the current SDU

		fragment
else {		
if (**Type** bit **Extended Type**)		Table 7
FSN	11 bits	Sequence number of the current SDU fragment. This field increments by one (modulo 2048) for each fragment, including unfragmented SDUs.
else		
FSN	3 bits	Sequence number of the current SDU fragment. This field increments by one (modulo 8) for each fragment, including unfragmented SDUs.
}		
Reserved	3 bits	Shall be set to zero
}		

6.1.3.3 Grant Management Subheader

Table 12 – Grant management subheader format

Syntax	Size	Notes
Grant_Management_Subheader_Format() {		
if (scheduling service type = UGS) {		
SI	1 bit	Slip Indicator 0 = No action 1 = Used by the CPE to indicate a slip of upstream grants relative to the upstream queue depth
PM	1 bit	Poll-Me 0 = No action 1 = Used by the CPE to request a bandwidth poll
Reserved	6 bits	Shall be set to zero
}		
}		

6.1.3.4 Packing Subheader

Table 13 – Packing subheader format

Syntax	Size	Notes
Packing_Subheader_Format() {		
FC	2 bits	Table 11
if (ARQ-enabled Connection)		
BSN	11 bits	Table 11
else {		
if (**Type** bit **Extended Type**)		Table 7
FSN	11 bits	Table 11
else		
FSN	3 bits	Table 11
}		
Length	11 bits	
}		

6.1.3.5 ARQ Feedback Payload

If the ARQ Feedback Payload bit in the MAC Type field (see Table 7) is set, the ARQ Feedback Payload shall be transported. If packing is used, it shall be transported as the first packed payload.

6.1.3.6 FAST-FEEDBACK Allocation Subheader

Table 14 – FAST-FEEDBACK allocation subheader format

Syntax	Size	Notes
FAST-FEEDBACK_allocation_Subheader_Format() {		
Allocation Offset	6 bits	Defines the offset, in units of slots, from the beginning of the FAST-FEEBACK upstream bandwidth allocation, of the slot in which the CPE servicing the CID appearing in the MAC generic header, must send a FAST-FEEBACK feedback message for the connection associated with the CID value. Range of values 0 to 63. The allocation applies to the UL subframe of the next frame.
Feedback Type	2 bits	00 = Fast downstream measurement
}		

7. Common IEs

In this section the common IEs are defined. These IEs have a global scope and are used throughput the MAC specification. Table 15 describes the common IEs used in CMAC.

Table 15 – Common IEs

Element ID	Name
149	HMAC (hashed message authentication code)
148	MAC version
147	Current transmit power
146	Downstream service flow
145	Upstream service flow
144	Vendor ID
143	Vender-specific information
142	Part 74 Acknowledgement

7.1 HMAC

Table 16 – HMAC definition

Element ID	Length (bytes)	Value	Scope
149	21	See Table 17	DSx-REQ, DSx-RSP, DSx-ACK, REG-REQ, REQ-RSP, RES-CMD, DREG-CMD, RFTP-CPLT

Table 17 – HMAC value field

Field	Length (bits)	Notes
Reserved	4	
HMAC Key Sequence Number	4	

HMAC-Digest	160 bits	HMAC with SHA-1

7.2 MAC Version

Table 18 – MAC version definition

Element ID	Length (bytes)	Value	Scope
148	1	802.22 version supported	SCH, DCD, RNG-REQ

7.3 Current Transmit Power

Table 19 – Current transmit power definition

Element ID	Length (bytes)	Value	Scope
147	1	Current transmitted power	SCH, CBC-REQ

7.4 Service Flow Descriptors

Table 20 – Downstream/Upstream service flow definition

Element ID	Length (bytes)	Value	Scope
146	*Variable*	Compound: Downstream service flow (8.8.10)	DSx-REQ, DSx-RSP, DSx-ACK
145	*Variable*	Compound: Upstream service flow (8.8.10)	DSx-REQ, DSx-RSP, DSx-ACK

7.5 Vendor ID

Table 21 – Vendor ID definition

Element ID	Length (bytes)	Value	Scope
144	3	Vendor ID information	REQ-REQ, REQ-RSP, DSx-REQ, DSx-RSP, DSx-ACK

7.6 Vendor-Specific Information

Table 22 – Vendor-specific information definition

Element ID	Length	Value
143	*Variable*	Vendor-specific information

7.7 Part 74 Acknowledgement

Table 23 – Part 74 acknowledgment definition

Element ID	Length (bytes)	Value	Scope
142	1	Channel number where notification has been received	SCH

8. Management Messages

As can be seen in Table 24, CMAC defines a collection of management messages to support and implement its basic functions. All these messages are carried in the payload of a MAC PDU, and share the same frame structure as depicted in Figure 7. Management messages begin with a Type field that uniquely identifies the message in question, while its payload varies according to the message type. As for transmission, management messages can only be transmitted in Initial Ranging, Basic, Primary, Multicast Management or Broadcast type of CIDs (see Table 227). No other types of CIDs shall carry management messages.

In the following sections we describe each of the management messages shown in Table 24 in detail.

Table 24 – Management messages

Type	Message	Description	Connection
0	DCD	Downstream Channel Descriptor	Broadcast
1	DS-MAP	Downstream Access Definition	Broadcast
2	UCD	Upstream Channel Descriptor	Broadcast
3	US-MAP	Upstream Access Definition	Broadcast
4	RNG-REQ	Ranging Request	Initial Ranging or Basic
5	RNG-RSP	Ranging Response	Initial Ranging or Basic
6	REG-REQ	Registration Request	Primary Management
7	REG-RSP	Registration Response	Primary Management
8		*Reserved*	
9	PKM-REQ	Privacy Key Management Request	Primary Management
10	PKM-RSP	Privacy Key Management Response	Primary Management
11	DSA-REQ	Dynamic Service Addition Request	Primary Management
12	DSA-RSP	Dynamic Service Addition Response	Primary Management
13	DSA-ACK	Dynamic Service Addition Acknowledge	Primary Management
14	DSC-REQ	Dynamic Service Change Request	Primary Management
15	DSC-RSP	Dynamic Service Change Response	Primary Management
16	DSC-ACK	Dynamic Service Change Acknowledge	Primary Management
17	DSD-REQ	Dynamic Service Deletion Request	Primary Management
18	DSD-RSP	Dynamic Service Deletion Response	Primary Management
19		*Reserved*	
20		*Reserved*	
21	MCA-REQ	Multicast Assignment Request	Primary Management
22	MCA-RSP	Multicast Assignment Response	Primary Management
23	DBPC-REQ	Downstream Burst Profile Change Request	Basic
24	DBPC-RSP	Downstream Burst Profile Change Response	Basic
25	RES-CMD	Reset Command	Basic
26	CBC-REQ	CPE Basic Capability Request	Basic
27	CBC-RSP	CPE Basic Capability Response	Basic
28		*Reserved*	
29	DREG-CMD	De/Re-register Command	Basic
30	DSX-RVD	DSx Receive Message	Primary Management
31	TFTP-CPLT	Config File TFTP Complete Message	Primary Management
32	TFTP-RSP	Config File TFTP Complete Response	Primary Management
33	ARQ-Feedback	Standalone ARQ Feedback	Basic

34	ARQ-Discard	ARQ Discard Message	Basic
35	ARQ-Reset	ARQ Reset Message	Basic
36	CPE-FPC	CPE Fast Power Control	Broadcast
37	DREG-REQ	CPE De-registration Message	Basic
38		*Reserved*	
39	BLM-REQ	Bulk Measurement Request	Primary Management, Multicast Management or Broadcast
40	BLM-RSP	Bulk Measurement Response	Primary Management
41	BLM-REP	Bulk Measurement Report	Primary Management
42	BLM-ACK	Bulk Measurement Acknowledgement	Primary Management
43	CHT-REQ	Channel Terminate Request	Basic, Multicast Management or Broadcast
44	CHT-RSP	Channel Terminate Response	Basic
45	CHA-REQ	Channel Add Request	Primary Management, Multicast Management or Broadcast
46	CHA-RSP	Channel Add Response	Primary Management
47	CHS-REQ	Channel Switch Request	Basic, Multicast Management or Broadcast
48	CHS-RSP	Channel Switch Response	Basic
49	CHQ-REQ	Channel Quiet Request	Basic, Multicast Management or Broadcast
50	CHQ-RSP	Channel Quiet Response	Basic
51	CHO-UPD	Channel Occupancy Update	Basic, Multicast Management or Broadcast
52	TRC-REQ	Traffic Constraint Request	Primary Management
53	TRC-REP	Traffic Constraint Report	Primary Management
54	TMO-REQ	Timeout Request	Primary Management
55	TMO-RSP	Timeout Response	Primary Management
56	FSL-REQ	Frame Slide Request	Basic, Multicast Management or Broadcast
57	FSL-RSP	Frame Slide Response	Basic
58	AAS-CFB-REQ	AAS Channel Feedback Request	Basic
59	AAS-CFB-RSP	AAS Channel Feedback Response	Basic

Figure 7 – The general management frame structure

8.1 Downstream Channel Descriptor (DCD)

The format of a DCD message is shown in Table 25. This message shall be transmitted by the BS at a periodic interval (Table 226) to define the characteristics of a downstream physical channel.

Table 25 – Message format

Syntax	Size	Notes
DCD_Message_Format() {		
Management Message Type = 0	8 bits	
Downstream Channel ID	8 bits	The identifier of the downstream channel to which this message refers. This identifier is arbitrarily

		chosen by the BS and is unique only within the MAC domain. This acts as a local identifier for transactions such as ranging.
Configuration Change Count	8 bits	Incremented by one (modulo 256) by the BS whenever any of the values of this channel descriptor change. If the value of this count in a subsequent DCD remains the same, the CPE can quickly decide that the remaining fields have not changed and may be able to disregard the remainder of the message.
Coexistence Backoff Start	8 bits	
Coexistence Backoff End	8 bits	
Information Elements (IEs) for the overall channel	*Variable*	Table 26
Begin PHY Specific Section {		
for (i = 1; i ≤ n; i++) {		
Downstream_Burst_Profile	*Variable*	PHY specific (Table 28)
}		
}		
}		

8.1.1 Channel IEs

Table 26 – DCD channel information elements

Name	Element ID (1 byte)	Length (bytes)	Value
Downstream_Burst_Profile	1		Value reserved for the burst profile (see Table 28)
BS EIRP	2	2	Singed in units of dBm
Reserved	3		
TTG	4	1	TTG in slots
RTG	5	1	RTG in slots
$RSS_{IR,max}$	6	2	Initial ranging maximum received signal strength at BS in units of 1dBm
BS ID	7	6	Base Station ID. This is needed in addition to the ID contained in the SCH, so that CPEs can distinguish messages from different collocated BSs.
Frame Duration Code	8	1	Time duration of the frame (Table 27)
Frame Number	9	1	The number of the frame containing the DCD message
Channel Action	10	1	Action to be taken by all CPEs in a cell. 0 = None 1 = Switch 2 = Add 3 = Remove 4 = Quiet
Action Frame Number	11	1	The starting frame number at which the Channel Action shall be performed by all CPEs.
Action Duration	12	2	This is valid only for quiet periods (Action = 4). • If this field is set to a value different from 0 (zero): It indicates the duration (expressed in slots), not including the Action Frame Number. Once this duration is over, normal operation resumes in the channel by the BS. If quiet periods are already scheduled after Action Frame Number, the value specified in

			this field shall override the length of the first quiet period after Action Frame Number. During this time, the CPE shall sense for incumbents only. If more detailed specification is needed, please see 8.21.7. • If this field is set to 0 (zero): it serves to indicate CPEs that the first quiet period after Action Frame Number is cancelled. Therefore, normal data transmission shall be carried out during this time.
Action Channel Number	13	1	
Action Number of Channels	14	1	
Channel Number for Backup	15	1	The backup channel to be used by CPEs in case of loss of communication with the BS due to incumbents. If possible, the backup channel(s) shall be a disjoint set with the current operating channels.
Number of Channels for Backup	16	1	The number of backup channels. To maximize the success probability that the backup channel is vacant when needed, this field should be set to 1.
Sensing RTG	17	1	If set to 1, the CPE shall perform fast sensing (i.e., energy detection) during RTG.
Channel Number for Sensing RTG	18	1	The channel number that is to be sensed during the Sensing RTG.
MAC version	148	1	Table 18

8.1.1.1 Frame Duration Codes

Table 27 illustrates the possible frame durations. Note that these frame durations typically are not integer multiples of one OFDMA symbol duration. Therefore, some time padding may be necessary between the last useful OFDMA symbol of a frame and the beginning of the next frame. In addition, in the TDD case, note that the RTG and TTG guard intervals must be included in a frame.

Table 27 – Frame duration codes

Code	Frame Duration (ms)	Frames per Second
0	20	50
1	40	25
2	80	12.5

8.1.2 Downstream Burst Profile

Table 28 – Downstream burst profile format

Syntax	Size	Notes
Downstream_Burst_Profile_Format() {		
Type = 1	8 bits	
Length	8 bits	
Reserved	4 bits	
DIUC	4 bits	Table 33
Information Elements (IEs)	*Variable*	Table 29
}		

Table 29 –Information elements

Name	Element ID (1 byte)	Length (bytes)	Value
Frequency	1	4	Downstream frequency (kHz)
FEC code type and modulation type	150	1	Combination of: • Spreading • (0ffset)QPSK;(0ffset)16-QAM; (0ffset)64-QAM; • Coding rates : ½;2/3; ¾ • RS+CC/CC; CTC codes Detailed specification TBD.
DIUC mandatory exit threshold	151	1	0-63.75 dB CINR at or below where this DIUC can no longer be used and where change to a more robust DIUC is required (in 0.25 dB units)
DIUC minimum entry threshold	152	1	0-63.75 dB The minimum CINR required to start using this DIUC when changing from a more robust DIUC is required (in 0.25 dB units)

8.2 Downstream Map (DS-MAP)

The format of a DS-MAP message is shown in Table 30. The DS-MAP message defines the access to the downstream information. If the length of the DS-MAP message is a non-integral number of bytes, the length field in the MAC header is rounded up to the next integral number of bytes. The message shall be padded to match this length, but the CPE shall disregard the four pad bits.

Table 30 – Message format

Syntax	Size	Notes
DS-MAP_Message_Format() {		
Management Message Type = 1	8 bits	
Synchronization Field	16 bits	Table 31
DCD Count	8 bits	Matches the value of the configuration change count of the DCD, which describes the downstream burst profiles that apply to this map.
BS ID	48 bits	See Table 26
Begin PHY Specific Section {		
for (i = 1; i ≤ n; i++) {		
DS-MAP_IE()	*Variable*	PHY specific (8.2.1)
}		
}		
If (!byte_boundary)		
Padding Nibble	4 bits	
}		

Table 31 – Synchronization field

Syntax	Size	Notes
Synchronization_field() {		
Frame Duration Code	8 bits	Table 26
Frame Number	8 bits	Table 26
}		

8.2.1 DS-MAP IE

Table 32 – DS-MAP information element

Syntax	Size	Notes
DS-MAP_IE() {		
DIUC	4 bits	0.0.08.2.1.1
If (DIUC == 15)		
Extended DIUC Dependent IE	*Variable*	805502985.512.58.2.1.2
else {		
If (INCLUDE_CID) {		The DS-MAP starts with INCLUDE_CID =0. INCLUDE_CID is toggled between 0 and 1 by the CID-SWITCH_IE() – see Table 36
N_CID	7 bits	Number of CIDs assigned for this IE
for (i=0; i<N_CID; i++)		
CID	16 bits	Unicast/Multicast/Broadcast CID. In case of a dummy allocation, the BS shall use the Padding CID – see Table 227
}		
Channel Offset	7 bits	
Number of Channels	7 bits	
Slot Offset	8 bits	
Number of Slots	8 bits	In number of PHY symbols. Specification of the duration in the downstream allows for the BS to introduce gaps in the downstream data transmission, and hence provide better support for coexistence.
Boosting	3 bits	000: normal (not boosted) 001: +6dB 010: -6dB 011: +9dB 100: +3dB 101: -3dB 110: -9dB 111: -12dB
}		
If (!byte_boundary)		
Padding Nibble	4 bits	
}		

8.2.1.1 DIUC Allocations

Table 33 illustrates the various DIUC values used in CMAC. In particular, it is important to highlight the use of the two self-coexistence DIUC intervals: the active and the passive mode. The coexistence DIUC in the active mode shall be used by CPEs for transmitting CBP packets, and hence its definition as being active. The coexistence DIUC interval in the passive mode, on the other hand, can be seen as an in-frame quiet period, where the intention is to sense for the presence of other co-located 802.22 cells in the area. More specifically, this interval is used by 802.22 stations to look for both CBP and BS beacons. This allows flexibility for the 802.22 system to efficiently coexist with other nearby 802.22 systems.

Table 33 – DIUC values

DIUC	Usage
0	Self-Coexistence (Active Mode)
1	Self-Coexistence (Passive Mode)
2-12	Burst Profiles
13	Gap/PAPR Reduction

14	End of Map
15	Extended DIUC

The DIUC = 13 may be used for allocation of channels for PAPR reduction schemes. DIUC = 13 may also be used by the BS to create coverage enhancing safety zones. This is intended to provide reduced interference zones within the coverage area of the BS. The reduced interference zones are useful when the BS interfere with other BS. In such situations, the reduced interference zones may be used by the interfered BS to transmit data to CPEs that are registered with it, which would otherwise suffer from interference.

8.2.1.2 DS-MAP Extended DIUC IE

Table 34 – DS-MAP extended IE general format

Syntax	Size	Notes
DS_Extended_IE() {		
Extended DIUC	4 bits	
Length	4 bits	
Unspecified Data	*Variable*	
}		

8.2.1.2.1 DS-MAP Dummy IE

Table 35 – DS-MAP dummy IE format

Syntax	Size	Notes
Dummy_IE() {		
Extended DIUC	4 bits	0x05
Length	4 bits	0..15 (in bytes)
Unspecified Data	*Variable*	
}		

8.2.1.2.2 CID Switch IE

In the DS-MAP, a BS may transmit DIUC=15 with the CID-Switch_IE() to toggle the inclusion of the CID parameter in DS-MAP allocations. The DS-MAP shall begin in the mode where CIDs are not included. The first appearance of the CID-Switch_IE() shall toggle the DS-MAP mode to include CIDs. Any subsequent appearance of the CID-Switch_IE() shall toggle the DS-MAP CID inclusion mode.

Table 36 – CID switch IE format

Syntax	Size	Notes
CID_Switch_IE() {		
Extended DIUC	4 bits	0x04
Length	4 bits	0
}		

8.3 Upstream Channel Descriptor (UCD)

The format of a UCD message is shown in Table 37. This message shall be transmitted by the BS at a periodic interval (Table 226) to define the characteristics of a upstream physical channel.

Table 37 – Message format

Syntax	Size	Notes

UCD_Message_Format() {		
Management Message Type = 2	8 bits	
Configuration Change Count	8 bits	Incremented by one (modulo 256) by the BS whenever any of the values of this channel descriptor change. If the value of this count in a subsequent UCD remains the same, the CPE can quickly decide that the remaining fields have not changed and may be able to disregard the remainder of the message. This value is also referenced from the US-MAP messages.
Initialization Backoff Start	8 bits	Initial backoff window size for initial ranging contention, expressed as a power of 2. Values of *n* range 0–15 (the highest order bits shall be unused and set to 0).
Initialization Backoff End	8 bits	Final backoff window size for initial ranging contention, expressed as a power of 2. Values of *n* range 0–15 (the highest order bits shall be unused and set to 0).
Request Backoff Start	8 bits	Initial backoff window size for contention BW requests, expressed as a power of 2. Values of *n* range 0–15 (the highest order bits shall be unused and set to 0).
Request Backoff End	8 bits	Final backoff window size for contention BW requests, expressed as a power of 2. Values of *n* range 0–15 (the highest order bits shall be unused and set to 0).
Coexistence Backoff Start	8 bits	Initial backoff window size for coexistence beacon transmission, expressed as a power of 2. Values of *n* range 0–15 (the highest order bits shall be unused and set to 0).
Coexistence Backoff End	8 bits	Final backoff window size for coexistence beacon transmission, expressed as a power of 2. Values of *n* range 0–15 (the highest order bits shall be unused and set to 0).
UCS Notification Backoff Start	8 bits	Initial backoff window size used by CPEs to contend to send UCS notifications to the BS. This is expressed as a power of 2. Values of *n* range 0–15 (the highest order bits shall be unused and set to 0).
UCS Notification Backoff End	8 bits	Final backoff window size used by CPEs to contend to send UCS notifications to the BS. This is expressed as a power of 2. Values of *n* range 0–15 (the highest order bits shall be unused and set to 0).
Information Elements (IEs) for the overall channel	*Variable*	8.3.1
Begin PHY Specific Section {		
for (i = 1; i ≤ n; i++) {		
Upstream_Burst_Profile	*Variable*	PHY specific (Table 40)
}		
}		
}		

8.3.1 Channel IEs

Common channel encodings are provided in Table 38.

Table 38 – UCD channel information elements

Name	Element ID (1 byte)	Length (bytes)	Value
Upstream_Burst_Profile	1		Value reserved for the burst profile (see Table 40)
Contention-based reservation timeout	2	1	Number of US-MAPs to receive before contention-based reservation is attempted again for the same connection
Bandwidth request	3	2	Size (in units of slots) of PHY payload that a CPE may use

opportunity size			to format and transmit a bandwidth request message in a contention request opportunity. The value includes all PHY overhead as well as allowance for the MAC data the message may hold.
Ranging request opportunity size	4	2	Size (in units of slots) of PHY bursts that a CPE may use to transmit a RNG-REQ message in a contention ranging request opportunity. The value includes all PHY overhead as well as the maximum CPE/BS round trip propagation delay
Frequency	5	4	Upstream center frequency (kHz)

8.3.1.1 Additional Channel IEs

In addition, specific channel encodings are shown in Table 39.

Table 39 – UCD channel information elements

Name	Element ID (1 byte)	Length (bytes)	Value
Initial ranging codes	150	1	Number of initial ranging CDMA codes. Possible values are 0–255.
Periodic ranging codes	151	1	Number of periodic ranging CDMA codes. Possible values are 0–255.
Bandwidth request codes	152	1	Number of bandwidth request CDMA codes. Possible values are 0–255.
Incumbent notification codes	153	1	Number of incumbent notification CDMA codes. Possible values are 0–255.
Periodic ranging backoff start	154	1	Initial backoff window size for periodic ranging contention, expressed as a power of 2. Range: 0–15 (the highest order bits shall be unused and set to 0).
Periodic ranging backoff end	155	1	Final backoff window size for periodic ranging contention, expressed as a power of 2. Range: 0–15 (the highest order bits shall be unused and set to 0).
Start of ranging codes group	156	1	Indicates the starting number, S, of the group of codes used for this upstream. All the ranging codes used on this upstream will be between S and ((S+N+M+L) mod 256). Where: • N is the number of initial-ranging codes • M is the number of periodic-ranging codes • L is the number of bandwidth-request codes The range of values is $0 \leq S \leq 255$.
Permutation base	157	1	Determines the US_IDcell parameter for the subcarrier permutation to be used on this upstream channel.
US allocated subchannels bitmap	158	9	This is a bitmap describing the subchannels allocated to the segment in the US, when using the upstream PUSC permutation. The LSB of the first byte shall correspond to subchannel 0. For any bit that is not set, the corresponding subchannel shall not be used by the CPE on that segment.
Optional permutation US Allocated subchannels bitmap	159	13	This is a bitmap describing the subchannels allocated to the segment in the US, when using the upstream optional PUSC permutation (see PHY spec). The LSB of the first byte shall correspond to subchannel 0. For any bit that is not set, the corresponding subchannel shall not be used by the CPE on that segment.
Band AMC Allocation Threshold	160	1	dB unit
Band AMC Release Threshold	161	1	dB unit
Band AMC Allocation Timer	162	1	Frame unit

Band AMC Release Timer	163	1	Frame unit
Band Status Reporting MAX Period	164	1	Frame unit
Band AMC Retry Timer	165	1	Frame unit
Safety Channel Allocation Threshold	166	1	dB unit
Safety Channel Release Threshold	167	1	dB unit
Safety Channel Allocation Timer	168	1	Frame unit
Safety Channel Release Timer	169	1	Frame unit
Bin Status Reporting MAX Period	170	1	Frame unit
Safety Channel Retry Timer	171	1	Frame unit
H-ARQ ACK delay for US burst	172	1	1 = one frame offset 2 = two frames offset 3 = three frames offset
CQICH Band AMC-Transition Delay	173	1	Frame unit

8.3.2 Upstream Burst Profile

Table 40 – Upstream burst profile format

Syntax	Size	Notes
Upstream_Burst_Profile_Format() {		
Type = 1	8 bits	
Length	8 bits	
Reserved	4 bits	
UIUC	4 bits	Table 44
Information Elements (IEs)	*Variable*	Table 41
}		

Table 41 – Information elements

Name	Element ID (1 byte)	Length (bytes)	Value
FEC code type and modulation type	150	1	Combination of: • Spreading • (0ffset)QPSK;(0ffset)16-QAM; (0ffset)64-QAM; • Coding rates : ½;2/3; ¾ • RS+CC/CC; CTC codes Detailed specification TBD.
Ranging data ration	151	1	Reducing factor, in units of 1 dB, between the power used for this burst and the power that should be used for CDMA Ranging.
Normalized C/N override	152	5	This is a list of numbers, where each number is encoded by one nibble, and interpreted as a signed integer. The nibbles are defined in the PHY spec. The number encoded by each nibble represents the difference in normalized C/N relative to the previous one.

8.4 Upstream Map (US-MAP)

The format of a US-MAP message is shown in Table 42. The US-MAP message defines the access to the upstream channel.

The CID field carried in the MAC header of the PDU where this message is transmitted represents the assignment of the IE to either a unicast, multicast, or broadcast address, or the padding CID. When specifically addressed to allocate a bandwidth grant, the CID shall be the Basic CID of the CPE. A UIUC shall be used to define the type of upstream access and the upstream burst profile associated with that access. An Upstream_Burst_Profile shall be included in the UCD for each UIUC to be used in the US-MAP.

Table 42 – Message format

Syntax	Size	Notes
US-MAP_Message_Format() {		
Management Message Type = 3	8 bits	
Upstream Channel ID	8 bits	The identifier of the upstream channel to which this message refers.
UCD Count	8 bits	Matches the value of the Configuration Change Count of the UCD, which describes the upstream burst profiles that apply to this map.
Allocation Start Time	8 bits // it was 32 bits	Effective start time (in MAC slots) of the upstream allocation defined by the US-MAP
Begin PHY Specific Section {		
for (i = 1; i ≤ n; i++) {		
US-MAP_IE()	*Variable*	PHY specific (8.4.1) Define upstream bandwidth allocations. Each US-MAP message shall contain at least one IE that marks the end of the last allocated burst.
}		
}		
If (!byte_boundary)		
Padding Nibble	4 bits	
}		

8.4.1 US-MAP IE

The US-MAP IE is shown in Table 43, and is used to define the upstream bandwidth allocations. The first US-MAP IE shall start at the lowest numbered non-allocated channel on the first non-allocated symbol defined by the allocation start time field of the US-MAP message that is not allocated with 0 ≤ UIUC ≤ 5. These IEs shall represent the number of slots provided for the allocation. Each allocation IE shall start immediately following the previous allocation and shall advance in the time domain. If the end of the US frame has been reached, the allocation shall continue at the next channel at first symbol (defined by the allocation start time field) that is not allocated with 0 ≤ UIUC ≤ 5. A Burst Descriptor shall be specified in the UCD for each UIUC to be used in the US-MAP.

Table 43 – US-MAP information element

Syntax	Size	Notes
US-MAP_IE() {		
CID	16 bits	Table 32
UIUC	4 bits	Table 44
If ((UIUC ≥ 0) && (UIUC ≤ 5)) {		

Channel Offset	7 bits	
Number of Channels	7 bits	
Slot Offset	8 bits	
Number of Slots	8 bits	
Reserved	2 bits	
} else if (UIUC == 15) {		
US_Extended_IE()	*Variable*	32966668.1357040.58.4.1.2
} else if (UIUC == 14) {		
CDMA_Allocation_IE()	32 bits	
} else {		
Duration	8 bits	In number of MAC slots
MDP	1 bit	Measurement Data Preferred Used by the BS to indicate to the CPE that this upstream allocation is to be preferably used by the CPE for the specific purpose of reporting back any measurement data. The measurement data to be reported is in connection to the specified Transaction ID. In case the CPE does not have anything to report, it can use this allocation for any other data. This is useful, for example, after a quiet period. 0 = Measurement data not required (default) 1 = Measurement data preferred
MRT	1 bit	Measurement Report Type In case MDP == 1, this field indicates which type of report the BS wants the CPE to send back. 0 = Detailed (see Table 162 thru Table 170) 1 = Consolidated (see Table 171)
CMRP	1 bit	Channel Management Response Preferred Used by the BS to indicate to the CPE that this upstream allocation is to be used for confirming or not the receipt of the channel management command with the Transaction ID specified. 0 = Channel management response not required (default) 1 = Channel management response required
Transaction ID	16 bits	The transaction ID of the corresponding channel management or measurement request for which a response/report is required by the BS.
Preamble Present	1 bit	Determines whether CPE shall send preamble before any data transmission 0 = Preamble shall not be used (default) 1 = Preamble shall be used
}		
If (!byte_boundary)		
Padding Nibble	4 bits	
}		

8.4.1.1 UIUC Allocations

Table 44 specifies the UIUC incorporated into CMAC. In particular, the self-coexistence UIUCs (in both modes) have the same applicability to their DIUC counterpart (see 217.0.58.2.1.1).

Table 44 – UIUC values

UIUC	Usage
0	Self-Coexistence (Active Mode)
1	Self-Coexistence (Passive Mode)
2	UCS Notification
3	BW Request, Ranging
4	CDMA UCS Notification
5	CDMA BW Request, CDMA Ranging
6-12	Burst Profiles
13	End of Map
14	CDMA Allocation IE
15	Extended UIUC

8.4.1.2 US-MAP Extended UIUC IE

Table 45 – US-MAP extended IE general format

Syntax	Size	Notes
US_Extended_IE() {		
Extended UIUC	4 bits	
Length	4 bits	
Unspecified Data	*Variable*	
}		

8.4.1.2.1 US-MAP Power Control IE

Table 46 – US-MAP power control IE format

Syntax	Size	Notes
Power_Control_IE() {		
Extended UIUC	4 bits	Fast power control = 0x00
Length	4 bits	
Power Control	8 bits	Signed integer which expresses the change in power level (in 0.25 dB units) that the CPE should apply to correct its current transmission power
}		

8.4.1.2.2 US-MAP Dummy IE

Table 47 – US-MAP dummy IE format

Syntax	Size	Notes
Dummy_IE() {		
Extended UIUC	4 bits	0x03
Length	4 bits	0..15 (in bytes)
Unspecified Data	*Variable*	
}		

8.4.1.3 CDMA Allocation IE

Table 48 – CDMA allocation IE format

Syntax	Size	Notes
CDMA_Allocation_IE() {		
Duration	6 bits	Indicates the duration, in units of MAC slots, of the allocation.
Repetition Coding Indication	2 bits	Indicates the repetition code used inside the allocated burst. 0b00 – No repetition coding 0b01 – Repetition coding of 2 used 0b10 – Repetition coding of 4 used 0b11 – Repetition coding of 6 used
Code	8 bits	Indicates the Code sent by the CPE.
Symbol	8 bits	Indicates the PHY symbol used by the CPE.
Channel	7 bits	Identifies the channel used by the CPE to send the Code.
Usage	1 bit	This field can mean two different things depending upon the type of Code used. • If ((Code = Ranging) or (Code = BW Request)) This field indicates whether the CPE shall include a Bandwidth (BW) Request in the allocation. 1 = yes; 0 = no • If (Code = Incumbent) This field indicates whether the CPE shall transmit only the MAC header with the notification. 1 = yes; 0 = no
}		

8.5 RNG-REQ

The format of an RNG-REQ message is shown in Table 49. An RNG-REQ shall be transmitted by the CPE at initialization and periodically to determine network delay and to request power and/or downstream burst profile change. The RNG-REQ message may be sent in Initial Ranging and data grant intervals.

The CID field carried in the MAC header of the PDU where this message is transmitted shall assume the following values when sent in Initial Ranging interval:

- Initial ranging CID if the CPE is attempting to join the network.
- Initial ranging CID if the CPE has not yet registered and is changing downstream (or both downstream and upstream) channels.
- In all other cases, the Basic CID is used as soon as one is assigned in the RNG-RSP message.

If sent in a data grant interval, the CID is always equal to the Basic CID.

Table 49 – Message format

Syntax	Size	Notes
RNG-REQ_Message_Format() {		
Management Message Type = 4	8 bits	
Downstream Channel ID	8 bits	The identifier of the downstream channel on which the CPE received the UCD describing the upstream on which this ranging request message is to be

		transmitted.
Information Elements (IEs)	*Variable*	Table 50
}		

Table 50 – Information elements

Name	Element ID (1 byte)	Length (bytes)	Value
Downstream burst profile	1	1	
CPE MAC address	2	6	
Ranging anomalies	3	1	A parameter indicating a potential error condition detected by the CPE during the ranging process. Setting the bit associated with a specific condition indicates that the condition exists at the CPE. Bit #0 — CPE already at maximum power. Bit #1 — CPE already at minimum power. Bit #2 — Sum of commanded timing adjustments is too large.
Number of Simultaneous Channels Supported	4	1	Indicates how many contiguous physical channels the CPE has capability of grouping. Shall be at least 3 in the basic mode.
AAS broadcast capability	5	1	0 = CPE can receive broadcast messages 1 = CPE cannot receive broadcast messages

8.6 RNG-RSP

The format of an RNG-RSP message is shown in Table 51. An RNG-RSP shall be transmitted by the BS in response to a received RNG-REQ. In addition, it may also be transmitted asynchronously to send corrections based on measurements that have been made on other received data or MAC messages. As a result, the CPE shall be prepared to receive an RNG-RSP at any time, not just following an RNG-REQ transmission.

Table 51 – Message format

Syntax	Size	Notes
RNG-RSP_Message_Format() {		
Management Message Type = 5	8 bits	
Upstream Channel ID	8 bits	The identifier of the upstream channel on which the BS received the RNG-REQ to which this response refers.
Information Elements (IEs)	*Variable*	Table 52
}		

Table 52 – Information elements

Name	Element ID (1 byte)	Length (bytes)	Value
Timing adjust	1	4	In units of system clock
Power level adjust	2	1	Signed in units of dBm
Offset frequency adjust	3	2	In Hertz
Ranging status	4	1	
Downstream operational burst profile	5	2	
CPE MAC address	6	6	A required parameter when the CID in the

			MAC header is the Initial Ranging CID.
Basic CID	7	2	A required parameter if the RNG-RSP message is being sent on the Initial Ranging CID in response to a RNG-REQ message that was sent on the Initial Ranging CID.
Primary Management CID	8	2	A required parameter if the RNG-RSP message is being sent on the Initial Ranging CID in response to a RNG-REQ message that was sent on the Initial Ranging CID.
AAS broadcast permission	9	1	0 = CPE may issue contention-based Bandwidth Request permission 1 = CPE shall not issue contention-based Bandwidth Request

8.7 REG-REQ/RSP

CPEs shall register with a BS before receiving or being provided any type of service. In the following subsections we present the registration process incorporated in CMAC, as well as a series of IEs that may be carried by these messages.

8.7.1 REG-REQ

The format of an REG-REQ message is shown in Table 53. This message shall be transmitted by CPEs at initialization phase.

The CID field carried in the MAC header of the PDU where this message is transmitted shall be the primary management CID for this CPE, which is assigned during the RNG-RSP message.

Table 53 –Message format

Syntax	Size	Notes
REG-REQ_Message_Format() {		
Management Message Type = 6	8 bits	
Information Elements (IEs)	*Variable*	8.7.3
}		

8.7.2 REG-RSP

The format of an REG-RSP message is shown in Table 54. This message shall be transmitted by the BS in response to a REG-REQ.

The CID field carried in the MAC header of the PDU where this message is transmitted shall be the primary management CID of the CPE for which this message is intended.

Table 54 –Message format

Syntax	Size	Notes
REG-RSP_Message_Format() {		
Management Message Type = 7	8 bits	
Response	8 bits	0: OK 1: Failure (e.g., authentication)
Information Elements (IEs)	*Variable*	8.7.3

}		

8.7.3 Information Elements

REG-REQ and REG-RSP management messages may carry a number of IEs that support the registration process. These IEs are described in detail in the following subsections.

8.7.3.1 ARQ Parameters

Table 55 – Information elements

Element ID	Length (bytes)	Value	Scope
1	*Variable*	Compound	REG-REQ, REG-RSP

8.7.3.2 CPE Management Support

Table 56 – Information elements

Element ID	Length (bytes)	Value	Scope
2	1	0: no secondary management connection 1: secondary management connection	REG-REQ, REG-RSP

8.7.3.3 IP Management Mode

Table 57 – Information elements

Element ID	Length (bytes)	Value	Scope
3	1	0: unmanaged mode 1: IP managed mode	REG-REQ, REG-RSP

8.7.3.4 IP Version

Table 58 – Information elements

Element ID	Length (bytes)	Value	Scope
4	1	Bit #0: 4 (default) Bit #1: 6	REG-REQ, REG-RSP

8.7.3.5 Secondary Management CID

Table 59 – Information elements

Element ID	Length (bytes)	Value	Scope
5	2	CID	REG-RSP

8.7.3.6 Number of Upstream CID Supported

Table 60 – Information elements

Element ID	Length (bytes)	Value	Scope
6	2	Number of upstream CID supported by CPE	REG-REQ, REG-RSP

8.7.3.7 CPE Capability

Through the registration process a BS shall become aware of the capabilities of the registering CPEs. The following subsections describe the IEs that convey the CPE capability information to the BS.

8.7.3.7.1 ARQ Support

Table 61 – Information elements

Element ID	Length (bytes)	Value	Scope
10	1	0: No ARQ support 1: ARQ supported	REG-REQ, REG-RSP

8.7.3.7.2 DSx Flow Control

Table 62 – Information elements

Element ID	Length (bytes)	Value	Scope
11	1	0: no limit	REG-REQ, REG-RSP

8.7.3.7.3 MCA Flow Control

Table 63 – Information elements

Element ID	Length (bytes)	Value	Scope
12	1	0: no limit	REG-REQ, REG-RSP

8.7.3.7.4 Maximum Number of Multicast Groups Supported

Table 64 – Information elements

Element ID	Length (bytes)	Value	Scope
13	1	0-255	REG-REQ, REG-RSP

8.7.3.7.5 PKM Flow Control

Table 65 – Information elements

Element ID	Length (bytes)	Value	Scope
14	1	0: no limit	REG-REQ, REG-RSP

8.7.3.7.6 Authorization Policy Support

Table 66 – Information elements

Element ID	Length	Value	Scope

	(bytes)		
15	1	Bit #0: 802.22 privacy supported	REG-REQ, REG-RSP

8.7.3.7.7 Maximum Number of Supported Security Associations

Table 67 – Information elements

Element ID	Length (bytes)	Value	Scope
16	1	0-255	REG-REQ, REG-RSP

8.7.3.7.8 Receiver Sensitivity

Table 68 – Information elements

Element ID	Length (bytes)	Value	Scope
17	1	Sensitivity in dB	REG-REQ, REG-RSP

8.7.3.7.9 Measurement Support

Table 69 – Information elements

Syntax	Size	Notes
MS_IE() {		
Element ID	8 bits	18
Length	16 bits	n*3 (in bytes)
for (i = 1; i ≤ n; i++) {		n = the number of system profiles
System_Profile	8 bits	Table 70
Dedicated Interface	2 bits	00: No dedicated radio interface available 01: Dedicated radio interface available 10-11: *reserved*
Sensitivity Threshold	14 bits	Only applies for incumbent profiles. Signed number which indicates the threshold (in dBm) (to be) used by CPEs for detection.
}		
}		

Table 70 – System profiles

System_Profile	Description
0	802.22
1	ATSC
2	NTSC
3	Part 74
4	DVB
5-255	*Reserved*

8.7.3.7.10 Antenna Gain

Table 71 – Information elements

Element ID	Length (bytes)	Value	Scope
19	1	Antenna gain in dB	REG-REQ, REG-RSP

8.8 Dynamic Service Messages (DSx-REQ/RSP/ACK)

To manage the various traffic flows between CPEs and the BS, the MAC protocol shall have the capability to dynamically manage the addition, deletion, and change of service flows. Therefore, in this section the messages related to dynamic service management are presented.

8.8.1 DSA-REQ

The format of an DSA-REQ message is shown in Table 72. This message is sent either by a CPE or BS and as to create a new service flow, and shall not contain parameters for more than one service flow.

The CID field carried in the MAC header of the PDU where this message is transmitted shall be the primary management CID of the CPE.

Table 72 – Message format

Syntax	Size	Notes
DSA-REQ_Message_Format() {		
Management Message Type = 11	8 bits	
Transaction ID	16 bits	Unique identifier for this transaction assigned by the sender.
Information Elements (IEs)	*Variable*	7.4
}		

8.8.2 DSA-RSP

Table 73 – Message format

Syntax	Size	Notes
DSA-RSP_Message_Format() {		
Management Message Type = 12	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
Information Elements (IEs)	*Variable*	7.4
}		

8.8.3 DSA-ACK

Table 74 – Message format

Syntax	Size	Notes
DSA-ACK_Message_Format() {		
Management Message Type = 13	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
Information Elements (IEs)	*Variable*	7.4
}		

8.8.4 DSC-REQ

Table 75 – Message format

Syntax	Size	Notes
DSC-REQ_Message_Format() {		
Management Message Type = 14	8 bits	
Transaction ID	16 bits	
Information Elements (IEs)	*Variable*	7.4
}		

8.8.5 DSC-RSP

Table 76 – Message format

Syntax	Size	Notes
DSC-RSP_Message_Format() {		
Management Message Type = 15	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
Information Elements (IEs)	*Variable*	7.4
}		

8.8.6 DSC-ACK

Table 77 – Message format

Syntax	Size	Notes
DSC-ACK_Message_Format() {		
Management Message Type = 16	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
Information Elements (IEs)	*Variable*	7.4
}		

8.8.7 DSD-REQ

Table 78 – Message format

Syntax	Size	Notes
DSD-REQ_Message_Format() {		
Management Message Type = 17	8 bits	
Transaction ID	16 bits	
Service Flow ID	32 bits	
Information Elements (IEs)	*Variable*	7.4
}		

8.8.8 DSD-RSP

Table 79 – Message format

Syntax	Size	Notes
DSD-RSP_Message_Format() {		
Management Message Type = 18	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
Service Flow ID	32 bits	
Information Elements (IEs)	*Variable*	7.4
}		

8.8.9 DSx-RVD

Table 80 – Message format

Syntax	Size	Notes
DSX-RVD_Message_Format() {		The DSX-RVD message shall be generated by the BS in response to a CPE-initiated DSx-REQ to inform the CPE that the BS has received the DSx-REQ message in a timelier manner than provided by the DSx-RSP message, which shall be transmitted only after the DSx-REQ is authenticated.
Management Message Type = 30	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
}		

8.8.10 Service Flow Encodings

Table 81 – Service flow encodings

Type	Parameter
1	Service Flow Identifier
2	CID
3	Service Class Name
4	*reserved*
5	QoS Parameter Set Type
6	Traffic Priority
7	Maximum Sustained Traffic Rate
8	Maximum Traffic Burst
9	Maximum Reserved Traffic Rate
10	Minimum Tolerable Traffic Rate
11	Service Flow Scheduling Type
12	Request/Transmission Policy
13	Tolerated Jitter
14	Maximum Latency
15	Fixed-length versus Variable-length SDU Indicator
16	SDU Size
17	Target SAID
18	ARQ Enable
19	ARQ_WINDOW_SIZE
20	ARQ_RETRY_TIMEOUT – Transmitter Delay
21	ARQ_RETRY_TIMEOUT – Receiver Delay
22	ARQ_BLOCK_LIFETIME
23	ARQ_SYNC_LOSS
24	ARQ_DELIVER_IN_ONRDER

25	ARQ_PURGE_TIMEOUT
26	ARQ_BLOCK_SIZE
27	*Reserved*
28	CS Specification
29	Maximum Tolerable Packet Loss Rate
143	Vendor-specific QoS Parameter
99-107	Convergence Sublayer Types

Table 82 – Confirmation Code (CC) values

CC	Status
0	OK/success
1	reject-other
2	reject-unrecognized-configuration-setting
3	reject-temporary / reject-resource
4	reject-permanent / reject-admin
5	reject-not-owner
6	reject-service-flow-not-found
7	reject-service-flow-exists
8	reject-required-parameter-not-present
9	reject-header-suppression
10	reject-unknown-transaction-id
11	reject-authentication-failure
12	reject-add-aborted
13	reject-exceed-dynamic-service-limit
14	reject-not-authorized-for-the-request-SAID
15	reject-fail-to-establish-the-requested-SA
16	reject-not-supported-parameter
17	reject-not-supported-parameter-value

8.8.10.1 SFID

Table 83 – SFID definition

Element ID	Length (bytes)	Value	Scope
[145/146].1	4	1-4 294 967 295	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.2 CID

Table 84 – CID definition

Element ID	Length (bytes)	Value	Scope
[145/146].2	2	CID	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.3 Service Class Name

Table 85 – Service class name definition

Element ID	Length (bytes)	Value	Scope

[145/146].3	2 to 128	Null-terminated string of ASCII characters. The length of the string, including null-terminator, may not exceed 128 bytes	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.4 QoS Parameter Set Type

Table 86 – QoS parameter set type definition

Element ID	Length (bytes)	Value	Scope
[145/146].5	1	Bit 0: Provisioned Set Bit 1: Admitted Set Bit 2: Active Set Bits 3-7: *Reserved*	DSx-REQ, DSx-RSP, DSx-ACK

Table 87 – Values used in Dynamic Service messages

Value	Messages
001	Apply to Provisioned set only
011	Apply to Provisioned and Admitted sets, and perform admission control
101	Apply to Provisioned and Admitted sets, perform admission control, and activate this service flow
111	Apply to Provisioned, Admitted, and Active sets, perform admission control, and activate this service flow
000	Set Active and Admitted sets to Null
010	Perform admission control and apply admitted set
100	Check against Admitted set in separate service flow encoding, perform admission control if needed, active this service flow, and apply to Active set
110	Perform admission control and activate this service flow, apply parameters to both Admitted and Active sets

8.8.10.5 Traffic Priority

Table 88 – Traffic priority definition

Element ID	Length (bytes)	Value	Scope
[145/146].6	1	0 to 7 – Higher numbers indicate higher priority Default 0	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.6 Maximum Sustained Traffic Rate

Table 89 – Maximum sustained traffic rate definition

Element ID	Length (bytes)	Value	Scope
[145/146].7	4	Rate (in bits per second)	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.7 Maximum Traffic Burst

Table 90 – Maximum traffic burst definition

Element ID	Length (bytes)	Value	Scope
[145/146].8	4	Burst size (bytes)	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.8 Minimum Reserved Traffic Rate

Table 91 – Minimum reserved traffic rate definition

Element ID	Length (bytes)	Value	Scope
[145/146].9	4	Rate (in bits per second)	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.9 Minimum Tolerable Traffic Rate

Table 92 – Minimum tolerable traffic rate definition

Element ID	Length (bytes)	Value	Scope
[145/146].10	4	Rate (in bits per second)	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.10 Vendor Specific QoS Parameters

Table 93 – Vendor specific QoS parameters definition

Element ID	Length (bytes)	Value	Scope
[145/146].143	*Variable*	Compound (7.6)	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.11 Service Flow Scheduling Type

Table 94 – Service flow scheduling type definition

Element ID	Length (bytes)	Value	Scope
[145/146].11	1	0: *Reserved* 1: for Undefined (BS implementation dependent) 2: for BE (Default) 3: for nrtPS 4: for rtPS 5: *Reserved* 6: for UGS 7-255: *Reserved*	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.12 Request/Transmission Policy

Table 95 – Request/transmission policy definition

Element ID	Length	Value	Scope

	(bytes)		
[145/146].12	1	Bit 0: Service flow shall not use broadcast bandwidth request opportunities (Upstream only) Bit 1: *Reserved* (shall be set to zero) Bit 2: The service flow shall not piggyback requests with data (Upstream only) Bit 3: The service flow shall not fragment data Bit 4: The service flow shall not suppress payload headers (CS parameters) Bit 5: The service flow shall not pack multiple SDUs (or fragments) into single MAC PDUs Bit 6: *Reserved* (shall be set to zero) Bit 7: *Reserved* (shall be set to zero)	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.13 Tolerated Jitter

Table 96 – Tolerated jitter definition

Element ID	Length (bytes)	Value	Scope
[145/146].13	4	ms	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.14 Maximum Latency

Table 97 – Maximum latency definition

Element ID	Length (bytes)	Value	Scope
[145/146].14	4	ms	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.15 Fixed-length versus Variable-length SDU Indicator

Table 98 – Fixed-length versus variable length SDU indicator definition

Element ID	Length (bytes)	Value	Scope
[145/146].15	1	0: variable-length SDU (default) 1: fixed-length SDU	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.16 SDU Size

Table 99 –SDU size definition

Element ID	Length (bytes)	Value	Scope
[145/146].16	1	Number of bytes Default = 49	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.17 Target SAID

Table 100 –Target SAID definition

Element ID	Length (bytes)	Value	Scope
[145/146].17	2	SAID onto which SF is mapped.	DSx-REQ, DSx-RSP

8.8.10.18 Maximum Tolerable Packet Loss Rate

Table 101 – Maximum tolerable packet loss rate

Element ID	Length (bytes)	Value	Scope
[145/146].29	1	Valid range: 0-100 Maximum percentage of packet loss rate tolerated before a flow is dropped.	DSx-REQ, DSx-RSP, DSx-ACK

8.8.10.19 ARQ IEs for ARQ-enabled Connections

8.8.10.19.1 ARQ Enable

Table 102 –ARQ enable definition

Element ID	Length (bytes)	Value	Scope
[145/146].18 1.18	1	0 = ARQ not requested/accepted 1 = ARQ requested/accepted	DSA-REQ, DSA-RSP, REG-REQ, REQ-RSP

8.8.10.19.2 ARQ_WINDOW_SIZE

Table 103 –ARQ_WINDOW_SIZE definition

Element ID	Length (bytes)	Value	Scope
[145/146].19 1.19	2	> 0 and ≤ (ARQ_BSN_MODULUS/2)	DSx-REQ, DSx-RSP, REG-REQ, REQ-RSP

8.8.10.19.3 ARQ_RETRY_TIMEOUT

Table 104 –ARQ_RETRY_TIMEOUT definition

Element ID	Length (bytes)	Value	Scope
[145/146].20 1.20	2	TRANSMITTER_DELAY 0-655350 (10μ granularity)	DSA-REQ, DSA-RSP, REG-REQ, REQ-RSP
[145/146].21 1.21	2	RECEIVER_DELAY 0-655350 (10μ granularity)	DSA-REQ, DSA-RSP, REG-REQ, REQ-RSP

8.8.10.19.4 ARQ_BLOCK_LIFETIME

Table 105 –ARQ_BLOCK_LIFETIME definition

Element ID	Length (bytes)	Value	Scope

[145/146].22 1.22	2	0 = Infinite 1-655350 (10μ granularity)	DSA-REQ, DSA-RSP, REG-REQ, REQ-RSP

8.8.10.19.5 ARQ_SYNC_LOSS_TIMEOUT

Table 106 –ARQ_SYNC_LOSS_TIMEOUT definition

Element ID	Length (bytes)	Value	Scope
[145/146].23 1.23	2	0 = Infinite 1-655350 (10μ granularity)	DSA-REQ, DSA-RSP, REG-REQ, REQ-RSP

8.8.10.19.6 ARQ_DELIVER_IN_ORDER

Table 107 –ARQ_DELIVER_IN_ORDER definition

Element ID	Length (bytes)	Value	Scope
[145/146].24 1.24	1	0 = Order o delivery is not preserved 1 = Order of delivery is preserved	DSA-REQ, DSA-RSP, REG-REQ, REQ-RSP

8.8.10.19.7 ARQ_RX_PURGE_TIMEOUT

Table 108 –ARQ_RX_PURGE_TIMEOUT definition

Element ID	Length (bytes)	Value	Scope
[145/146].25 1.25	2	0 = Infinite 1-655350 (10μ granularity)	DSA-REQ, DSA-RSP, REG-REQ, REQ-RSP

8.8.10.19.8 ARQ_BLOCK_SIZE

Table 109 –ARQ_BLOCK_SIZE definition

Element ID	Length (bytes)	Value	Scope
[145/146].26 1.26	2	0 = *Reserved* 1-2040 = Desired/Agreed size in bytes 2041-65535 = *Reserved*	DSA-REQ, DSA-RSP, REG-REQ, REQ-RSP

8.9 CPE Fast Power Control (CPE-FPC)

In addition to controlling the power through the ranging procedure (see 17), the BS may also employ the CPE-FPC message (see Table 110) to adjust the power level of multiple CPEs simultaneously. Since the CPE-FPC message may apply to multiple CPEs, this message is transmitted with the broadcast CID. Upon receiving a CPE-FCP message from the BS, the CPE shall adjust its power within a pre-determined time, which is a system parameter.

Table 110 – Message format

Syntax	Size	Notes
CPE-FPC_Message_Format() {		
Management Message Type = 36	8 bits	
Number of CPEs	8 bits	Specifies the length of the

		message
for (i = 0; i < Number of CPEs; i++) {		
Basic CID	16 bits	The Basic CID which identifies a particular CPE
Power Adjustment	8 bits	Signed integer which indicates the Power Adjustment (in units of 0.5 dB) that the CPE shall perform to its current power level
}		
}		

8.10 Multicast Assignment Request (MCA-REQ)

Table 111 – Message format

Syntax	Size	Notes
MCA-REQ_Message_Format() {		
Management Message Type = 21	8 bits	
Transaction ID	16 bits	
Information Elements (IEs)	*Variable*	Table 112
}		

Table 112 – Information elements

Name	Element ID (1 byte)	Length (bytes)	Value
Multicast CID	1	2	Either transport or multicast management CID (see Table 227)
Assignment	2	1	0x00 = Leave multicast group 0x01 = Join multicast group
Multicast group type	3	1	0 = regular, default 1 = *reserved*
Periodic allocation parameters	4	4	Valid only for transport CIDs. Byte 0 (LS byte) = m Byte 1 = k Byte 2 = n Byte 3 = *reserved* Multicast group gets an upstream multicast allocation at the end of the frame #N if: N mod k = m; size of allocation is n.
Operation	5	1	0 = allocate 1 = deallocate
Reserved	6-255		

8.11 Multicast Assignment Response (MCA-RSP)

Table 113 – Message format

Syntax	Size	Notes
MCA-RSP_Message_Format() {		
Management Message Type = 22	8 bits	
Transaction ID	16 bits	

Confirmation Code	8 bits	Table 82
}		

8.12 Downstream Burst Profile Change Request (DBPC-REQ)

Table 114 – Message format

Syntax	Size	Notes
DBPC-REQ_Message_Format() {		
Management Message Type = 23	8 bits	
reserved	4 bits	
DIUC	4 bits	Table 33
Configuration Change Count	8 bits	
}		

8.13 Downstream Burst Profile Change Response (DBPC-RSP)

Table 115 – Message format

Syntax	Size	Notes
DBPC-RSP_Message_Format() {		
Management Message Type = 24	8 bits	
reserved	4 bits	
DIUC	4 bits	Table 33
Configuration Change Count	8 bits	
}		

8.14 Reset Command (RES-CMD)

Table 116 – Message format

Syntax	Size	Notes
RES-CMD_Message_Format() {		
Management Message Type = 25	8 bits	
Information Elements (IEs)	*Variable*	The IE shall include the HMAC tuple (7.1)
}		

8.15 CPE Basic Capability Request/Response (CBC-REQ/RSP)

8.15.1 CBC-REQ

Table 117 – Message format

Syntax	Size	Notes
CBC-REQ_Message_Format() {		
Management Message Type = 26	8 bits	
Information Elements (IEs)	*Variable*	See 8.7.3.7, 8.15.3.1, and 8.15.3.3
}		

8.15.2 CBC-RSP

Table 118 – Message format

Syntax	Size	Notes
CBC-RSP_Message_Format() {		
Management Message Type = 27	8 bits	
Information Elements (IEs)	*Variable*	8.15.3.1 8.15.3.3
}		

8.15.3 Information Elements

8.15.3.1 Bandwidth Allocation Support

Table 119 – Information elements

Element ID	Length (bytes)	Value	Scope
1	1	Bits 0-7: *reserved* (set to zero)	CBC-REQ, CBC-RSP

8.15.3.2 Capabilities for Construction and Transmission of MAC PDUs

Table 120 – Information elements

Element ID	Length (bytes)	Value	Scope
4	1	Bit 0: Ability to receive requests piggybacked with data Bit 1: Specifies the maximum size of FSN values used when forming MAC PDUs on non-ARQ connections 0: Only 3-bit supported 1: Only 11-bit supported Bits 2-7: *reserved* (set to zero)	CBC-REQ, CBC-RSP

8.15.3.3 Physical Parameters Supported

8.15.3.3.1 CPE Transition Gaps

Table 121 – Information elements

Element ID	Length (bytes)	Value	Scope
2	2	Bits 0-7: CPETTG (μs) Bits 8-15: CPERTG (μs) Allowed values: TDD: 0…50	CBC-REQ, CBC-RSP

8.15.3.3.2 Maximum Transmit Power

The maximum available power for QPSK, 16-QAM, and 64-QAM, 256-QAM, and O-QAM constellations. The maximum power parameters are reported in dBm and quantized in 1 dBm steps ranging from -64 dBm (encoded 0x00) to 64 dBm (encoded 0xFF). Values outside this range shall be assigned the closest extreme. CPEs that do not support a specific modulation shall report the value of 0x00 in the maximum power field.

Table 122 – Information elements

Element ID	Length (bytes)	Value	Scope
3	5	Byte 0: Maximum transmitted power for QPSK Byte 1: Maximum transmitted power for 16-QAM Byte 2: Maximum transmitted power for 64-QAM. Byte 3: Maximum transmitted power for 256-QAM. CPEs that do not support 256-QAM shall report the value 0x00. Byte 4: Maximum transmitted power for O-QAM. CPEs that do not support O-QAM shall report the value 0x00.	CBC-REQ

8.15.3.3.3 Current Transmit Power

Table 123 – Information elements

Element ID	Length (bytes)	Value	Scope
147	1	Current transmit power (7.3)	CBC-REQ

8.15.3.3.4 PHY-Specific Parameters

8.15.3.3.4.1 CPE FFT Sizes

This field indicates the FFT sizes supported by the CPE. For each FFT size, a bit value of 0 indicates "not supported" while 1 indicates "supported." Note that 2K, 4K, 6K FFT modes are mandatory to be supported. See PHY spec for further details.

Table 124 – Information elements

Element ID	Length (bytes)	Value	Scope
150	1	Bit 0: FFT-12K Bits 1-7: *reserved* (set to zero)	CBC-REQ, CBC-RSP

8.15.3.3.4.2 CPE Demodulator

This field indicates the different demodulator options supported by a CPE for downstream reception. A bit value of 0 indicates "not supported" while 1 indicates "supported."

Table 125 – Information elements

Element ID	Length (bytes)	Value	Scope
151	1	Bit 0: 256-QAM Bit 1: O-QAM Bit 2: CTC	CBC-REQ, CBC-RSP

		Bit 3: RS Bits 4-7: *reserved* (set to zero)	

8.15.3.3.4.3 CPE Modulator

This field indicates the different modulator options supported by a CPE for upstream transmission A bit value of 0 indicates "not supported" while 1 indicates "supported."

Table 126 – Information elements

Element ID	Length (bytes)	Value	Scope
152	1	Bit 0: 256-QAM Bit 1: O-QAM Bit 2: CTC Bit 3: RS Bit 4: H-ARQ Bits 5-7: *reserved* (set to zero)	CBC-REQ, CBC-RSP
153	1	The number of HARQ ACK channel	CBC-REQ, CBC-RSP

8.15.3.3.4.4 CPE Permutation Support

This field indicates the different optional permutation modes (optional PUSC, optional FUSC and AMC) supported by a CPE. A bit value of 0 indicates "not supported" while 1 indicates "supported."

Table 127 – Information elements

Element ID	Length (bytes)	Value	Scope
154	1	Bit# 0: Optional PUSC support Bit# 1: Optional FUSC support Bit# 2: AMC support Bits# 3–7: *Reserved*, shall be set to zero	CBC-REQ, CBC-RSP

8.16 De/Re-register Command (DREG-CMD)

Table 128 – Message format

Syntax	Size	Notes
DREG-CMD_Message_Format() {		
Management Message Type = 29	8 bits	
Action Code	8 bits	Table 129
Information Elements (IEs)	*Variable*	The IE shall include the HMAC tuple (7.1) as the last attribute in the message
}		

Table 129 – Action codes and actions

Action Code	Action
0x00	CPE shall leave the current channel and attempt to access another channel.
0x01	CPE shall listen to the current channel but shall not transmit until an RES-CMD message or DREG-CMD with an Action Code that allows transmission is received.

0x02	CPE shall listen to the current channel but only transmit on the Basic, Primary Management, and Secondary Management Connections.
0x03	CPE shall return to normal operation and may transmit on any of its active connections.
0x04	CPE shall terminate current Normal Operations with the BS; the BS shall transmit this action code only in response to any CPE DREG-REQ message.
0x05-0xFF	*Reserved*

8.17 CPE De-registration Request (DREG-REQ)

Table 130 – Message format

Syntax	Size	Notes
DREG-REQ_Message_Format() {		
Management Message Type = 49	8 bits	
De-registration Request Code	8 bits	• 0x00 = CPE de-registration request from BS and network • 0x01-0xFF = *reserved*
Information Elements (IEs)	*Variable*	The IE shall include the HMAC tuple (7.1) as the last attribute in the message
}		

8.18 ARQ-Feedback

Table 131 – Message format

Syntax	Size	Notes
ARQ_Feedback_Message_Format() {		
Management Message Type = 33	8 bits	
ARQ Feedback Payload	*Variable*	See 9.4.3
}		

8.19 ARQ-Discard

Table 132 – Message format

Syntax	Size	Notes
ARQ_Discard_Message_Format() {		This message is sent when the transmitter wants to skip a certain number of ARQ blocks
Management Message Type = 34	8 bits	
Connection ID	16 bits	CID to which message refers
reserved	5 bits	Shall be set to zero
BSN	11 bits	Sequence number of the last block in the transmitter window that the transmitter wants to discard
}		

8.20 ARQ-Reset

Table 133 – Message format

Syntax	Size	Notes
ARQ_Reset_Message_Format() {		The transmitter or the receiver may send this message. The message is used in a dialog to reset the parent connection's ARQ transmitter and receiver state machines.
Management Message Type = 35	8 bits	
Connection ID	16 bits	CID to which message refers
Type	2 bits	00 = Original message from Initiator 01 = Acknowledgment from Responder 10 = Confirmation from Initiator 11 = *Reserved*
reserved	6 bits	Shall be set to zero
}		

8.21 Channel Management

CMAC provides a comprehensive set of messages that allows the BS to dynamically manage the channel operations, and so support many essential features such as effective coexistence and measurements. In this section we present the channel management messages supported by CMAC.

All channel management messages possess a Transaction ID field that uniquely identifies the message in a single 802.22 network. If two or more management messages are received with the same Transaction ID, the parameters of the last message received shall override those of all previously received management messages with the same Transaction ID.

8.21.1 Channel Terminate Request (CHT-REQ)

This message (Table 134) is sent by the BS in order to terminate operation in channel, which can be due to several reasons such as protection of incumbent services. Note that termination of operation in a channel can also be implemented through the SCH by having the BS specify a different value for Channel Number and Number of Channels. However, since a superframe may be quite large and contain many frames, this message reduces the latency in the response time which is of paramount importance for the protection of incumbent services. In other words, this message allows termination of operation in a channel to be immediate or else to be scheduled at the earliest.

Table 134 – Message format

Syntax	Size	Notes
CHT-REQ_Message_Format() {		
Management Message Type = 43	8 bits	
Transaction ID	16 bits	
Starting Channel Number	8 bits	
Number of Channels	8 bits	
Confirmation Needed	1 bit	Indicates whether the CPE is required by the BS to confirm the receipt of this message. 0 = No confirmation needed (default) 1 = Confirmation needed
Terminate Mode	1 bit	Indicates any restrictions on transmission until

		termination of operation in a channel. The BS shall set the Terminate Mode field to either 0 or 1 on transmission. A value of 1 means that the CPE to which the frame containing the element is addressed shall transmit no further frames until the scheduled channel termination. A channel Terminate Mode set to 0 does not impose any requirement on the receiving CPE.
Terminate Count	8 bits	This field either shall be set to the number of frames until the BS sending the Channel Terminate message terminates the operation in the specified channels or shall be set to 0. A value of 1 indicates that termination of operation will occur immediately before the next frame. A value of 0 indicates that termination will occur at any time after the frame containing the element is transmitted.
}		

8.21.2 Channel Terminate Response (CHT-RSP)

This message (Table 135) is sent by the CPE in response to the receipt of a CHT-REQ. This message shall only be transmitted by the CPE if the Confirmation Needed field in the received CHT-REQ is set.

Table 135 – Message format

Syntax	Size	Notes
CHT-RSP_Message_Format() {		
Management Message Type = 44	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
}		

8.21.3 Channel Add Request (CHA-REQ)

This message (Table 136) is sent by the BS in order to add channel(s) to the regular operation of the BS, which can be due to several reasons such as some channels being detected as available for use by 802.22. Note that addition of channel(s) to the BS operation can also be implemented through the SCH by having the BS specify a different and larger value for Channel Number and Number of Channels. However, since a superframe may be quite large and contain many frames, this message reduces the latency in the response time which is of paramount importance for increasing the cell performance. In other words, this message allows the addition of channel(s) to be immediate or else to be scheduled at the earliest.

Table 136 – Message format

Syntax	Size	Notes
CHA-REQ_Message_Format() {		
Management Message Type = 45	8 bits	
Transaction ID	16 bits	
Starting Channel Number	8 bits	
Number of Channels	8 bits	
Confirmation Needed	1 bit	Indicates whether the CPE is required by the BS to confirm the receipt of this message. 0 = No confirmation needed (default) 1 = Confirmation needed
Addition Count	8 bits	This field either shall be set to the number of

		frames until the BS sending the Channel Add message adds the new channels or shall be set to 0. A value of 1 indicates that the addition will occur immediately before the next frame. A value of 0 indicates that the addition will occur at any time after the frame containing the element is transmitted.
}		

8.21.4 Channel Add Response (CHA-RSP)

This message (Table 137) is sent by the CPE in response to the receipt of a CHA-REQ. This message shall only be transmitted by the CPE if the Confirmation Needed field in the received CHA-REQ is set.

Table 137 – Message format

Syntax	Size	Notes
CHA-RSP_Message_Format() {		
Management Message Type = 46	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
}		

8.21.5 Channel Switch Request (CHS-REQ)

This message (Table 138) is sent by the BS in order to switch the entire cell operation (BS and CPEs) to different channel(s). Transmission of this message may be due to several reasons such as protection of incumbent services or availability of larger number or better quality channel(s).

Table 138 – Message format

Syntax	Size	Notes
CHS-REQ_Message_Format() {		
Management Message Type = 47	8 bits	
Transaction ID	16 bits	
Starting Channel Number	8 bits	
Number of Channels	8 bits	
Confirmation Needed	1 bit	Indicates whether the CPE is required by the BS to confirm the receipt of this message. 0 = No confirmation needed (default) 1 = Confirmation needed
Switch Mode	1 bit	Indicates any restrictions on transmission until a channel switch. The BS shall set the Switch Mode field to either 0 or 1 on transmission. A value of 1 means that the CPE to which the frame containing the element is addressed shall transmit no further frames until the scheduled channel switch. A Channel Switch Mode set to 0 does not impose any requirement on the receiving CPE.
Switch Count	8 bits	This field either shall be set to the number of frames until the BS sending the Channel Switch message switches to the new channel or shall be set to 0. A value of 1 indicates that the switch will occur immediately before the next frame. A value of 0 indicates that the

		switch will occur at any time after the frame containing the element is transmitted.
}		

8.21.6 Channel Switch Response (CHS-RSP)

This message (Table 139) is sent by the CPE in response to the receipt of a CHS-REQ. This message shall only be transmitted by the CPE if the Confirmation Needed field in the received CHS-REQ is set.

Table 139 – Message format

Syntax	Size	Notes
CHS-RSP_Message_Format() {		
Management Message Type = 48	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
}		

8.21.7 Channel Quiet Request (CHQ-REQ)

This message (Table 140) is sent by the BS in order to quiet any transmission activity in channel(s) currently used by the BS for communication with its associated CPEs. Transmission of this message may be due to several reasons such as the need to quiet the channel(s) to perform measurements. The transmission of a CHQ-REQ message should be preceded by the transmission of a BLM-REQ message.

Table 140 – Message format

Syntax	Size	Notes
CHQ-REQ_Message_Format() {		
Management Message Type = 49	8 bits	
Transaction ID	16 bits	
Confirmation Needed	1 bit	Indicates whether the CPE is required by the BS to confirm the receipt of this message. 0 = No confirmation needed (default) 1 = Confirmation needed
Quiet Count	8 bits	Shall be set to the number of frames until the next quiet interval shall start. A value of 1 indicates the quiet interval will start at the next frame. A value of 0 is reserved.
Quiet Offset	8 bits	Shall be set to the offset of the start of the quiet interval from the start of the frame specified by the Quiet Count field, expressed in slots. The value of the Quiet Offset field shall be less than one frame length.
Duration	16 bit	Shall be set to the duration of the quiet interval, expressed in slots. This duration shall be commensurate with all the measurement durations of all BLM-REQ in effect (see 8.22.1). This is an aggregate duration for all measurement requests. • If this field is set to a value different from 0 (zero): If quiet periods are already scheduled after Quiet Count and Quiet Offset and Transaction ID = Transaction ID of the already scheduled quiet

		period, then the value specified in this field shall override the length of all following quiet periods after Quiet Count and Quiet Offset. Otherwise, if quiet periods are already scheduled after Quiet Count and Quiet Offset but Transaction ID != Transaction ID of the already scheduled quiet period, then the value specified in this field shall override the length of only the first quiet periods after Quiet Count and Quiet Offset. • If this field is set to 0 (zero): If quiet periods are already scheduled after Quiet Count and Quiet Offset and Transaction ID = Transaction ID of the already scheduled quiet period, then all quiet periods after Quiet Count and Quiet Offset are cancelled. Otherwise, if quiet periods are already scheduled after Quiet Count and Quiet Offset but Transaction ID != Transaction ID of the already scheduled quiet period, then only the first quiet period after Quiet Count and Quiet Offset is cancelled. Normal data transmission should be carried out in the absence of a quiet period.
Quiet Period	8 bits	Shall be set to the number of frames between the start of regularly scheduled quiet intervals defined by this Quiet element. A value of 0 indicates that no periodic quiet interval is defined.
Number of Quiet Period Purposes	3 bits	
Quiet Period Purpose IEs	24 bits	See Table 141
}		

Table 141 – Encoding of quiet period purpose

Syntax	Size	Notes
Quiet_Period_Purpose_Format() {		
Element ID	8 bits	
Length	8 bits	
Type of Purpose	1 bit	Indicates the type of the purpose. 0 = Incumbent (default) 1 = 802.22
Purpose	3 bits	Indicates the purpose of the quiet period. The meaning of this field depended upon the value of the Type of Purpose field. If Type of Purpose = 0 Bit #0 (MSB): If set, the CPE shall perform TV service measurement (see Table 70) Bit #1: If set, the CPE shall perform Part 74 measurements Bit #2: If set, the CPE shall perform WMB measurements If Type of Purpose = 1 Bit #0 (MSB): If set, the CPE shall perform 802.22 CBP measurements Bit #1: If set, the CPE shall perform 802.22 BS beacon measurements Bit #2: Undefined

Fraction	4 bits	In ascending order, indicates the fraction of the total Quiet Period Duration (Table 140) that is dedicated for this specific purpose. If case of conflict between this time allocation and the one specified in the BLM-REQ (see 8.22.1), this allocation shall prevail. The sum of the Fractions of all quiet periods purposes shall not exceed 100%. The CPE is free to use any unassigned fraction as it sees fit. 0000 0% 0001 10% 0010 20% 0011 30% 0100 40% 0101 50% 0110 60% 0111 70% 1000 80% 1001 90% 1011 100% Others Undefined
}		

8.21.8 Channel Quiet Response (CHQ-RSP)

This message (Table 142) is sent by the CPE in response to the receipt of a CHQ-REQ. This message shall only be transmitted by the CPE if the Confirmation Needed field in the received CHQ-REQ is set.

Table 142 – Message format

Syntax	Size	Notes
CHQ-RSP_Message_Format() {		
Management Message Type = 50	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
}		

8.21.9 Channel Occupancy Update (CHO-UPD)

This message (Table 143) is sent by the BS in order to inform CPEs about the consolidated channel occupancy information in the overall cell. Typically, this would be transmitted by the BS once it has received enough measurement reports from its associated CPEs that allow it to make a reliable decision on overall channel occupancy throughout the cell. This allows several features to be implemented including better channel management, optimization of a CPE's measurement activities, and improved recovery procedure in case of in-band detection of incumbents.

Table 143 – Message format

Syntax	Size	Notes
CHO-UPD_Message_Format() {		
Management Message Type = 51	8 bits	
Transaction ID	16 bits	
Starting Channel Number	8 bits	
Number of Channels	8 bits	

for (i = 1; i ≤ **Number of Channels**; i++) {		
Channel State	3 bits	See Table 144
}		
Padding Nibble		
}		

Table 144 – Channel state information

Value	State
000	Unmeasured
001	Vacant
010	Occupied
011	Occupied by Incumbent
100	Occupied by 802.22
101-111	*Reserved*

8.22 Measurements Management

In this section we present the measurements management component of CMAC, which is a critical component for many features of the protocol including for guaranteeing incumbent system protection at all times[6].

We start this discussion in Section 8.22.1 by presenting the bulk measurement request (BLM-REQ) management message. This message is transmitted (e.g., via unicast/multicast/broadcast) by the BS to one or multiple CPEs, and contains instructions on the type of measurements to be performed, when to perform, the measurement duration, on which channels, and so on. Since the correct receipt of these management messages may be critical to the correct system behaviour (especially for in-band measurements – see 21.1.1), the BS may require CPEs to acknowledge the receipt of BLM-REQ messages. This is done through bulk measurement response (BLM-RSP) messages, and these are covered in Section 8.22.2. Next, Section 8.22.3 deals with bulk measurement report (BLM-REP) messages which, as the name suggests, allows CPEs to report back to the BS all the measurement data they have collected as per requested by the BS in the corresponding BLM-REQ message.

8.22.1 Bulk Measurement Request (BLM-REQ)

Table 145 illustrates the format of a BLM-REQ message. BLM-REQ messages can be comprised of a multitude of single measurement messages (see 9.512.58.22.1.1). Each of these single measurement requests can be associated with a different type of measurement, and hence provides a high degree of flexibility to the system.

Table 145 – Message format

Syntax	Size	Notes
BLM-REQ_Message_Format() {		
Management Message Type = 39	8 bits	
Transaction ID	16 bits	Shall be set to a nonzero value chosen by the BS sending the measurement request to identify the request/report transaction.
Starting Channel Number	8 bits	
Number of Channels	8 bits	
Confirmation Needed	1 bit	Indicates whether or not the CPE is required by the BS to confirm, with a BLM-RSP message, the receipt of this message. 0 = No confirmation needed (default)

[6] The protocol for using the spectrum measurement messages is described in Section 21.1.

		1 = Confirmation needed
Number of Single Measurement Requests	3 bits	The number of single measurement requests contained in this message
Single Measurement Requests	*Variable*	A series of single measurement requests. See Table 146.
}		

8.22.1.1 Single Measurement Request

Table 146 illustrates an example of the format of single measurement requests, which are carried in the body of BLM-REQ management messages. These single measurement requests can be of various types as shown in Table 150. Also, we can see from Table 146 that various timing parameters are associated with measurement requests. Figure 8 depicts how these parameters are related to a measurement activity (the Randomization Interval and Duration parameters are introduced in the next subsections).

Table 146 – Message format

Syntax	Size	Notes
Single_Measurement_Request_Format() {		
Element ID	8 bits	Table 150
Length	8 bits	
Transaction ID	16 bits	
Number of Repetitions	16 bits	Contains the requested number of repetitions per channel for the periodic measurement request elements in this frame. A value of zero in the Number of Repetitions field indicates measurement request elements are executed only once.
Report Frequency	8 bits	This field indicates how often a CPE shall report measurements back to the BS 0: This field is not used to request a measurement report. That is, the CPE will report measurements either autonomously or whenever requested by the BS. 1: The CPE shall either report immediately to the BS (if this is in regards to an existing Transaction ID) or will report to the BS at the end of each repetition interval (in the case of a new Transaction ID). Note that in the case of an existing Transaction ID and Report Frequency == 1, the local information maintained by the CPE shall only be updated for this transaction if Number of Repetitions is not zero. 2-127 The CPE shall send a report to the BS at the end of every X number of repetitions.
Restart Delay	16 bits	This field indicates the delay between two measurement repetitions. As shown in Table 147, the Measurement Period is divided into two subfields: Time Scale and Restart Delay. The Time Scale subfield defines the scale for the Restart Delay subfield as shown in Table 148. The subfield consists of a 15 bit unsigned integer number representing the fixed time delay between the completion of the last periodic measurement until the measurement activity is restarted.
Request Mode	3 bits	Table 149

Request Element	*Variable*	Table 150
}		

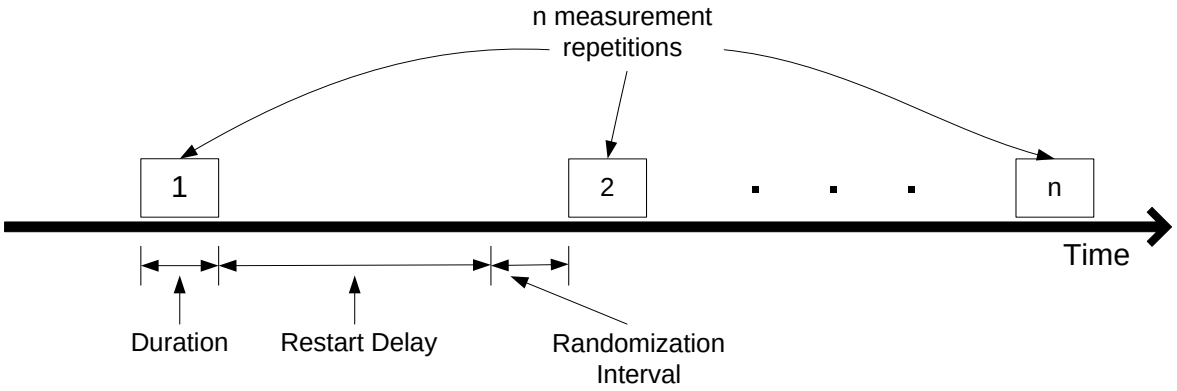

Figure 8 – Illustration of the timing parameters used in measurement requests

Table 147 – Repetition delay field

Bits: 1	15
Time Scale	Restart delay

Table 148 – Time unit (TU) and time scale definitions

Time Unit	Time Scale value
TU	0
1000 TU	1

Table 149 – Request mode

Syntax	Size	Notes
Request_Mode_Format() {		
Parallel	1 bit	Indicates whether the measurement should start in series or in parallel with the measurement described by any immediately previous Measurement Request element in the same Measurement Request frame. A value of 0 shall mean that the measurement shall start after the previous measurement request has completed. A value of 1 shall mean the measurement shall start at the same time as the previous measurement. The Parallel bit shall be set to 0 in the first or only measurement request element in the frame.
Autonomous Report	1 bit	Indicates whether the CPE receiving the request shall enable or disable autonomous measurement reports for the measurements specified in this request. The Report bit shall be set to 1 when enabling autonomous measurement report. The Report bit shall be set to 0 when disabling an autonomous measurement report.
Duration Mandatory	1 bit	Indicates whether the measurement duration contained within the Measurement Request should be interpreted as mandatory by the CPE receiving the request. A value of 0 shall indicate that the duration requested is target duration, but the requesting BS shall accept measurement results taken over a shorter duration. A value of 1 shall indicate that the duration requested is a mandatory duration.
}		

Table 150 –Request Information Elements

Element ID (1 byte)	Length (bytes)	Description
0	*Variable*	TV System Related Measurement Request – Table 151
1	*Variable*	Part 74 System Related Measurement Request – Table 151
2	*Variable*	Beacon (802.22 BS and CPE Related) Measurement Request – Table 152
3	*Variable*	Stop Measurement Request – Table 153
4	*Variable*	CPE Statistics Measurement Request – Table 155
5	*Variable*	Location Configuration Measurement Request – Table 157
6	*Variable*	Beacon (Part 74) Measurement Request – Table 152
7	*Variable*	Channel Feedback Measurement Request – Table 217
8-128		*Reserved*

8.21.1.1.1 Signal Specific Measurement Request

This refers to a particular incumbent signal specific measurement. Note that the incumbent detection thresholds are not specified here, but during the registration procedure (see 8.7.3.7.9).

Table 151 – Message format

Syntax	Size	Notes
Signal_Specific_Measurement_Request_Format() {		
System Profile	8 bits	See Table 70
Randomization Interval	16 bits	This field only applies to out-of-band measurements, as in-band measurements are driven by quiet periods. It specifies the upper bound of the random delay that can be used by the CPE prior to making the measurement. It is specified in units of TU (see Table 148).
Duration	16 bits	Shall be set to the preferred duration of the requested measurement, expressed in TUs.
}		

8.21.1.1.2 Beacon Measurement Request

Table 152 – Message format

Syntax	Size	Notes
Beacon_Measurement_Request_Format() {		
Randomization Interval	16 bits	This field only applies to out-of-band measurements, as in-band measurements are driven by quiet periods. It specifies the upper bound of the random delay that can be used by the CPE prior to making the measurement. It is specified in units of TU (see Table 148).
Duration	16 bits	Shall be set to the preferred duration of the requested measurement, expressed in TUs.
ID	48 bits	Specifies the ID (e.g., MAC ID) to listen to. Can be a broadcast ID or a specific station ID.
}		

8.21.1.1.3 Measurement Stop Request

Table 153 – Message format

Syntax	Size	Notes
Stop_Measurement_Request_Format() {		
Stop Time	16 bits	Consists of an unsigned integer number representing the time at which the CPE shall stop conducting all measurements activities. The Stop Time field consists of a Time Scale subfield and a Stop Time subfield as shown in Table 154. The Time Scale subfield is defined in Table 148 and represents the time units for the integer in the Stop Time subfield.
}		

Table 154 – Pause Time field

Bits: 1	15
Time Scale	Pause time

8.21.1.1.4 CPE Statistics Measurement Request

Table 155 – Message format

Syntax	Size	Notes
CPE_Statistics_Measurement_Request_Format() {		
Group Identity	8 bits	Table 156
Randomization Interval	16 bits	This field only applies to out-of-band measurements, as in-band measurements are driven by quiet periods. It specifies the upper bound of the random delay that can be used by the CPE prior to making the measurement. It is specified in units of TU (see Table 148).
Duration	16 bits	Shall be set to the preferred duration of the requested measurement, expressed in TUs.
}		

Table 156 – Group Identity

Statistics Name	Group Identity
CPE Counters from dot22CountersTable	0
Reserved	1-255

8.21.1.1.5 Location Configuration Measurement Request

Table 157 – Message format

Syntax	Size	Notes
Location_Configuration_Measurement_Request_Format() {		
LCI Discovery Mode	8 bits	Indicates the means by which the CPE should acquire its location information. • 0 – CPE should infer its own location information from other CPEs/BSs. • 1 – External ways can be used (e.g., GPS). • 2 – Either internal or external means can be used. • 3-255 – *Reserved*
}		

8.22.2 Bulk Measurement Response (BLM-RSP)

A BLM-RSP management message (shown in Table 158) is sent in response to a BLM-REQ and serves to confirm the receipt of the BLM-REQ message by the CPE. The need to send a BLM-RSP message is indicated by the BS in the corresponding BLM-REQ message, through the use of the Confirmation Needed field.

Table 158 – Message format

Syntax	Size	Notes
BLM-RSP_Message_Format() {		
Management Message Type = 40	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
}		

8.22.3 Bulk Measurement Report (BLM-REP)

A BLM-REP management message (see Table 159) is sent from a CPE to a BS, and contains the measurement data collected by the CPE as per requested by the BS in a preceding BLM-REQ message.

Table 159 – Message format

Syntax	Size	Notes
BLM-REP_Message_Format() {		
Management Message Type = 41	8 bits	
Transaction ID	16 bits	
Number of Single Measurement Reports	8 bits	The number of single measurement reports contained in this message
Single Measurement Reports	*Variable*	Table 160
}		

8.22.3.1 Single Measurement Report

Table 160 – Message format

Syntax	Size	Notes
Single_Measurement_Report_Format() {		
Element ID	8 bits	Table 161
Length	8 bits	
Transaction ID	16 bits	
Report Information Element	*Variable*	Table 161
}		

Table 161 –Report Information Elements

Element ID (1 byte)	Length (bytes)	Description
129	*Variable*	TV Measurement Report – Table 162
130	*Variable*	Part 74 Measurement Report – Table 162
131	*Variable*	Beacon (802.22 Related) Measurement Report – Table 164
132	*Variable*	CPE Statistics Measurement Report – Table 167
133	*Variable*	Location Configuration Measurement Report – Table 169
134	*Variable*	Beacon (Part 74) Measurement Report – Table 170
133	*Variable*	Consolidated Spectrum Occupancy Measurement Report – Table 171

132	*Variable*	Channel Feedback Measurement Report – Table 219
131-255		*Reserved*

8.21.3.1.1 Signal Specific Measurement Report

Table 162 – Message format

Syntax	Size	Notes
Signal_Specific_Measurement_Report_Format() {		
System Profile	8 bits	See Table 70
Report Mode	4 bits	Table 163
Start Frame	8 bits	Frame number (see Table 26) in which the channel measurement started
Duration	16 bits	The actual duration of the measurement
Channel Number	8 bits	
Value	10 bits	The value (e.g., output SINR) of the measurement
Precision	6 bits	Indicates the accuracy (significance) of the measured value
}		

Table 163 – Report mode

Syntax	Size	Notes
Report_Mode_Format() {		
Late	1 bit	Indicates whether this CPE is unable to carry out a measurement request because it received the request after the requested measurement time. The Late bit shall be set equal to 1 to indicate the request was too late. The Late bit shall be set to 0 to indicate the request was received in time for the measurement to be executed, or if no start time was specified.
Incapable	1 bit	Indicates whether this CPE is incapable of generating this report requested by the BS. The Incapable bit shall be set to 1 to indicate the CPE is incapable. The Incapable bit shall be set to 0 to indicate the CPE is capable or the report is autonomous.
Refused	1 bit	Indicates whether this CPE is refusing to generate this report requested by the BS. The Refused bit shall be set to 1 to indicate the CPE is refusing. The Refused bit shall be set to 0 to indicate the CPE is not refusing or the report is autonomous.
Unmeasured	1 bit	CPE did not measure the channel
}		

8.21.3.1.2 Beacon Measurement Report

A beacon measurement report (see Table 164) is sent from a CPE to its corresponding BS, and conveys information about one single overhead SCH (transmitted by other BSs) and/or CBP packet (transmitted by other CPEs) originated at other collocated 802.22 cells. A CPE shall never report on beacons that are originated at its own cell.

Table 164 – Message format

Syntax	Size	Notes
Beacon_Measurement_Report_Format() {		
Element ID	8 bits	
Length	8 bits	
Report Mode	4 bits	Table 163
Start Frame	16 bits	Table 162
Duration	16 bits	Table 162
Frame Number	8 bits	The frame number in which the beacon was received. See definition in Table 26
Reception Offset	8 bits	Indicates the offset (in units of slot duration) relative to the start of the first slot of the PHY PDU (including preamble) frame where the beacon was received. The time instants indicated by the Reception Offset values are the reception times of the first slot of the beacon including preamble (if present).
Channel Number	8 bits	
Number of Channels	8 bits	
FDC	8 bits	See Table 27
FS	7 bits	See Table 1
TTQP	16 bits	See Table 1
DQP	16 bits	See Table 1
BS ID	48 bits	ID/Address that uniquely identifies the BS
BS/CPE IEs	*Variable*	Table 165 and Table 166
}		

Table 165 – BS IE

Syntax	Size	Notes
BS_IE_Format() {		
Element ID	8 bits	
Length	8 bits	
RCPI	8 bits	Received Carrier Power Indicator (in dBm)
Link Margin	8 bits	In dBm
}		

Table 166 – CPE IE

Syntax	Size	Notes
CPE_IE_Format() {		
Element ID	8 bits	
Length	8 bits	
Frame Number	8 bits	Frame number, with respect to the transmitter, where the beacon was sent. See Table 26
Transmission Offset	8 bits	See Table 8
CPE ID	48 bits	Address/ID that uniquely identifies the CPE
Channel Number for Backup	8 bits	See definition in Table 26
Number of Channels for Backup	8 bits	See definition in Table 26
Starting DS Allocation Channel	8 bits	See Table 8
Ending DS Allocation Channel	8 bits	See Table 8
Ending DS Allocation Slot	7 bits	See Table 8
Starting US Allocation Channel	8 bits	See Table 8
Ending US Allocation Channel	8 bits	See Table 8
Starting US Allocation Slot	7 bits	See Table 8
RCPI	8 bits	See Table 165

Link Margin	8 bits	See Table 165
Beacon IEs	*Variable*	See Table 9
}		

8.21.3.1.3 CPE Statistics Measurement Report

Table 167 – Message format

Syntax	Size	Notes
CPE_Statistics_Measurement_Report_Format() {		
Report Mode	4 bits	Table 163
Start Frame	16 bits	Frame number (see Table 26) in which the channel measurement started
Duration	16 bits	
Group Statistics Data	*Variable*	Table 168
}		

Table 168 – Group statistics data

Group Identity Requested	Statistics Returned (possibly stored in the MIB) (32 bits)
0	dot22TransmittedCoexistenceBeaconCount dot22TransmittedFragmentCount dot22TransmittedFrameCount dot22MulticastTransmittedFrameCount dot22FailedCount dot22RetryCount dot22MultipleRetryCount dot22FrameDuplicateCount dot22ReceivedFragmentCount dot22ReceivedCoexistenceBeaconCount dot22ReceivedBSBeaconCount dot22MulticastReceivedFragmentCount dot22CRCErrorCount
1-255	None

8.21.3.1.4 Location Configuration Measurement Report

A Location Configuration report (see Table 169), as described in IETF RFC 3825 (“Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information”), includes latitude, longitude and altitude. The report format shall be as described in RFC 3825, and the length shall be 16 octets.

Table 169 – Message format

Syntax	Size	Notes
Location_Configuration_Measurement_Report_Format() {		An LCI with Latitude resolution, Longitude resolution and Altitude resolution set to zero shall indicate that the location is not known.
Element ID	8 bits	
Length	8 bits	Shall be 16 octets
Report Mode	4 bits	Table 163
Latitude Resolution	6 bits	Latitude resolution indicates the number of valid bits in the fixed-point value of Latitude
Latitude	34 bits	Latitude is a fixed point value consisting of 9 bits of integer and 25 bits of fraction
Longitude Resolution	6 bits	Longitude resolution indicates the number of

		valid bits in the fixed-point value of Longitude
Longitude	34 bits	Longitude is a fixed point value consisting of 9 bits of integer and 25 bits of fraction
Altitude Type	4 bits	Altitude Type encodes the type of altitude. Codes defined are: • 1: Meters – in 2s-complement fixed-point 22-bit integer part with 8-bit fraction; • 2: Floors – in 2s-complement fixed-point 22-bit integer part with 8-bit fraction; Altitude type = 2 for Floors enables representing altitude in a form more relevant in buildings which have different floor-to-floor dimensions.
Altitude Resolution	6 bits	Altitude resolution indicates the number of valid bits in the altitude
Altitude	30 bits	Altitude is a value defined by the Altitude type field
Datum	8 bits	Datum is encodes the horizontal and vertical references used for the coordinates. The Datum octet has 256 possibilities, of which 3 have been registered with the Internet Assigned Numbers Authority (IANA): • 1: WGS 84 (Geographical 3D) – World Geodesic System 1984, Coordinate Reference System (CRS) Code 4327, Prime Meridian Name: Greenwich; • 2: NAD83 – North American Datum 1983, CRS Code 4269, Prime Meridian Name: Greenwich; The associated vertical datum is the North American Vertical Datum of 1988 (NAVD88); • 3: NAD83 – North American Datum 1983, CRS Code 4269, Prime Meridian Name: Greenwich; The associated vertical datum is Mean Lower Low Water (MLLW). The WGS 84 datum shall be used when referencing locations anywhere. The GeoConf Option referred to in IETF RFC 3825 defines two fields for which the IANA maintains a registry: The Altitude type (AT) field and the Datum field. The initial values of the Altitude registry are as follows: • AT = 1 meters of Altitude defined by the vertical datum specified; • AT = 2 building Floors of Altitude.
}		

8.21.3.1.5 Part 74 Beacon Measurement Report

A WMB measurement report (see Table 170) is transmitted from a CPE to its corresponding BS, and conveys information about any overhead WMB (transmitted by class B CPEs). This report shall not be sent if the CPE did not receive any WMB. Furthermore, the report shall only be sent if the Security Key of the WMB device is valid. See 21.1.7 for more information on this feature.

Table 170 – Message format

Syntax	Size	Notes
Part74_Beacon_Measurement_Report_Format() {		

Element ID	8 bits	
Length	8 bits	
Report Mode	4 bits	Table 163
Start Frame	16 bits	Table 162
Duration	16 bits	Table 162
Reserved	4 bits	
Class B CPE ID	48 bits	ID/Address that uniquely identifies the class B CPE who transmitted the WMB.
Transaction ID		See Table 225
Part 74 Operation Start Time		See Table 225
Part 74 Operation Duration		See Table 225
IEs		See Table 225
}		

8.21.3.1.6 Consolidated Spectrum Occupancy Measurement Report

A consolidated spectrum occupancy measurement report (see Table 171) is sent from a CPE to its corresponding BS, and conveys a brief summary about the overall spectrum occupancy from the viewpoint of the CPE.

Table 171 – Message format

Syntax	Size	Notes
Consolidated_Spectrum_Occupancy_Measurement_Report_Format() {		
Start Frame	8 bits	Frame number (see Table 26) in which the channel measurement started
Duration	16 bits	The actual duration of the measurement
Starting Channel Number	8 bits	
Number of Channels	8 bits	
for (i = 1; i ≤ **Number of Channels**; i++) {		
Channel State	3 bits	See Table 144
}		
Padding Nibble		
}		

8.22.4 Bulk Measurement Acknowledgement (BLM-ACK)

A BLM-ACK management message (shown in Table 172) shall be sent from the BS to the CPE in response to a received BLM-REP. It serves to confirm to the CPE the reception of the BLM-REP message by the BS.

Table 172 – Message format

Syntax	Size	Notes
DSA-ACK_Message_Format() {		
Management Message Type = 42	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
}		

8.23 Scheduling Constraint

8.23.1 Traffic Constraint Request (TRC-REQ)

Table 173 – Message format

Syntax	Size	Notes
TRC-REQ_Message_Format() {		
Management Message Type = 52	8 bits	
Transaction ID	16 bits	
Downstream	1 bit	
Upstream	1 bit	
CoS	3 bits	
}		

8.23.2 Traffic Constraint Response (TRC-RSP)

Table 174 – Message format

Syntax	Size	Notes
TRC-RSP_Message_Format() {		
Management Message Type = 53	8 bits	
Transaction ID	16 bits	
DS/US Traffic Constraint Information Elements	Variable	
}		

8.23.2.1 DS/US Traffic Constraint IE

Table 175 – Message format

Syntax	Size	Notes
DS/US_Traffic_Constraint_Format() {		
Element ID	8 bits	
Length	8 bits	
US/DS	1 bits	Indicates whether this is a Upstream (set to 0) or Downstream (set to 1) constraint
CoS	3 bits	
Spectrum Usage Bitmap (SUB)	*Variable*	Compound of Time SUB (Table 176) and Frequency SUB (Table 177)
}		

8.23.2.1.1 Time SUB

Table 176 – Message format

Syntax	Size	Notes
Time_SUB_Format() {		
Type = 0	1 bit	Indicates whether this is a Time SUB (set to 0) or Frequency SUB (set to 1).
Length	8 bits	
for (i = 1; i ≤ Number_Slots_Per_Frame; i++)		
Slot I	1 bit	Set to 1, if slot *i* is currently in use. Set to 0, otherwise.
}		

8.23.2.1.2 Frequency SUB

Table 177 – Message format

Syntax	Size	Notes
Frequency_SUB_Format() {		
Type = 1	1 bit	Indicates whether this is a Time SUB (set to 0) or Frequency SUB (set to 1).
Starting Channel Number	8 bits	Logical channel number
Number of Channels	8 bits	Logical channel number
}		

8.24 Timeout

Timeout (TMO) messages may be sent by a CPE to the BS (or vice-versa) and, in this case, refers to a solicitation from the CPE to be excused from any communication with the BS for up to a specified duration of time (see Table 178). During this time, the BS shall retain the connection identifiers (basic, primary, and secondary) of the excused CPE for when it returns to the cell. If a CPE returns within the negotiated time with the BS, it shall not need to re-register with the BS. Alternatively, a BS may de-register a CPE which does not return in the negotiated time. In this case, the CPE would have to go through the initialization steps all over again before gaining access to the network.

Timeout messages may have several uses. A CPE may request a timeout to look for and associate with multiple BSs in its neighbourhood. In this case, load balancing algorithms could be implemented, and a CPE could determine their location information provided it can obtain the location from at least three other BSs. Another use for timeout messages is for CPEs to perform out-of-band measurements (see 21.1.1) and attempt assess the radio spectrum usage in its vicinity (this is in addition to the quiet periods scheduled by the BS). A third use of timeout intervals would be for power conservation.

8.24.1 Timeout Request (TMO-REQ)

Table 178 – Message format

Syntax	Size	Notes
TMO-REQ_Message_Format() {		
Management Message Type = 54	8 bits	
Transaction ID	16 bits	
Frame Number Index	8 bits	Identifies the frame number (see Table 26) in which the CPE desires to disconnect from the BS.
Duration	16 bit	Duration for which the station is requesting a timeout
}		

8.24.2 Timeout Response (TMO-RSP)

Table 179 – Message format

Syntax	Size	Notes
TMO-RSP_Message_Format() {		

Management Message Type = 55	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
}		

8.25 Frame Slide

Frame slide messages are transmitted by the BS only, and are used for the purpose sliding the frame in time. The BS shall only use these messages when frame synchronization amongst overlapping BS is needed. A BS shall always synchronize its frame with other detected overlapping BSs in order to improve self-coexistence. The method to achieve synchronization is described in 21.3. Note that the BS shall always start a new superframe transmission at the time the frame slide is to be effected.

8.25.1 Frame Slide Request (FSL-REQ)

Table 180 – Message format

Syntax	Size	Notes
FSL-REQ_Message_Format() {		
Management Message Type = 56	8 bits	
Transaction ID	16 bits	
Confirmation Needed	1 bit	Indicates whether the CPE is required by the BS to confirm the receipt of this message. 0 = No confirmation needed (default) 1 = Confirmation needed
BS ID	48 bits	Address that uniquely identifies the BS to which synchronization is being sought
Slide Count	8 bits	This field either shall be set to the number of frames, including the current one, until the BS sending this message slides the frame or shall be set to 0. A value of 1 indicates that the frame slide will occur immediately after the current frame. A value of 0 indicates that the addition will occur at any time after the frame containing the element is transmitted.
Slide Amount	8 bits	Indicates (in units of symbol duration) how much the frame slide is going to be.
Direction	1 bit	Indicates the direction of the frame slide 0 = Right (default) 1 = Left
}		

8.25.2 Frame Slide Response (FSL-RSP)

Table 181 – Message format

Syntax	Size	Notes
FSL-RSP_Message_Format() {		
Management Message Type = 57	8 bits	
Transaction ID	16 bits	
Confirmation Code	8 bits	Table 82
}		

8.26 Config File TFTP Complete (TFTP-CPLT)

The Config File TFTP-CPLT message shall be generated by the CPE whenever it has successfully retrieved its configuration file from the provisioning server (see 15). If the CPE does not need a config file, it shall send the TFTP-CPLT message to the BS anyway to indicate that it has completed secondary management connection initialization and is ready to accept services. The format of the TFTP-CPLT shall be as shown in Table 182.

Table 182 – Message format

Syntax	Size	Notes
TFTP-CPLT_Message_Format() {		The CID in the MAC header shall be set to the CPE's primary management CID
Management Message Type = 31	8 bits	
Information Elements (IEs)	*Variable*	Shall include at least the HMAC tuple (see 7.1), which shall be the last IE in the message.
}		

8.26.1 Configuration File Encodings

These settings are found only in the configuration file. They shall not be forwarded to the BS in the Registration Request.

8.26.1.1 CPE MIC Configuration Setting

This value field contains the CPE MIC code (see Table 183). This is used to detect unauthorized modification or corruption of the configuration file.

Table 183 – Information element

Name	Element ID (1 byte)	Length (bytes)	Value
CPE MIC	1	20	d1 d2 … d20

8.26.1.2 Software Upgrade Descriptors

This field defines the parameters associated with software upgrades (see Table 184). It is composed of one or more upgrade descriptors. An upgrade descriptor is defined by the set of all encapsulated tags defined in 8.26.1.2.1 through 8.26.1.2.4, occurring in order in the TFTP file. A new upgrade descriptor begins with the occurrence of the Vendor ID IE.

When a managed CPE decodes a descriptor with a matching Vendor ID, Hardware ID, and Software version different than the one currently running, it may initiate a TFTP transfer to upgrade its software.

Table 184 – Information element

Name	Element ID (1 byte)	Length (bytes)	Value
Software upgrade descriptor	2	Variable	Compound

8.26.1.2.1 Vendor ID

This value identifies the managed CPE vendor to which the software upgrade is to be applied (see Table 185). Its format and value is described in 7.5.

Table 185 – Information element

Name	Element ID (1 byte)	Length (bytes)	Value
Vendor ID	2.144	3	v1, v2, v3

8.26.1.2.2 Hardware ID

This value identifies the hardware version to which the software upgrade is to be applied (see Table 186). This value is administered by the vendor identified by the Vendor ID field.

Table 186 – Information element

Name	Element ID (1 byte)	Length (bytes)	Value
Hardware ID	2.1	n	String

8.26.1.2.3 Software Version

This value identifies the software version of the software upgrade file (see Table 187). The value is administered by the vendor identified in the Vendor ID field. It should be defined by the vendor to be unique with respect to a given Hardware ID.

Table 187 – Information element

Name	Element ID (1 byte)	Length (bytes)	Value
Software version	2.2	n	String

8.26.1.2.4 Upgrade Filename

The filename of the software upgrade file for the managed CPE (see Table 188). The filename is a fully qualified directory-path name that is in a format appropriate to the server. There is no requirement that the character string be null-terminated; the length field always identifies the end of the string. The file is expected to reside on a TFTP server identified in a configuration setting option defined in 8.26.1.3.

Table 188 – Information element

Name	Element ID (1 byte)	Length (bytes)	Value
Upgrade filename	2.3	n	String

8.26.1.3 Software Upgrade TFTP Server

This object is the IP address of the TFTP server on which the software upgrade file for the CPE resides (see Table 189).

Table 189 – Information element

Name	Element ID (1 byte)	Length (bytes)	Value
Software upgrade TFTP server	3	4 or 16	IP address

8.26.1.4 TFTP Server Timestamp

This is the sending time of the configuration file in seconds (see Table 190). The definition of time is as in IETF RFC 868. Note that the purpose of this parameter is to prevent replay attacks with old configuration files.

Table 190 – Information element

Name	Element ID (1 byte)	Length (bytes)	Value
TFTP server timestamp	4	4	Number of seconds since 00:00 1 January 1900

8.27 Config File TFTP Complete Response (TFTP-RSP)

The Config File TFTP-RSP message shall be generated by the BS in response to a TFTP-CPLT message from the CPE (see 15). The format of the TFTP-RSP shall be as shown in Table 191.

Table 191 – Message format

Syntax	Size	Notes
TFTP-RSP_Message_Format() {		The CID in the MAC header shall be set to the CPE's primary management CID
Management Message Type = 32	8 bits	
Response	8 bits	0 = OK 1 = Message authentication failure
}		

8.28 Privacy Key Management (PKM) Messages (PKM-REQ/PKM-RSP)

PKM employs two MAC message types: PKM Request (PKM-REQ) and PKM Response (PKM-RSP), as described in Table 192.

Table 192 – PKM MAC messages

Type Value	Message name	Message description
9	PKM-REQ	Privacy Key Management Request [CPE -> BS]
10	PKM-RSP	Privacy Key Management Response [BS -> CPE]

These MAC management message types distinguish between PKM requests (CPE–to–BS) and PKM responses (BS–to–CPE). Each message encapsulates one PKM message in the Management Message Payload.

PKM protocol messages transmitted from the CPE to the BS shall use the form shown in Table 193. They are transmitted on the CPEs Primary Management Connection.

Table 193 – PKM request (PKM-REQ) message format

Syntax	Size	Notes
PKM-REQ_Message_Format() {		
Management Message Type = 9	8 bits	
Code	8 bits	Identifies the type of PKM packet. When a packet is received with an invalid Code, it shall be silently discarded. The code values are defined in Table 195.
PKM Identifier	8 bits	A CPE uses the identifier to match a BS response to the CPE's requests.

		The CPE shall increment (modulo 256) the Identifier field whenever it issues a new PKM message. A "new" message is an Authorization Request or Key Request that is not a retransmission being sent in response to a Timeout event. For retransmissions, the Identifier field shall remain unchanged. The Identifier field in Authentication Information messages, which are informative and do not effect any response messaging, shall be set to zero. The Identifier field in a BS's PKM-RSP message shall match the Identifier field of the PKM-REQ message the BS is responding to. The Identifier field in TEK Invalid messages, which are not sent in response to PKM-REQs, shall be set to zero. The Identifier field in unsolicited Authorization Invalid messages shall be set to zero. On reception of a PKM-RSP message, the CPE associates the message with a particular state machine (the Authorization state machine in the case of Authorization Replies, Authorization Rejects, and Authorization Invalids; a particular TEK state machine in the case of Key Replies, Key Rejects, and TEK Invalids). An CPE shall keep track of the identifier of its latest, pending Authorization Request. The CPE shall discard Authorization Reply and Authorization Reject messages with Identifier fields not matching that of the pending Authorization Request. A CPE shall keep track of the identifiers of its latest, pending Key Request for each SA. The CPE shall discard Key Reply and Key Reject messages with Identifier fields not matching those of the pending Key Request messages.
TLV Encoded Attributes	*variable*	TLV specific PKM attributes carry the specific authentication, authorization, and key management data exchanged between client and server. Each PKM packet type has its own set of required and optional attributes. Unless explicitly stated, there are no requirements on the ordering of attributes within a PKM message. The end of the list of attributes is indicated by the Length field of the MAC PDU header.
}		

PKM protocol messages transmitted from the BS to the CPE shall use the form shown in Table 194. They are transmitted on the CPEs Primary Management Connection.

Table 194 – PKM response (PKM-RSP) message format

Syntax	Size	Notes
PKM- RSP_Message_Format() {		
Management Message Type = 10	8 bits	
Code	8 bits	**See** Table 193
PKM Identifier	8 bits	**See** Table 193
TLV Encoded Attributes	*variable*	TLV specific See Table 193
}		

Table 195 – PKM message codes

Code	PKM message type	MAC Management Message Name
0-2	*reserved*	
3	PKM RSA-Request	PKM-REQ
4	PKM RSA-Reply	PKM-RSP
5	PKM RSA-Reject	PKM-RSP
6	PKM RSA-Acknowledgement	PKM-REQ
7	PKM EAP Start	PKM-REQ
8	PKM EAP-Transfer	PKM-REQ/PKM-RSP
9	PKM Authenticated EAP-Transfer	PKM-REQ/PKM-RSP
10	PKM SA TEK Challenge	PKM-RSP
11	PKM SA TEK Request	PKM-REQ
12	PKM SA TEK Response	PKM-RSP
13	PKM Key-Request	PKM-REQ
14	PKM Key-Reply	PKM-RSP
15	PKM Key-Reject	PKM-RSP
16	PKM SA-Addition	PKM-RSP
17	PKM TEK-Invalid	PKM-RSP
18	PKM Group-Key-Update-Command	PKM-RSP
19	PKM EAP Complete	PKM-RSP
20	PKM Authenticate EAP Start	PKM-REQ
21	PKM Auth Invalid	PKM-RSP
22	PKM Auth Info	PKM-REQ
23—255	*reserved*	

Formats for each of the PKM messages are described in the following subclauses.

8.28.1 PKM RSA-Request

A client CPE sends a PKM RSA-Request message to the BS in order to request mutual authentication in the RSA-based authorization.

Code: 3

Table 196 – PKM RSA-Request attributes

Attribute	Contents
CPE_Random	A 64-bit random number generated in the CPE
CPE_Certificate	Contains the CPE s X.509 user certificate
SAID	CPE s primary SAID equal to the Basic CID

SigCPE	An RSA signature over all the other attributes in the message

The CPE-certificate attribute contains an X.509 CPE certificate (see 7.6) issued by the CPE s manufacturer. The CPE s X.509 certificate is as defined in 6.3.2.3.9.2.

The SigCPE indicates an RSA signature over all the other attributes in this message, and the CPE's private key is used to make an RSA signature.

8.28.2 PKM RSA-Reply

Sent by the BS to a client CPE in response to a PKM RSA-Request message, the PKM RSA-Reply message contains an encrypted pre-PAK, the key s lifetime, and the key s sequence number. The pre-PAK shall be encrypted with the CPE s public key. The CPE_Random number is returned from the PKM RSA-Request message, along with a random number supplied by the BS, thus enabling assurance of key liveness.

Code: 4

Table 197 – PKM RSA-Reply attributes

Attribute	Contents
CPE_Random	A 64-bit random number generated in the CPE
BS_Random	A 64-bit random number generated in the BS
Encrypted pre-PAK	RSA-OAEP-Encrypt(PubKey(CPE), pre-PAK | CPE MAC Address)
Key Lifetime	PAK Aging timer
Key Sequence Number	PAK sequence number
BS_Certificate	Contains the BS's X.509 certificate
SigBS	An RSA signature over all the other attributes in the message

The SigBS indicates an RSA signature over all the other attributes in this message, and the BS's private key is used to make an RSA signature.

8.28.3 PKM RSA-Reject

The BS responds to a CPE's authorization request with an PKM RSA-Reject message if the BS rejects the CPE's authorization request.

Code: 5

Table 198 – PKM RSA-Reject attributes

Attribute	Contents
CPE_Random	A 64 bit random number generated in the CPE
BS_Random	A 64 bit random number generated in the BS
Error-Code	Error code identifying reason for rejection of authorization request
BS_Certificate	Contains the BS s X.509 certificate

Display-String (optional)	Display string providing reason for rejection of authorization request
SigBS	An RSA signature over all the other attributes in the message

The Error-Code and Display-String attributes describe to the requesting CPE the reason for the RSA-based authorization failure.

The SigBS indicates an RSA signature over all the other attributes in this message, and the BS's private key is used to make an RSA signature.

8.28.4 PKM RSA-Acknowledgement

The CPE sends the PKM RSA-Acknowledgement message to BS in response to a PKM RSA-Reply message or a PKM RSA-Reject message. Only if the value of Auth Result Code is failure, then the Error-Code and Display-String can be included in this message.

Code: 6

Table 199 – PKM RSA-Acknowledgement attributes

Attribute	Contents
BS_Random	A 64 bit random number generated in the BS
Auth Result Code	Indicates result (Success or Failure) of authorization procedure.
Error-Code	Error code identifying reason for rejection of authorization request
Display-String (optional)	Display string providing reason for rejection of authorization request
SigCPE	An RSA signature over all the other attributes in the message

The SigCPE indicates an RSA signature over all the other attributes in this message, and the CPE's private key is used to make an RSA signature.

8.28.5 PKM EAP Start

In the case of EAP re-authentication HMAC Digest/CMAC Digest and Key Sequence Number attributes shall be included. At initial EAP authentication, these attributes are omitted.

Code: 7

Table 200 – EAP-Start attributes

Attribute	Contents
Key Sequence Number	AK sequence number
HMAC Digest/CMAC Digest	Message digest calculated using AK

8.28.6 PKM EAP Transfer

When a CPE has an EAP payload received from an EAP method for transmission to the BS or when a BS has an EAP payload received from an EAP method for transmission to the CPE, it encapsulates it in a PKM EAP Transfer message. In the case of re-authentication, HMAC Digest/CMAC Digest and Key Sequence Number attributes shall be included.

Code: 8

Table 201 – PKM EAP Transfer attributes

Attribute	Contents
EAP Payload	Contains the EAP authentication data, not interpreted in the MAC
Key Sequence Number	AK sequence number
HMAC Digest/ CMAC Digest	Message Digest calculated using AK

The EAP Payload field carries data in the format described in section 4 of RFC 3748.

8.28.7 PKM Authenticated EAP Transfer

This message is used for Authenticated EAP-based authorization (if this was specified by Authorization Policy Support negotiated in the SBC-REQ/RSP exchange). Specifically, when a CPE or BS has an EAP payload received from an EAP method for transmission after an authentication established EIK, it encapsulates the EAP payload in a PKM Authenticated EAP Transfer message.

Code: 9

Table 202 – PKM Authenticated EAP message attributes

Attribute	Contents
Key Sequence Number	PAK Sequence Number (optional)
EAP Payload	Contains the EAP authentication data, not interpreted in the MAC
HMAC/CMAC Digest	Message Digest calculated using EIK

The EAP Payload field carries EAP data in the format described in RFC 3748

The CMAC-Digest s or HMAC Digest s attribute shall be the final attribute in the message s attribute list.

Inclusion of the CMAC or HMAC Digest allows the CPE and BS to cryptographically bind previous authorization and following EAP authentication by authenticating the EAP payload. The key for the CMAC Value or HMAC Digest is derived from the EIK.

PAK Sequence Number attribute carries PAK sequence number only if CPE and BS negotiate Authenticated EAP after RSA mode.

8.28.8 PKM SA-TEK-Challenge

The BS transmits the PKM SA-TEK-Challenge message as a first step in the 3-way SA-TEK handshake at initial network entry and at reauthorization. The BS shall send this message to the CPE after finishing authorization procedure(s) selected by the negotiated Authorization Policy Support included in the SBC-REQ/RSP messages. It identifies an AK to be used, and includes a random number challenge to be returned by the CPE in the PKM SA-TEK-Request message.

Code: 10

Table 203 – PKM SA-TEK-Challenge message attributes

Attribute	Contents
BS_Random	A freshly generated random number of 64 bits.
Key Sequence Number	AK sequence number
AKID	AKID of the AK (this is the AKID of the *new* AK in the case of re-authentication)
Key lifetime	PMK lifetime, this attribute shall include only follows EAP-based authorization or EAP-based re-authorization procedures
HMAC/CMAC Digest	Message authentication digest for this message

The HMAC Digest attribute or the CMAC-Digest attribute shall be the final attribute in the message's attribute list.

Inclusion of the HMAC Digest or the CMAC-Digest allows the CPE and BS to authenticate a PKM SA-TEK-Challenge message. The HMAC or the CMAC authentication keys are derived from the AK.

8.28.9 PKM SA-TEK-Request

The CPE transmits the PKM SA-TEK-Request message after receipt and successful HMAC Digest or CMAC value verification of an SA-Challenge tuple or PKM SA-TEK-Challenge message from the BS. The PKM SA-TEK-Request proves liveliness of the CPE and its possession of the AK to the BS. If this message is being generated during initial network entry, then it constitutes a request for SA-Descriptors identifying the primary and static SAs and GSAs the requesting CPE is authorized to access and their particular properties (e.g., type, cryptographic suite).

If this message is being generated following HO, then it constitutes a request for establishment (in the target BS) of TEKs, GTEKs and GKEKs for the CPE and renewal of active primary, static and dynamic SAs and associated SAIDs used by the CPE in its previous serving BS.

Code: 11

Table 204 – PKM SA-TEK-Request message attributes

Attribute	Contents
CPE_Random	A 64-bit number chosen by the CPE freshly for every new handshake.
BS_Random	The 64-bit random number used in the PKM SA-TEK-Challenge message or SA-Challenge Tuple
Key Sequence Number	AK sequence number
AKID	Identifies the AK that was used for protecting this message

Security-Capabilities	Describes requesting CPE s security capabilities
Security Negotiation Parameters	ConfirCPE requesting CPE s security capabilities (see 11.8.4)
PKM configuration settings	PKM configuration defined in 11.9.36.
HMAC/CMAC Digest	Message authentication digest for this message.

Receipt of a new BS Random value in SA-TEK-Challenge or SA-Challenge tuple indicates the beginning of a new handshake

8.28.10 PKM SA-TEK-Response

The BS transmits the PKM SA-TEK-Response message as a final step in the 3-way SA-TEK handshake.

Code: 12

Table 205 – PKM SA-TEK-Response message attributes

Attribute	Contents
CPE_Random	The number received from the CPE
BS_Random	The random number included in the PKM SA-TEK-Challenge message or SA-Challenge TLV.
Key Sequence Number	AK sequence number
AKID	This identifies the AK to the CPE that was used for protecting this message.
SA_TEK_Update	A compound TLV list each of which specifies an SA identifier (SAID) and additional properties of the SA that the CPE is authorized to access. This compound field may be present at the reentry only. For each active SA in previous serving BS, corresponding TEK, GTEK and GKEK parameters are included.
Frame Number	An absolute frame number in which the old PMK and all its associate AKs should be discarded.
(one or more) SA-Descriptor(s)	Each compound SA-Descriptor attribute specifies an SA identifier (SAID) and additional properties of the SA. This attribute is present at the initial net-work entry only.
Security Negotiation Parameters	ConfirCPE the authentication and message integrity parameters to be used (see 11.8.4)
HMAC Digest/CMAC Digest	Message authentication digest for this message.

8.28.11 PKM Key-Request

A CPE sends a PKM Key-Request message to the BS to request new TEK and TEK-related parameters (GTEK and GTEK-related parameters for the multicast or broadcast service) or GKEK and GKEK-related parameters for the multicast or broadcast service.

Code: 13

Table 206 – PKM Key Request attributes

Attribute	Contents
Key Sequence Number	AK sequence number
SAID	Security association identifier - GSAID for multicast or broadcast service
Nonce	A random number generated in a CPE
HMAC Digest/CMAC Digest	Message Digest calculated using AK

The HMAC Digest attribute or the CMAC-Digest attribute shall be the final attribute in the message's attribute list.

Inclusion of the HMAC Digest or the CMAC digest allows the CPE and BS to authenticate the PKM Key-Request message. The HMAC Digest or the CMAC-Digest's authentication key is derived from the AK.

8.28.12 PKM Key-Reply

The BS responds to a CPE's PKM Key-Request message with a PKM Key-Reply message.

Code: 14

Table 207 – PKM Key-Reply attributes

Attribute	Contents
Key Sequence Number	AK sequence number
SAID	Security association identifier - GSAID for multicast or broadcast service.
TEK-Parameters	Older generation of key parameters relevant to SAID - GTEK-Parameters for the multicast or broadcast service
TEK-Parameters	Newer generation of key parameters relevant to SAID
GKEK-Parameters	Older generation of GKEK-related parameters for multicast or broadcast service.
GKEK-Parameters	Newer generation of GKEK-related parameters for multicast or broadcast service.
Nonce	A same random number included in the PKM Key Request message
HMAC/CMAC Digest	Message Digest calculated using AK

The GKEK-Parameters attribute is a compound attribute containing all of the GKEK-related parameters corresponding to a GSAID. This would include the GKEK, the GKEK's remaining key lifetime, and the GKEK's key sequence number. The older generation of GKEK-Parameters is valid within the current lifetime and the newer generation of GKEK-Parameters is valid within the next lifetime.

The HMAC Digest or the CMAC-Digest attribute shall be the final attribute in the message's attribute list.

Inclusion of the HMAC Digest or the CMAC digest allows the CPE and BS to authenticate the PKM Key-Reply message. The HMAC Digest or the CMAC-Digest's authentication key is derived from the AK.

8.28.13 PKM Key-Reject

The BS responds to a CPE's PKM Key-Request message with a PKM Authorization-Reject message if the BS rejects the CPE's traffic keying material request.

Code: 15

Table 208 – PKM Key-Reject attributes

Attribute	Contents
Key Sequence Number	AK sequence number
SAID	Security association identifier
Error-Code	Error code identifying reason for rejection of the PKM Key-Request message
Display-String (optional)	Display string containing reason for the PKM Key-Request message
Nonce	A same random number included in the PKM Key Request message
HMAC/CMAC Digest	Message Digest calculated using AK

The HMAC Digest or the CMAC-Digest attribute shall be the final attribute in the message's attribute list.

Inclusion of the HMAC Digest or the CMAC digest allows the CPE and BS to authenticate the PKM Key-Reject message. The HMAC Digest or the CMAC-Digest's authentication key is derived from the AK.

8.28.14 PKM SA-Addition

This message is sent by the BS to the CPE to establish one or more additional SAs.

Code: 16

Table 209 – PKM SA-Addition attributes

Attribute	Contents
Key Sequence Number	AK sequence number
(one or more) SA-Descriptor(s)	Each compound SA-Descriptor attribute specifies an SA identifier (SAID) and additional properties of the SA
HMAC/CMAC Digest	Message Digest calculated using AK

The HMAC Digest or the CMAC-Digest attribute shall be the final attribute in the message's attribute list.

Inclusion of the HMAC Digest or the CMAC digest allows the CPE and BS to authenticate the PKM SA-Add message. The HMAC Digest or the CMAC-Digest's authentication key is derived from the AK.

8.28.15 PKM TEK-Invalid

The BS sends a PKM TEK-Invalid message to a client CPE if the BS determines that the CPE encrypted an upstream PDU with an invalid TEK (i.e., an SAID's TEK key sequence number), contained within the received packet's MAC header, is out of the BS's range of known, valid sequence numbers for that SAID.

Code: 17

Table 210 – PKM TEK-Invalid attributes

Attribute	Contents
Key Sequence Number	AK sequence number
SAID	Security Association Identifier
Error-Code	Error code identifying reason for PKM TEK-Invalid message
Display-String (optional)	Display string containing reason for the PKM TEK-Invalid message
HMAC/CMAC Digest	Message Digest calculated using AK

The HMAC Digest or the CMAC-Digest attribute shall be the final attribute in the message's attribute list.

Inclusion of the HMAC Digest or the CMAC digest allows the CPE and BS to authenticate the PKM SA-Add message. The HMAC Digest or the CMAC-Digest's authentication key is derived from the AK.

8.28.16 PKM Group-Key-Update-Command

This message is sent by BS to push the GTEK and/or GKEK parameters to CPEs served with the specific multicast service or broadcast service.

Code: 18

Table 211 – PKM Group Key update command attributes

Attribute	Contents
Key-Sequence-Number	AK sequence number
GSAID	Security Association ID
Key Push Modes	Usage code of Key Update Command message.
Key Push Counter	Counter one greater than that of older generation.
GTEK-Parameters	Newer generation of GTEK-related parameters relevant to GSAID. The GTEK-Parameters is the TEK-Parameters for multicast or broadcast service.
GKEK-Parameters	Newer generation of GKEK-related parameters for multicast or broadcast service.
HMAC/CMAC Digest	Message integrity code of this message.

Key Sequence Number is the sequence number of the synchronized AK (Authorization Key) between a CPE and a BS.

GSAID is SAID for the multicast group or the broadcast group. The type and length of the GSAID is equal to ones of the SAID.

There are two types in a PKM Group Key Update Command message, GKEK update mode and GTEK update mode. The former is used to update GKEK and the latter is used to update GTEK for the multicast service or the broadcast service. Key Push Modes indicates the usage code of a PKM Group Key Update Command message.

The PKM Group Key Update Command message for the GKEK update mode is carried on the Primary Management connection, but one for the GTEK update mode is carried on the Broadcast connection. A few attributes in a PKM Group Key Update Command message shall not be used according this Key Push Modes attribute's value.

Key Push Counter is used to protect for replay attack. This value is one greater than that of older generation.

A PKM Group Key Update Command message contains only newer generation of key parameters, because this message inform a CPE next traffic key material. The GTEK-Parameters attribute is a com-pound attribute containing all of the keying material corresponding to a newer generation of a GSAID's GTEK. This would include the GTEK, the GTEK's remaining key lifetime, the GTEK's key sequence number, the associated GKEK sequence number, and the cipher block chaining (CBC) initialization vector. The GTEK is TEK for the multicast group or the broadcast group. The type and length of the GTEK is equal to ones of the TEK. The GKEK (Group Key Encryption Key) can be randomly generated from a BS or a certain network node (i.e., an ASA server). The GKEK should be identically shared within the same multicast group or the broadcast group. The GTEK is encrypted with GKEK for the multicast service or the broadcast service. GKEK parameters contain the GKEK encrypted by the KEK, GKEK sequence number, and GKEK lifetime.

The HMAC/CMAC Digest attribute shall be the final attribute in the message's attribute list. Inclusion of the keyed digest allows the receiving client to authenticate the Group Key Update Command message. The HMAC/CMAC Digest's authentication key is derived from the AK for the GKEK update mode and GKEK for the GTEK update mode.

8.28.17 PKM EAP Complete

In double EAP mode (EAP after EAP), BS sends the PKM EAP Complete message to CPE with EAP-Success to inform CPE of completing 1st EAP conversation.

This message is used only if CPE and BS negotiate EAP in EAP mode.

The Key Sequence Number and HMAC/CMAC Digest attributes of this message appear only in re-authentication.

Code: 19

Table 212 – PKM EAP Complete attributes

Attribute	Contents
EAP Payload	Contains the EAP authentication data, not interpreted in the MAC layer
Key Sequence Number	AK sequence number appear only if AK is available from previous double EAP
HMAC Digest/CMAC Digest	Message Digest calculated using AK only if AK is available from previ-ous double EAP Message Digest calculated using EIK when initial authentication

8.28.18 PKM Authenticated EAP Start

In double EAP mode (EAP after EAP), CPE sends the PKM EAP Authenticated EAP Start message to BS in order to initiate 2nd round EAP. This message is signed by EIK which is generated by 1st EAP.

This message is used only for initial authentication of double EAP.

Code: 20

Table 213 – PKM Authenticated EAP Start attribute

Attribute	Contents
CPE_Random	Random number generated by CPE.
HMAC Digest/CMAC Digest	Message Digest calculated using EIK

8.28.19 PKM Authentication Invalid

The BS may send an Authorization Invalid message to a client CPE as:

1. An unsolicited indication, or
2. A response to a message received from that CPE.

In either case, the Authorization Invalid message instructs the receiving CPE to reauthorize with its BS.

The BS sends an Authorization Invalid in response to a Key Request if (1) the BS does not recognize the CPE as being authorized (i.e., no valid AK associated with the requesting CPE) or (2) verification of the Key Request's keyed message digest (in HMAC-Digest attribute) failed, indicating a loss of AK synchronization between CPE and BS.

Code: 21

Table 214 – Authorization Invalid attributes

Attribute	Contents
Error-Code	Error code identifying reason for Authorization Invalid.
Display-String (optional)	Display String describing failure condition.

8.28.20 PKM Authentication Information (Auth Info)

The Auth Info message contains a single CA-Certificate attribute, containing an X.509 CA certificate for the manufacturer of the CPE. The CPE's X.509 user certificate shall have been issued by the CA identified by the X.509 CA certificate.

Auth Info messages are strictly informative; while the CPE shall transmit Auth Info messages as indicated by the Authentication state model, the BS may ignore them.

Code: 22

Table 215 – Auth Info attributes

Attribute	Contents
CA-Certificate	Certificate of manufacturer CA that issued CPE certificate.

The CA-certificate attribute contains an X.509 CA certificate for the CA that issued the CPE's X.509 user certificate. The external CA issues these CA certificates to CPE manufacturers.

8.29 AAS Channel Feedback Request/Response (AAS-CFB-REQ/RSP)

The AAS Channel Feedback Request message shall be used by a system supporting AAS. This message serves to request channel measurement that will help in adjusting the direction of the adaptive array. See Table 216.

Table 216 – Message format

Syntax	Size	Notes
AAS-CFB-REQ_Message_Format() {		
Management Message Type = 58	8 bits	
Message Body	*Variable*	See Table 217
}		

Table 217 – Message format

Syntax	Size	Notes
AAS-CFB-REQ_Message_Body() {		
Frame Number	8 bits	The least significant 8 bits of the frame number in which to start the measurement.
Number of Frames	7 bits	The number of frames over which to measure.
Measurement Data Type	1 bit	Increases every time an AAS-CFB-REQ is sent to the CPE. Individual counters shall be maintained for each CPE. The value 0 shall not be used. • 0 = measure on downstream preamble only • 1 = measure on downstream data (for this CPE) only
Feedback Request Counter	3 bits	Indicates the frequency measurement points to report on. Measurement points shall be on the frequencies corresponding to the negative subcarrier offset indices $-N_{used}/2 + n$ times the indicated subcarrier resolution and corresponding to the positive subcarrier offset indices $N_{used}/2 - n \times$ the indicated subcarrier resolution, where n is a positive integer.
Frequency Measurement Resolution	2 bits	0b00 = 32 subcarriers 0b01 = 64 subcarriers 0b10 = 128 subcarriers 0b11 = 256 subcarriers
reserved	3 bits	Shall be set to zero
}		

The AAS Channel Feedback Response message shall be sent as a response to the AAS-CFB-REQ message after the indicated measurement period has expired. See Table 218.

Table 218 – Message format

Syntax	Size	Notes
AAS-CFB-RSP_Message_Format() {		
Management Message Type = 59	8 bits	
Message Body	*Variable*	See Table 219
}		

Table 219 – Message format

Syntax	Size	Notes
AAS-CFB-RSP_Message_Body() {		
Reserved	2 bits	Shall be set to zero

Measurement Data Type	1 bit	See Table 217
Feedback Request Counter	3 bits	Counter from the AAS-CFB-REQ messages to which this is the response. The value 0 indicates that the response is unsolicited. In this case, the measurement corresponds to the preceding frame.
Frequency Measurement Resolution	2 bits	See Table 217
for (i = 0; i < Number of Frequencies; i++) {		
Re(Frequency_value[i])	8 bits	The real (Re) and imaginary (Im) part of the measured amplitude on the frequency measurement point (low to high frequency) in signed integer fixed point format ([±][2 bits].[5 bits]).
Im(Frequency_value[i])	8 bits	
}		
RSSI Mean Value	8 bits	The mean RSSI as measured on the element pointed to by data measurement type, frame number and number of frames in the corresponding request. The RSSI is quantized as described in the PHY spec. When the AAS feedback response is unsolicited, this value corresponds to preceding frame.
CINR Mean Value	8 bits	The mean CINR as measured on the element pointed to by data measurement type, frame number, and number of frames in the corresponding request. The CINR is quantized as described in the PHY spec. When the AAS feedback response is unsolicited, this value corresponds to preceding frame.
}		

9. Management of MAC PDUs

9.1 Conventions

Data shall be transmitted in accordance with the following rules:

1. Data fields of messages are transmitted in the same order as they appear in the corresponding tables and figures in this proposal.
2. Data fields messages which are specified as binary numbers, are transmitted as a sequence of their binary digits, starting from MSB (here, bit masks are also considered numerical fields). For signed numbers, MSB is allocated for the sign. Length field in the "definite form" of ITU-T X.690 is also considered a numerical field.
3. Data fields specified as SDUs or SDU fragments (e.g., MAC PDU payloads) are transmitted in the same order of bytes as received from upper layers.
4. Data fields specified as strings are transmitted in the order of symbols in the string.

In (3) and (4), bits within a byte are transmitted in the order "MSB first."

9.2 Concatenation

Multiple MAC PDUs may be concatenated into a single transmission in either the upstream or downstream directions, as depicted in Figure 9 for an upstream burst transmission. Since each MAC PDU is identified by a unique CID, the receiving MAC entity is able to present the MAC SDU (after reassembling the MAC SDU from one or more received MAC PDUs) to the correct instance of the MAC SAP. MAC PDUs containing management messages or user data may be concatenated into the same transmission.

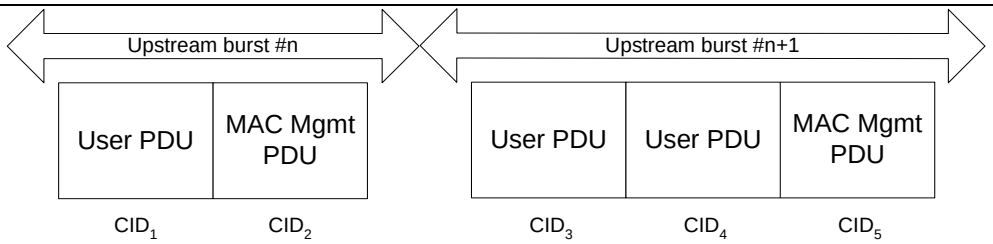

Figure 9 – Concatenation of MAC PDUs

9.3 Fragmentation

Fragmentation and packing (discussed in 9.4) are features proposed to be mandatory. Fragmentation is the process by which a MAC SDU is divided into one or more MAC PDUs. This process is undertaken to allow efficient use of available bandwidth relative to the QoS requirements of a connection's service flow. Upon the creation of a connection by the MAC SAP, fragmentation capability is defined. Fragmentation may be initiated by a BS for downstream connections and by a CPE for upstream connections.

Fragments are tagged with their position in their parent SDU in accordance with Table 220.

Table 220 – Fragmentation rules

Fragment	Fragmentation Control
First Fragment	10
Continuing Fragment	11
Last Fragment	01
Unfragmented	00

9.3.1 Non-ARQ Connections

For non-ARQ connections, fragments are transmitted once and in sequence. The sequence number assigned to each fragment allows the receiver to recreate the original payload and to detect the loss of any intermediate packets. A connection may be in only one fragmentation state at any given time.

Upon loss, the receiver shall discard all MAC PDUs on the connection until a new first fragment or a non-fragmented MAC PDU is detected.

9.3.2 ARQ-Enabled Connections

For ARQ-enabled connections, fragments are formed for each transmission by concatenating sets of ARQ blocks with adjacent sequence numbers (see 10). The BSN value carried in the fragmentation subheader is the BSN for the first ARQ block appearing in the segment.

9.4 Packing

In CMAC, the transmitting side has full discretion whether or not to pack a group of MAC SDUs into a single MAC PDU. BSs and CPEs shall both have the capability of unpacking. If packing is turned on for a connection,

the MAC may pack multiple MAC SDUs into a single MAC PDU. Also, packing makes use of the connection attribute indicating whether the connection carries fixed-length or variable-length packets.

The construction of PDUs varies for ARQ and non-ARQ connections with respect to packing and fragmentation syntax. The packing and fragmentation mechanisms for both the ARQ and non-ARQ connections are specified in the subsections below.

9.4.1 Non-ARQ Connections

For connections that do not utilize ARQ, the packing procedure described in 805512728.1332372.59.4.1.1 may be used when the connection carries fixed-length MAC SDUs. For all other non-ARQ connections, the variable length packing algorithm described in 805502985.512.59.4.1.2 shall be employed. Please refer to 8.8.10.15 for more information on the indication whether a connection carries fixed-length or variable-length SDUs.

9.4.1.1 Fixed-length MAC SDUs

For packing with fixed-length blocks, the Request/Transmission Policy (8.8.10.12) shall be set to allow packing and prohibit fragmentation, and the SDU size (8.8.10.16) shall be included in DSA-REQ message when establishing the connection. The length field of the MAC header implicitly indicates the number of MAC SDUs packed into a single MAC PDU. If the MAC SDU size is *n* bytes, the receiving side can unpack simply by knowing that the length field in the MAC header will be $n \times k + j$, where *k* is the number of MAC SDUs packed into the MAC PDU and *j* is the size of the MAC header and any prepended MAC subheaders. A MAC PDU containing a packed sequence of fixed-length MAC SDUs would be constructed as in Figure 10. Note that there is no added overhead due to packing in the fixed-length MAC SDU case.

9.4.1.2 Variable-length MAC SDUs

When packing variable-length SDU connections (e.g., 802.3/Ethernet), the $n \times k + j$ relationship between the MAC header's length field and the higher-layer MAC SDUs no longer holds. Therefore, it is necessary to indicate where one MAC SDU ends and another begins. In the variable-length MAC SDU case, the MAC attaches a Packing subheader (PSH) to each MAC SDU. This subheader is described in 6.1.3.

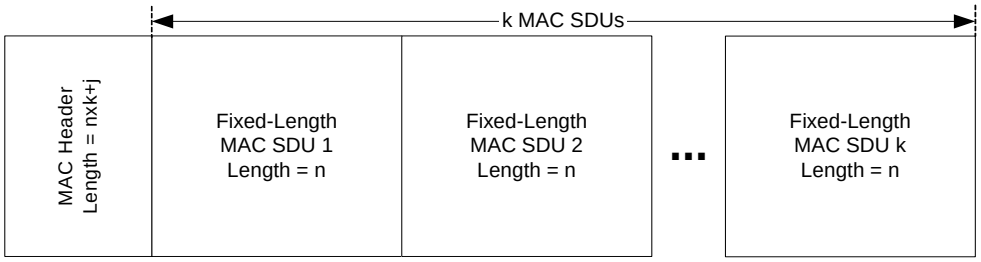

Figure 10 – Packing fixed-length MAC SDUs into a single MAC PDU

A MAC PDU containing a packed sequence of variable-length MAC SDUs is constructed as shown in Figure 11. If more than one MAC SDU is packed into the MAC PDU, the type field in the MAC header indicates the presence of PSHs. Note that unfragmented MAC SDUs and MAC SDU fragments may both be present in the same MAC PDU (see Figure 12).

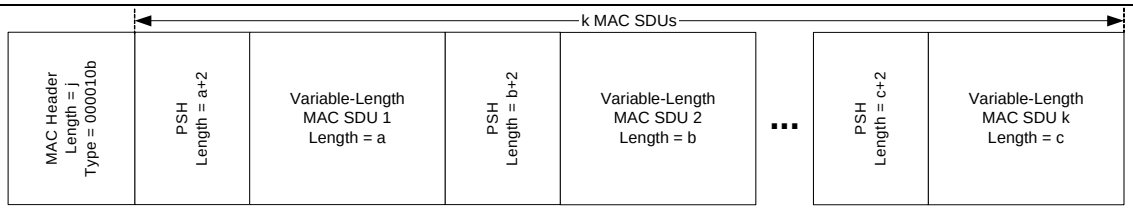

Figure 11 – Packing variable-length MAC SDUs into a single MAC PDU

Simultaneous fragmentation and packing allows efficient use of the wireless channel, but requires guidelines to be followed so it is clear which MAC SDU is currently in a state of fragmentation. To accomplish this, when a PSH is present, the fragmentation information for individual MAC SDUs or MAC SDU fragments is contained in the corresponding PSH. If no PSH is present, the fragmentation information for individual MAC SDU fragments is contained in the corresponding Fragmentation subheader (FSH). This procedure is shown in Figure 12.

Finally, note that while it is legal to have continuation fragments packed with other fragments, the circumstances for creating continuation fragments would preclude this from happening.

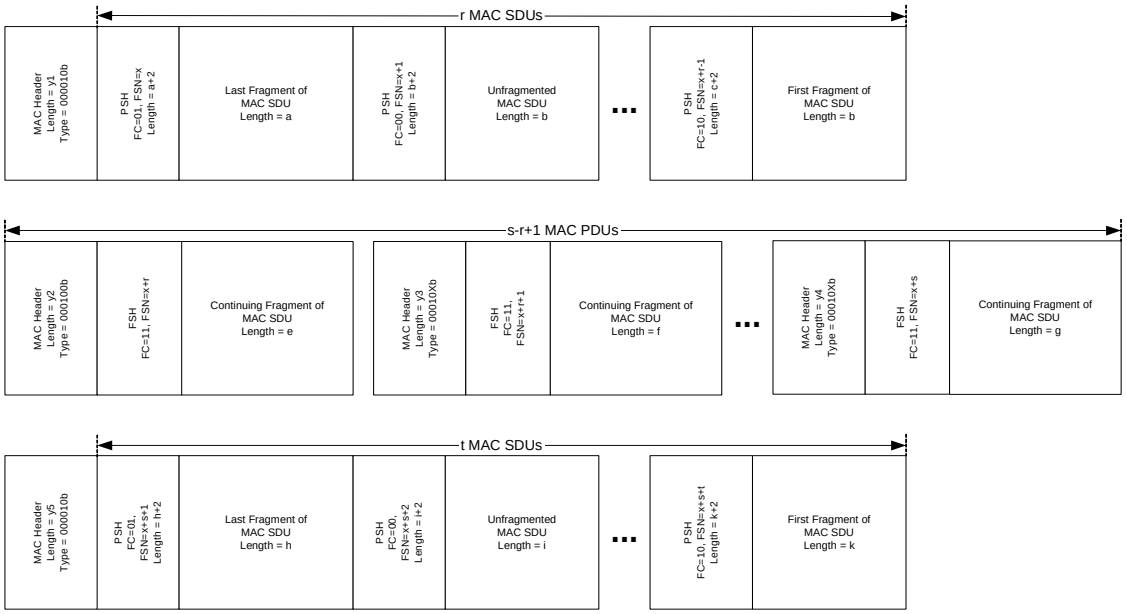

Figure 12 – Packing with fragmentation

9.4.2 ARQ-Enabled Connections

The use of PSH for ARQ-enabled connections is similar to that for non-ARQ connections as described in 9.512.59.4.1.2, except that ARQ-enabled connections shall set the Extended Type bit (Table 7) in the generic MAC header to 1. The packing of variable-length MAC SDUs for the ARQ-enabled connections is similar to that of non-ARQ connections, when fragmentation is enabled. The BSN of the PSH shall be used by the ARQ protocol to identify and retransmit lost fragments.

For ARQ-enabled connections, when the type field indicates that PSHs are in use, fragmentation information for each individual MAC SDU or MAC SDU fragment is contained in the associated PSH. When the type field indicates that packing is not in use, fragmentation information for the MAC PDU's single payload (MAC SDU or MAC SDU fragment) is contained in the FSH appearing in the message. Figure 13 illustrates the use of FSH without packing.

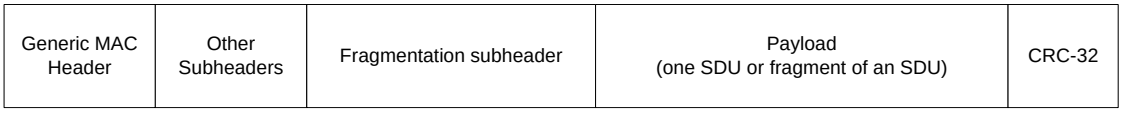

Figure 13 – Example of a MAC PDU with extended fragmentation subheader

Figure 14 depicts the structure of a MAC PDU with ARQ PSHs. Each of the packed MAC SDU or MAC SDU fragments or ARQ feedback payload requires its own PSH, and some of them may be transmissions while others are retransmissions.

A MAC SDU may be partitioned into multiple fragments that are then packed into the same MAC PDU for the first transmission. MAC PDUs may have fragments from the same or different SDUs, including a mix of first transmissions and retransmissions. The 11-bit BSN and 2-bit FC fields uniquely identify each fragment or non-fragmented SDU.

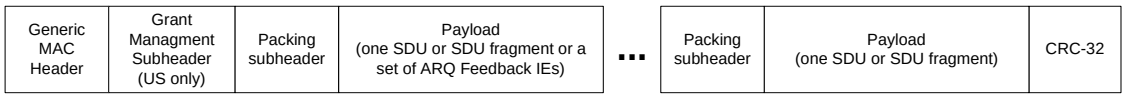

Figure 14 – Example of a MAC PDU with ARQ packing subheader

9.4.3 ARQ Feedback IEs

The ARQ Feedback Payload (see Table 221) may be sent on an ARQ or non-ARQ connection, and consists of one or more ARQ Feedback IEs (see 10). However, policies based on implementation and/or QoS constraints may restrict the use of certain connections for transporting ARQ Feedback Payload. The ARQ Feedback Payload is treated like any other payload (SDU or fragments) from the packing perspective, except that only one ARQ Feedback Payload shall be present within a single MAC PDU.

The presence of an ARQ Feedback Payload in a MAC PDU is indicated by the value of the ARQ Feedback Payload bit in the Type field (see Table 7) in the generic MAC header. When present, the first packed payload shall be the ARQ Feedback Payload. The PSH preceding the ARQ Feedback Payload indicates the total length of the payload including the PSH and all ARQ Feedback IEs within the payload. The FSN/BSN field of the PSH shall be ignored for the ARQ Feedback Payload and the FC bits shall be set to 00.

Table 221 – Message format

Syntax	Size	Notes
ARQ_Feedback_Payload_Format() {		
do		
ARQ_Feedback_IE(last)	*variable*	Include as many as needed until last == TRUE. See 10.
until(last)		
}		

9.5 CRC Calculation

CRC shall be calculated as defined in IEEE Std 802.3. The CRC shall be appended to the payload of the MAC PDU. MAC PDUs which do not contain a payload may choose not to use CRC and be unprotected. The CRC shall cover the MAC header and the Payload of the MAC PDU. In addition, the CRC shall be calculated after encryption, that is, it protects both the Header and the ciphered Payload.

9.6 Padding

Allocated space within a data burst that is unused shall be initialized to a known state. This may be accomplished by setting each unused byte to the stuff byte value (0xFF). If the size of the unused region is at least the size of a MAC header, the unused space may also be initialized as a MAC PDU. In this case, the MAC header CID field shall be set to the value of the Padding CID (see Table 227), the UCS, EC, and Type fields shall be set to zero, the length field shall be set to the number of unused bytes (including the size of the MAC header created for the padding MAC PDU) in the data burst, and the HCS shall be computed in the normal way.

10. The ARQ Mechanism

CMAC adopts the same ARQ mechanism as specified in IEEE 802.16 (IEEE Std 802.16^{TM}-2004). For further information, please refer to [5].

11. Scheduling Services

Scheduling services represent the data handling mechanisms supported by the MAC scheduler for data transport on a connection. Each connection is associated with a single data service. Each data service is associated with a set of QoS parameters that quantify aspects of its behavior (these parameters are managed using the DSA and DSC messages). Four services (8.8.10.11) are supported: Unsolicited Grant Service (UGS), Real-time Polling Service (rtPS), Non-real-time Polling Service (nrtPS), and Best Effort (BE). Below we provide a description of each of these services and some of the applications they aim at supporting. Mandatory QoS parameters associated with each of these services are also identified. A detailed description of all supported QoS parameters can be found in Section 8.8.10.

The UGS is designed to support real-time data streams consisting of fixed-size data packets sent at periodic intervals, such as T1/E1 and Voice over IP without silence suppression. The mandatory QoS service flow parameters for this scheduling service are Maximum Sustained Traffic Rate, Maximum Latency, Tolerated Jitter, and Request/Transmission Policy. If present, the Minimum Reserved Traffic Rate parameter shall have the same value as the Maximum Sustained Traffic Rate parameter.

The rtPS is designed to support real-time data streams consisting of variable-sized data packets that are issued at periodic intervals, such as MPEG video. The mandatory QoS service flow parameters for this scheduling service are Minimum Reserved Traffic Rate, Maximum Sustained Traffic Rate, Maximum Latency, and Request/Transmission Policy.

The nrtPS is designed to support delay-tolerant data streams consisting of variable-sized data packets for which a minimum data rate is required, such as FTP. The mandatory QoS service flow parameters for this scheduling service are Minimum Reserved Traffic Rate, Maximum Sustained Traffic Rate, Traffic Priority, and Request/Transmission Policy.

The BE service is designed to support data streams for which no minimum service level is required and therefore may be handled on a space-available basis. The mandatory QoS service flow parameters for this scheduling service are Maximum Sustained Traffic Rate, Traffic Priority, and Request/Transmission Policy.

11.1 Data Transmission Scheduling

Data transmission scheduling selects the data for transmission in a particular frame/bandwidth allocation and is performed by the BS for downstream, and by the CPE for upstream. In addition to whatever other factors the scheduler may deem pertinent, the following items shall be taken into account for each active service flow:

- Coexistence with other overlapping 802.22 cells (through "interference-free" scheduling and traffic constraints – see Section 8.23)
- The scheduling service specified for the service flow.
- The values assigned to the service flow's QoS parameters.
- The availability of data for transmission.
- The capacity of the granted bandwidth.

11.2 Upstream Request/Grant Scheduling

Upstream request/grant scheduling is performed by the BS with the intent of providing each associated CPE with bandwidth for upstream transmissions or opportunities to request bandwidth. By specifying a scheduling service and its associated QoS parameters, the BS scheduler can anticipate the throughput and latency needs of the upstream traffic and provide polls and/or grants at the appropriate times. Table 222 summarizes the scheduling services and the poll/grant options available for each. The following subsections define service flow scheduling services for upstream operations.

Table 222 – Scheduling services and corresponding poll/grant options

Scheduling Type	PiggyBack Request	Bandwidth Stealing	Polling
UGS	Not Allowed	Not Allowed	PM bit is used to request a unicast poll for bandwidth needs of non-UGS connections.
rtPS	Allowed	Allowed	Scheduling only allows unicast polling.
nrtPS	Allowed	Allowed	Scheduling may restrict a service flow to unicast polling via the transmission/request policy; otherwise all forms of polling are allowed.
BE	Allowed	Allowed	All forms of polling allowed.

11.2.1 UGS

The UGS service offers fixed-size grants on a real-time periodic basis, which eliminate the overhead and latency of CPE requests and assure that grants are available to meet the flow's real-time needs. The BS shall provide Data Grant Burst IEs to the CPE at periodic intervals based upon the Maximum Sustained Traffic Rate of the service flow. The size of these grants shall be sufficient to hold the fixed-length data associated with the service flow (with associated generic MAC header and Grant management subheader) but may be larger at the discretion of the BS scheduler. In order for this service to work correctly, the Request/Transmission Policy (see 8.8.10.12) setting shall be such that the CPE is prohibited from using any contention request opportunities for this connection.

The Grant Management subheader (805502985.512.56.1.3.3) is used to pass status information from the CPE to the BS regarding the state of the UGS service flow. The most significant bit of the Grant Management field is the Slip Indicator (SI) bit. The CPE shall set this flag once it detects that this service flow has exceeded its transmit queue depth. Once the CPE detects that the service flow's transmit queue is back within limits, it shall clear the SI flag. The flag allows the BS to provide for long term compensation for conditions, such as lost maps or clock rate mismatches, by issuing additional grants. The poll-me (PM) bit may be used to request to be polled for a different, non-UGS connection.

The BS shall not allocate more bandwidth than the Maximum Sustained Traffic Rate parameter of the Active QoS Parameter Set, excluding the case when the SI bit of the Grant Management field is set. In this case, the BS may grant up to 1% additional bandwidth for clock rate mismatch compensation.

11.2.2 rtPS

The rtPS service offers real-time, periodic, unicast request opportunities, which meet the flow's real-time needs and allows the CPE to specify the size of the desired grant. This service requires more request overhead than UGS, but supports variable grant sizes for optimum data transport efficiency.

The BS shall provide periodic unicast request opportunities. In order for this service to work correctly, the Request/Transmission Policy setting (see 8.8.10.12) shall be such that the CPE is prohibited from using any contention request opportunities for that connection. The BS may issue unicast request opportunities as prescribed by this service even if prior requests are currently unfulfilled. This results in the CPE using only unicast request opportunities in order to obtain upstream transmission opportunities (the CPE could still use unsolicited Data Grant Burst Types for upstream transmission as well). All other bits of the Request/Transmission Policy are irrelevant to the fundamental operation of this scheduling service and should be set according to network policy.

11.2.3 nrtPS

The nrtPS offers unicast polls on a regular basis, which assures that the service flow receives request opportunities even during network congestion. The BS typically polls nrtPS CIDs on an interval on the order of one second or less.

The BS shall provide timely unicast request opportunities. In order for this service to work correctly, the Request/Transmission Policy setting (see 8.8.10.12) shall be set such that the CPE is allowed to use contention request opportunities. This results in the CPE using contention request opportunities as well as unicast request opportunities and unsolicited Data Grant Burst Types. All other bits of the Request/Transmission Policy are irrelevant to the fundamental operation of this scheduling service and should be set according to network policy.

11.2.4 BE

The intent of the BE service is to provide efficient service for best effort traffic. In order for this service to work correctly, the Request/Transmission Policy setting shall be set such that the CPE is allowed to use contention request opportunities. This results in the CPE using contention request opportunities as well as unicast request opportunities and unsolicited Data Grant Burst Types. All other bits of the Request/Transmission Policy are irrelevant to the fundamental operation of this scheduling service and should be set according to network policy.

12. Bandwidth Management

During network entry and initialization, every CPE is assigned up to three dedicated CIDs for the purpose of sending and receiving control messages. These connection pairs are used to allow differentiated levels of QoS to be applied to the different connections carrying MAC management traffic. Increasing (or decreasing) bandwidth requirements is necessary for all services except incompressible constant bit rate UGS connections. The needs of incompressible UGS connections do not change between connection establishment and termination. The requirements of compressible UGS connections, 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.

There are numerous methods by which a CPE can get a bandwidth request message to the BS, and these are described in the following sections.

12.1 Bandwidth Requests

Bandwidth Requests (or simply, Requests) refer to the mechanism that CPEs use to indicate to the BS that they need upstream bandwidth allocation. Two types of bandwidth requests are available in CMAC (with proper PHY support).

12.1.1 Contention-based Request

In this case, a Request comes as a subheader appended to the general MAC header (see 805512728.1332372.56.1.3.1), which may or may not contain payload. Typically, a Request will not contain a payload if it is the first Request made for the connection. It may contain a payload otherwise.

For self-coexistence purposes, a Bandwidth Request may be followed by one or more IE that indicate any traffic constraints a CPE may have (see 8.23.2.1). These traffic constraints are the result of any coexistence beacons received by the CPE, which uses this information when requesting upstream bandwidth allocation. The scheduler (which is implementation dependent) at the BS shall do its best to allocate upstream bandwidth to the CPE that respects the CPE's traffic constraints.

Because the upstream burst profile can change dynamically, all requests for bandwidth shall be made in terms of the number of bytes needed to carry the MAC header and payload, but not the PHY overhead. The Bandwidth Request message may be transmitted during any upstream allocation, except during any initial ranging interval, UCS notification slots, and SSS.

Bandwidth Requests may be incremental or aggregate. When the BS receives an incremental Bandwidth Request, it shall add the quantity of bandwidth requested to its current perception of the bandwidth needs of the connection. When the BS receives an aggregate Bandwidth Request, it shall replace its perception of the bandwidth needs of the connection with the quantity of bandwidth requested. The Type field in the bandwidth request header indicates whether the request is incremental or aggregate. The self-correcting nature of the request/grant protocol requires that CPEs shall periodically use aggregate Bandwidth Requests. The period may be a function of the QoS of a service and of the link quality. Due to the possibility of collisions, Bandwidth Requests transmitted in broadcast or multicast Request IEs should be aggregate requests.

12.1.2 Contention-based CDMA Request

In addition to the transmission of bandwidth requests by the CPE, the PHY also supports the use a CDMA-based mechanism for the purpose of upstream bandwidth allocation.

As detailed in the PHY spec, the PHY has available a subset of Ranging codes that shall be used for contention-based CDMA Bandwidth Requests. The CPE, upon needing to request bandwidth, shall select, with equal probability, a Ranging Code from the code subset allocated to Bandwidth Requests. This Ranging Code shall be modulated onto a Ranging Subchannel and transmitted during the appropriate upstream allocation. The Ranging Subchannel shall be selected amongst the ones reserved by the MAC for the upstream transmission.

Upon detection, the BS shall provide (an implementation dependent) upstream allocation for the CPE, but instead of indicating a Basic CID, the broadcast CID shall be sent in combination with a CDMA_Allocation_IE, which specifies the transmit region and Code that were used by the CPE. This allows a CPE to determine whether it has been given an allocation by matching these parameters with the parameters it used. The CPE shall use the

allocation to transmit a MAC PDU with the bandwidth request subheader and/or data (this is indicated by the Usage field – see Table 48). The CPE may only omit the Bandwidth Request PDU when the BS indicated so in the CDMA_Allocation_IE (see Table 48).

If the BS does not issue the upstream allocation described above, or the MAC PDU with the bandwidth request subheader does not result in a subsequent allocation of any bandwidth, the CPE shall assume that the Ranging Code transmission resulted in a collision and follow the contention resolution as specified in 14.

12.2 Grants

For a CPE, bandwidth requests reference individual connections while each bandwidth grant is addressed to the CPE's Basic CID, not to individual CIDs. Since it is nondeterministic which request is being honored, when the CPE receives a shorter transmission opportunity than expected (scheduler decision, request message lost, etc.), no explicit reason is given. In all cases, based on the latest information received from the BS and the status of the request, the CPE may decide to perform backoff and request again or to discard the SDU.

A CPE may use Request IEs (that are broadcast) directed at a multicast polling group it is a member of or directed at its Basic CID. In all cases, the Request IE burst profile is used, even if the BS is capable of receiving the CPE with a more efficient burst profile. To take advantage of a more efficient burst profile, the CPE should transmit in an interval defined by a Data Grant IE directed at its Basic CID. Because of this, unicast polling of a CPE would normally be done by allocating a Data Grant IE directed at its Basic CID. Also note that, in a Data Grant IE directed at its Basic CID, the CPE may make bandwidth requests for any of its connections.

The procedure followed by CPEs is shown in Figure 15. Note that it is the CPE's local scheduler which decides which connections get the granted bandwidth.

12.3 Polling

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 connections and/or CPEs actually define bandwidth request contention IEs. The allocations are not in the form of an explicit message, but are contained as a series of IEs within the US-MAP.

Note that polling is done on CPE basis. Bandwidth is always requested on a CID basis and bandwidth is allocated on a CPE basis.

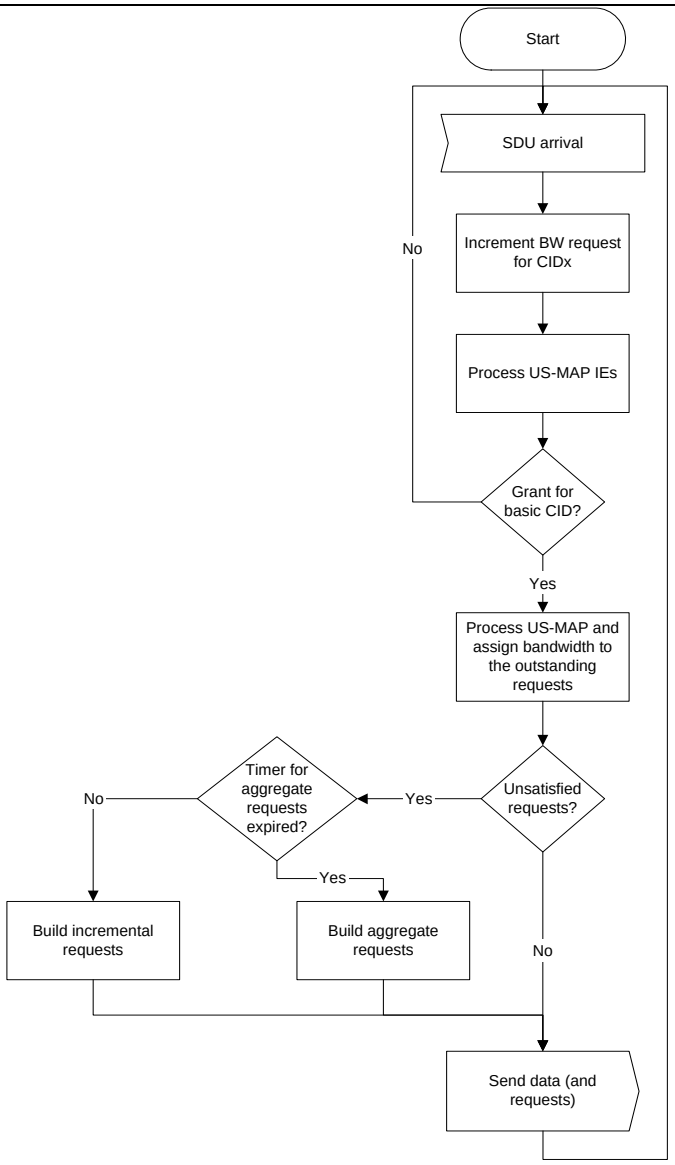

Figure 15 – Request/grant mechanism

12.3.1 Unicast

When a CPE is polled individually, no explicit message is transmitted to poll the CPE. Rather, the BS reserves in the US-MAP enough bandwidth for the CPE to respond with a Bandwidth Request. If the CPE does not need bandwidth, the allocation is padded in accordance with 9.6. CPEs that have an active UGS connection of sufficient bandwidth shall not be polled individually unless they set the PM bit in the header of a packet on the UGS connection. This saves bandwidth over polling all CPEs individually. Note that unicast polling would normally be done on a per-CPE basis by allocating a Data Grant IE directed at its Basic CID.

12.3.2 Multicast and Broadcast

If insufficient bandwidth is available to individually poll many inactive CPEs, some CPEs may be polled in multicast groups or a broadcast poll may be issued. Certain CIDs are reserved for multicast groups and for

broadcast messages, as described in Table 227. As with individual polling, the poll is not an explicit message, but bandwidth allocated in the US-MAP. The difference is that, rather than associating allocated bandwidth with a CPE's Basic CID, the allocation is to a multicast or broadcast CID.

When the poll is directed at a multicast or broadcast CID, a CPE belonging to the polled group may request bandwidth during any request interval allocated to that CID in the US-MAP by a Request IE. In order to reduce the likelihood of collision with multicast and broadcast polling, only CPEs needing bandwidth reply; they shall apply the contention resolution algorithm as defined in Section 14 to select the slot in which to transmit the initial bandwidth request. Zero-length bandwidth requests shall not be used in multicast or broadcast Request Intervals.

The CPE shall assume that the transmission has been unsuccessful if no grant has been received in the number of subsequent US-MAP messages specified by the parameter Contention-based reservation timeout (see 8.3.1). If the retransmission of the request is made in a multicast or broadcast opportunity, the CPE continues to run the contention resolution algorithm in Section 14. Note that the CPE is not restricted to retransmitting the request in a multicast or broadcast Request Interval only.

12.3.3 PM Bit

CPEs with currently active UGS connections may set the PM bit (see 805502985.512.56.1.3.3) in a MAC packet of the UGS connection to indicate to the BS that they need to be polled to request bandwidth for non-UGS connections. To reduce the bandwidth requirements of individual polling, CPEs with active UGS connections need be individually polled only if the PM bit is set (or if the interval of the UGS is too long to satisfy the QoS of the CPE's other connections). 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 16. To minimize the risk of the BS missing the PM bit, the CPE may set the bit in all UGS MAC Grant Management subheaders in the upstream scheduling interval.

13. PHY Support

In this section, we discuss aspects of CMAC that may impact the design of the PHY layer.

13.1 Duplexing

Given that CMAC overcomes some major limitations of TDD while supporting some aspects of FDD (through its proposed architecture), TDD is the duplexing mode of choice in this proposal (see Section 4 for a more comprehensive discussion on this issue).

In TDD, the upstream and downstream transmissions occur at different times and usually share the same frequency. As depicted in Figure 17, a TDD frame in CMAC typically has a fixed duration and contains two subframes: a predominantly downstream and an upstream (see 4 for further details). The frame is divided into an integer number of MAC slots, which help to partition the bandwidth easily. The TDD framing is adaptive in that the bandwidth allocated to the downstream versus the upstream can vary. The split between upstream and downstream is a parameter that is controlled at higher layers within the system.

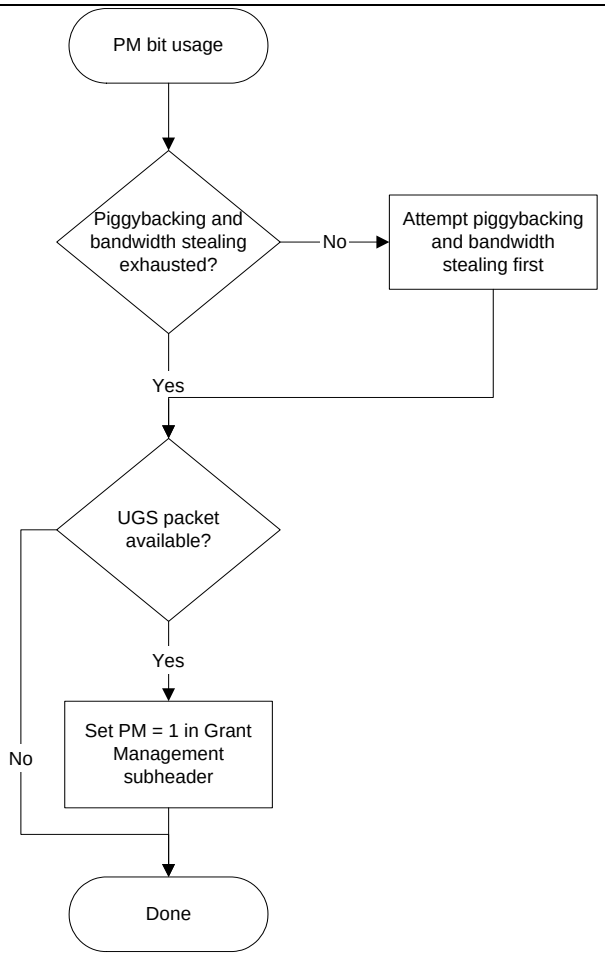

Figure 16 – PM bit usage

13.2 DS-MAP

The broadcast message DS-MAP defines the usage of the downstream intervals.

13.3 US-MAP

The broadcast message US-MAP defines the usage of the upstream intervals in terms of the offset of the upstream burst relative to the Allocation Start Time parameter.

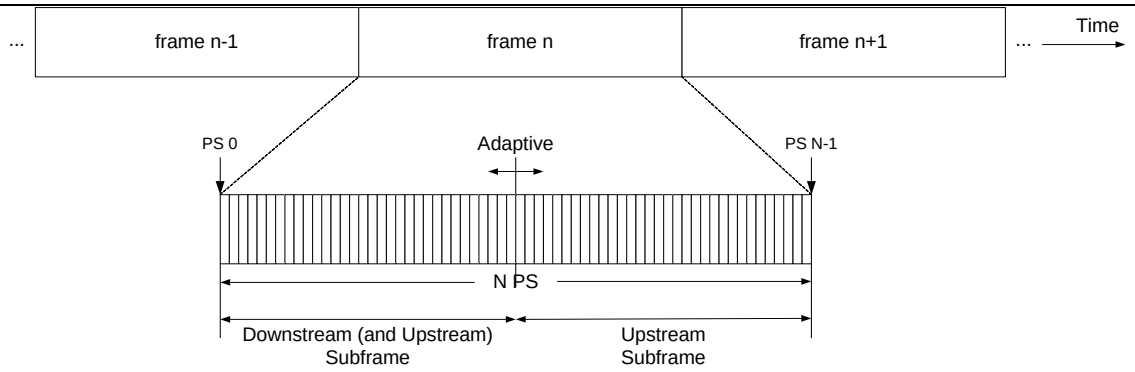

Figure 17 – A TDD frame

13.3.1 Timing

In CMAC, upstream timing is relative to the beginning of the downstream subframe. The Allocation Start Time in the US-MAP is relative to the start of the downstream subframe, and is such that the US-MAP references some point in the current frame (see 13.4). The CPE shall always adjust its concept of upstream timing based upon the Timing Adjustments sent in the RNG-RSP messages.

13.3.2 Allocations

The upstream bandwidth allocation map (US-MAP) employs units of MAC slots and channels to manage resource allocation amongst CPEs.

13.4 Map Timing

The DS-MAP and US-MAP messages provide relative timing information. The following time instants are used as a reference for timing information:

- DS-MAP: The start of the first symbol (including the preamble if present) of the frame in which the message was transmitted.
- US-MAP: The start of the first symbol (including the preamble if present) of the frame in which the message was transmitted plus the value of the Allocation Start Time.

Information contained in both the DS-MAP and US-MAP messages pertain to the current frame (i.e., the frame in which this message was received)[7] as shown in Figure 18. In addition, information carried in the US-MAP pertains to a time interval starting at the Allocation Start Time measured from the beginning of the current frame and ending after the last specified allocation. The Allocation Start Time value shall be equal to the Adaptive TDD (ATDD) split, where ATDD split is the time instant within the frame that divides the downstream subframe and the upstream subframe.

[7] It is important to note that, from a scheduling perspective, allocations specified in a DS-MAP and US-MAP messages shall remain, as much as possible, valid across multiple frames. This way, self-coexistence issues can be best managed.

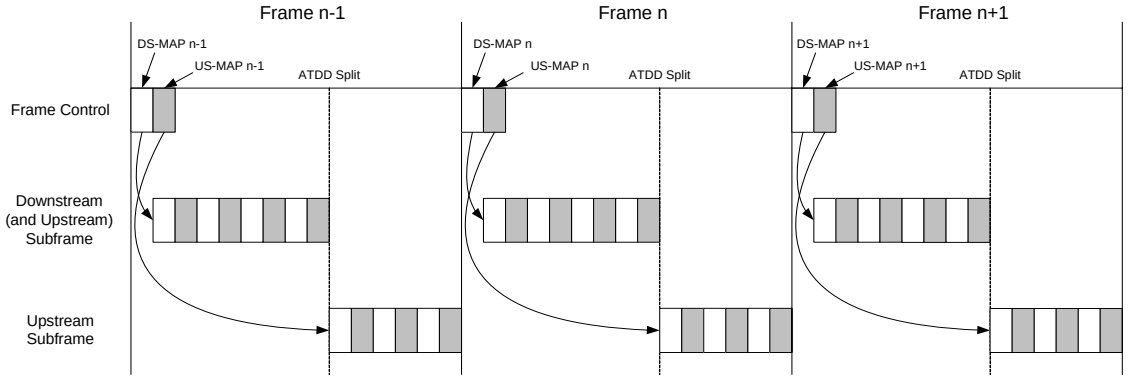

Figure 18 – Time relevance of DS-MAP and US-MAP

13.5 Individual CPE Maximum TPC for the Protection of TV Incumbents

A CPE is not allowed to operate on the first adjacent channel of a TV operation within the Grade B or noise-protected contour. However it can operate co-channel, provided that it is located farther than 10 km from the noise-protected contour of an ATSC TV operation, or 4.7 km away from the Grade B contour of an NTSC TV operation. On other adjacent channels, the CPE must meet a transmit power constraint, defined by the EIRP profile, which constraints the power radiated on channel adjacent to the channel of a TV operation, up to +/- 15. The present contribution presents the method to determine the power constraint for a single CPE transmitting in 6 MHz in the presence of multiple TV operations in adjacent channels, according to the EIRP profile and the estimated or known distance of the CPE to the noise-protected or Grade B contours of nearby TV stations. The method is applied at the base station, from the collective knowledge of channel sensing, CPE locations, TV operation database information. Note that this transmit power constraint is a maximum transmit power constraint. Other constraints can be build up on top of that constraint to decrease the maximum transmit power, but in no case can the maximum transmit power determined by this method be exceeded.

13.5.1 Description of the Interference Management Module

The maximum transmit power is determined sequentially as follows:

- Determine the maximum transmit power for each CPE on each TV band from the constraint of a single TV operation: fill in Table 223 cell by cell, using the flowchart of Figure 19.
- Determine the maximum transmit power for each CPE on each TV band from the constraints of all TV operations: fill in Table 224 using Table 223.

13.5.2 Individual CPE Power Control

Table 223 shows an example of how the individual maximum transmit power constraint for a single CPE is computed from the knowledge of TV operations at the WRAN base station. Each column is filled in turn. Given NTSC operation on the 6 MHZ TV band 2, and given that the CPE is located within the Grade B contour plus a margin area, it is determined that TV bands 2 and 3 are not allowed. The maximum transmit power of that CPE if it was transmitting alone in 6 MHz is determined by the EIRP profile at +2 to +5 on TV bands 3 to 6. It is assumed that the out-of-band emission mask meets the constraints on adjacent TV bands when a CPE is transmitting in a given TV band, so this functional requirement does not have to be addressed in this table (Section 15.1.7 of [3]). The values in bold font in Table 223 show how the individual maximum transmit power constraint for a single CPE is computed from constraints on all TV bands, by taking the minimum of all constraints on each row.

The flowchart of the decisions made to fill in one cell of Table 223 is shown in Figure 19. In [4], calculations show that a CPE transmitting at 4W with TV operation on channel N should be:

- At least 10 km away from noise-protected contour co-channel to DTV operation
- At least 123 m away from noise-protected contour on N-1 of DTV operation
- At least 155 m away from noise-protected contour on N+1 of DTV operation
- At least 4.7 km away from Grade B contour co-channel to NTSC operation
- At least 44 m away from Grade B contour on N-1 of NTSC operation
- At least 31 m away from Grade B contour on N+1 of NTSC operation

In particular, this means that no co-channel and first adjacent channel operation is allowed within the noise-protected/Grade B contours by any CPE. But operation outside the noise-protected/Grade B contours is allowed with a constraint on the minimum distance between the CPE and the contour.

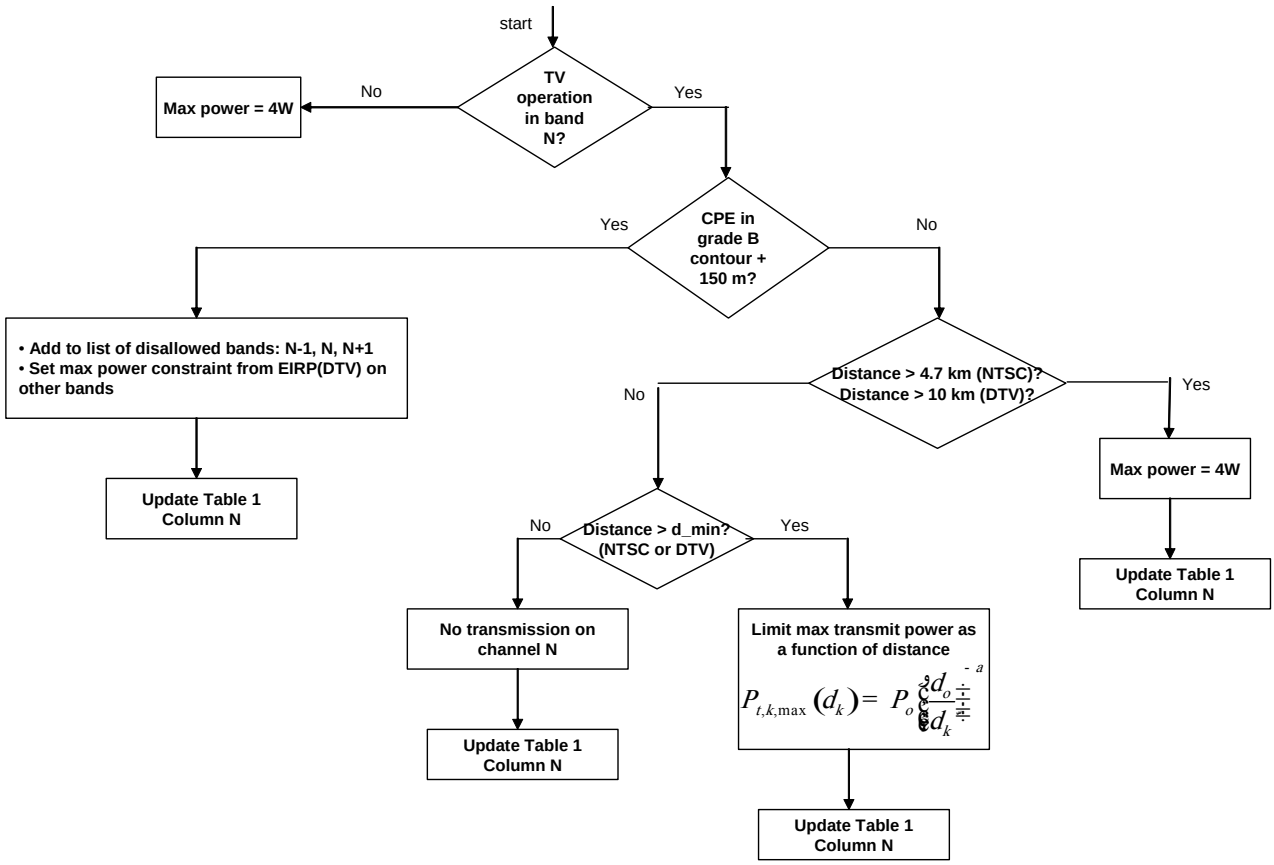

Figure 19 – Flowchart of the decision on the first layer on individual transmit power constraint for each CPE from a possible TV operation on channel N

The parameters in Figure 19 are:

- $P_o = 4\,\mathrm{W}$
- $d_o = 10\,\mathrm{km}$ for DTV, $4.7\,\mathrm{km}$ for NTSC
- d_k distance of CPE to Grade B/noise-protected contour on channel N.
- a is the path loss exponent

Table 223 – Individual CPE maximum transmit power constraint from each individual TV operation, assuming 6 MHz CPE signal bandwidth. The values in bold are then reported in Table 224.

TV band index	1	2	3	4	5	6	7	8
TV band operation	No	Yes	No	Yes	No	No	Yes	No
CPE #1 location	Outside grade B	Inside grade B	Outside grade B	Outside grade B	Inside grade B	Outside grade B	Inside grade B	Outside grade B
CPE #1 operation on TV band 1	4 W	**Not allowed (adjacent)**	4 W	4 W	4 W	4 W	EIRP(-6)	4 W
CPE #1 operation on TV band 2	4 W	**Not allowed**	4 W	4 W	4 W	4 W	EIRP(-5)	4 W
CPE #1 operation on TV band 3	4 W	**Not allowed (adjacent)**	4 W	4 W @ 123 m DTV @ 44 m NTSC	4 W	4 W	EIRP(-4)	4 W
CPE #1 operation on TV band 4	4 W	EIRP(+2)	4 W	4 W @ 10 km DTV @ 4.7 km NTSC	4 W	4 W	**EIRP(-3)**	4 W
CPE #1 operation on TV band 5	4 W	**EIRP(+3)**	4 W	4 W @ 155 m DTV @ 31 m NTSC	4 W	4 W	EIRP(-2)	4 W
CPE #1 operation on TV band 6	4 W	EIRP(+4)	4 W	4 W	4 W	4 W	**Not allowed (adjacent)**	4 W
CPE #1 operation on TV band 7	4 W	EIRP(+5)	4 W	4 W	4 W	4 W	**Not allowed**	4 W

Table 224 summarizes the outcome of this calculation. The values in Table 224 must be scaled to the actual bandwidth occupied by the signal of the individual CPE, if it is transmitting alone within a 6 MHz band. If the CPE is only using 1.5 MHz, then its maximum transmit power can be increased by $6/1.5 = 4$ times (up to 4W EIRP). Otherwise, if the whole band is occupied by multiple CPEs over separate sub-bands, the power constraint applies to the sub-band occupied by each CPE, so that the total power transmitted by all CPEs is 4W in 6 MHz.

Table 224 – Individual CPE maximum transmit power constraint from all TV operations

CPE #1 operation on	Maximum transmit power
TV band 1	Not allowed
TV band 2	Not allowed
TV band 3	Not allowed
TV band 4	**EIRP(-3)**
TV band 5	**EIRP(+3)**
TV band 6	Not allowed
TV band 7	Not allowed

13.6 Optional MAC AAS Support

AAS, through the use of more than one antenna element, can improve range and system capacity by adapting the antenna pattern and concentrating its radiation to each individual CPE. The spectral efficiency can be increased linearly with the number of antenna elements. This is achieved by steering beams to multiple users simultaneously so as to realize an inter-cell frequency reuse of one and an in-cell reuse factor proportional to the number of antenna elements. An additional benefit is the SNR gain realized by coherently combining multiple signals, and the ability to direct this gain to particular users. Another possible benefit is the reduction in interference achieved by steering nulls in the direction of co-channel interferers. Combining the benefits of increasing the SNR of certain CPEs and steering nulls to others enables bursts to be concurrently transmitted to spatially separated CPEs. For the upstream direction the same principle can be applied in a reciprocal fashion. A concurrent transmission of bursts does not necessarily increase the system's range but may enhance system capacity. Support mechanisms for AAS are specified, which allow a system to deliver the benefits of adaptive arrays while maintaining compatibility for non-AAS CPEs.

The design of the AAS option provides a mechanism to migrate from a non-AAS system to an AAS enabled system in which the initial replacement of the non-AAS capable BS by an AAS capable BS should cause the only service interruption to (non-AAS) CPEs. This is achieved by dedicating part of the frame to non-AAS traffic and part to AAS traffic. The allocation is performed dynamically by the BS. Non-AAS CPEs shall ignore AAS traffic, which they can identify based on the DS-MAP/US-MAP messages. The AAS part of the DS frame begins with an AAS specific Preamble. Figure 20 Illustrates the frame structure with AAS zones.

For bandwidth request/allocation, AAS enabled CPEs may use dedicated private DS-MAP/US-MAP messages as well as tools specific for AAS (see specific PHY sections), which can be used to facilitate avoidance of collisions with non-AAS traffic.

Special considerations apply to those parts of the frame that are not scheduled (e.g., initial-ranging and bandwidth request), as discussed in this section.

13.6.1 MAC Control Functions

The control of the AAS part of the frame may be done by unicasting private management messages to individual CPEs. These messages shall be the same as the broadcast management messages, except that the basic CID assigned to the CPE is used instead of the Broadcast CID.

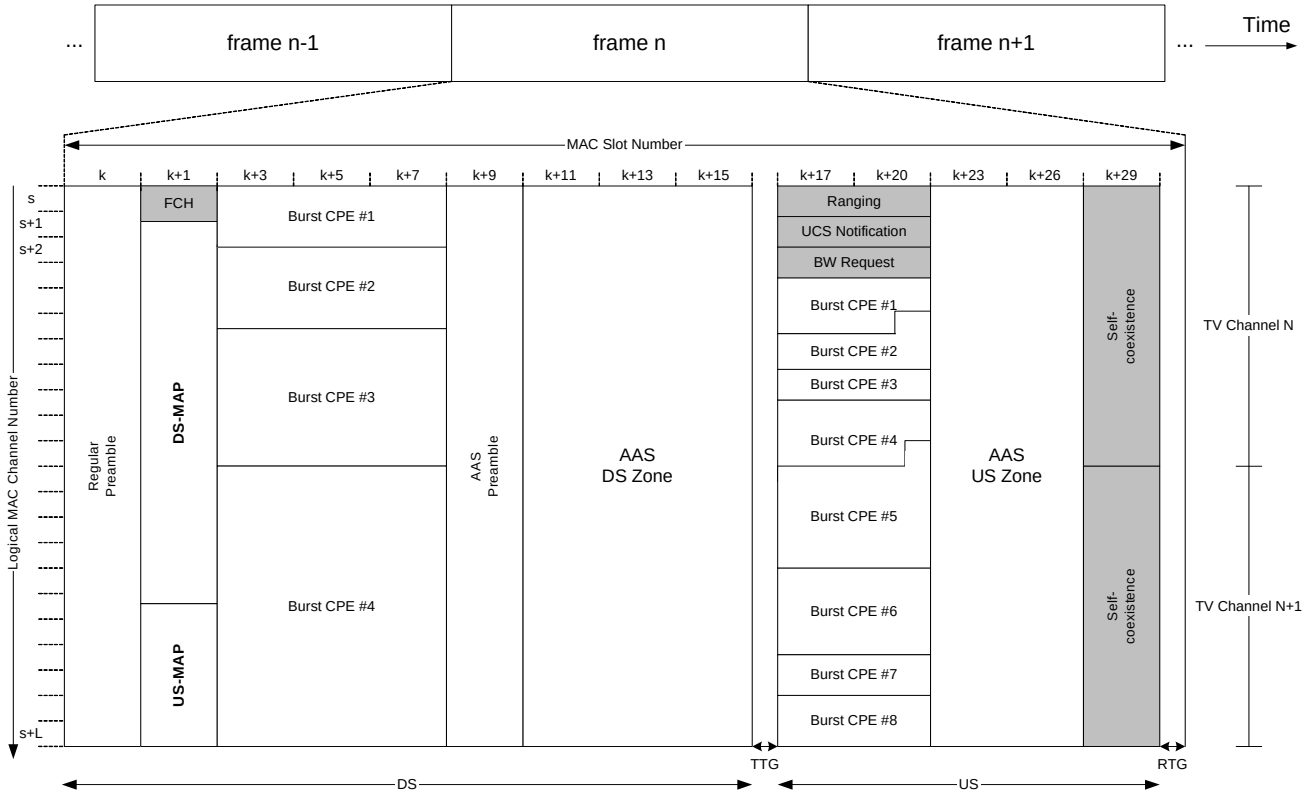

Figure 20 – Frame structure with AAS zones

13.6.2 AAS Downstream Synchronization

When the CPE first attempts to synchronize to the downstream transmission, the BS is unaware of its presence, and therefore is not aiming the adaptive array at its direction. Nevertheless, the frame start preamble is a repetitive well-known pattern, and CPEs may utilize the inherent processing gain associated with it in order to synchronize timing and frequency parameters with the BS.

13.6.3 Alerting the BS about the Presence of a new CPE

In a non-AAS system, after synchronizing to the downstream, a CPE attempts to obtain the downstream parameters by decoding the DS-MAP and DCD messages. In an AAS system, a CPE may be able to obtain the downstream parameters if it receives the broadcast channel with enough energy so it can decode the DS-MAP and DCD messages. If this is the case, the CPE can continue with the network entry process just like the non-AAS case, and the BS will get the chance to tune the adaptive array to it during the ranging process.

Alternatively, an AAS CPE may use the following procedure to alert the BS to its presence, so the BS can adapt its antenna array to the CPE position. An AAS BS may reserve a fixed, pre-defined part of the frame as initial-ranging contention slots for this alert procedure. The number of contention slots and their location in the frame is PHY specific (see PHY spec). These contention slots shall be called Alert-Window, which is depicted in Figure 21.

When an AAS CPE has synchronized to the downstream, yet is unable to obtain the downstream parameters because it cannot decode the DS-MAP and DCD messages, it shall attempt initial ranging on the Alert-Window. Unlike usual initial ranging, the CPE shall use all available contention slots, in order to allow the BS adaptive array enough time and processing gain to shape the beam for it. After such an attempt, the CPE shall wait for a transmission containing DS-MAP and DCD messages from the BS, and shall continue the network entry process like a non-AAS CPE.

If the DS-MAP and DCD messages fail to arrive, the CPE shall use an exponential backoff algorithm for selecting the next frame in which to attempt alerting the BS to its presence. This algorithm shall be the same as that used for initial ranging by non-AAS stations (see 14).

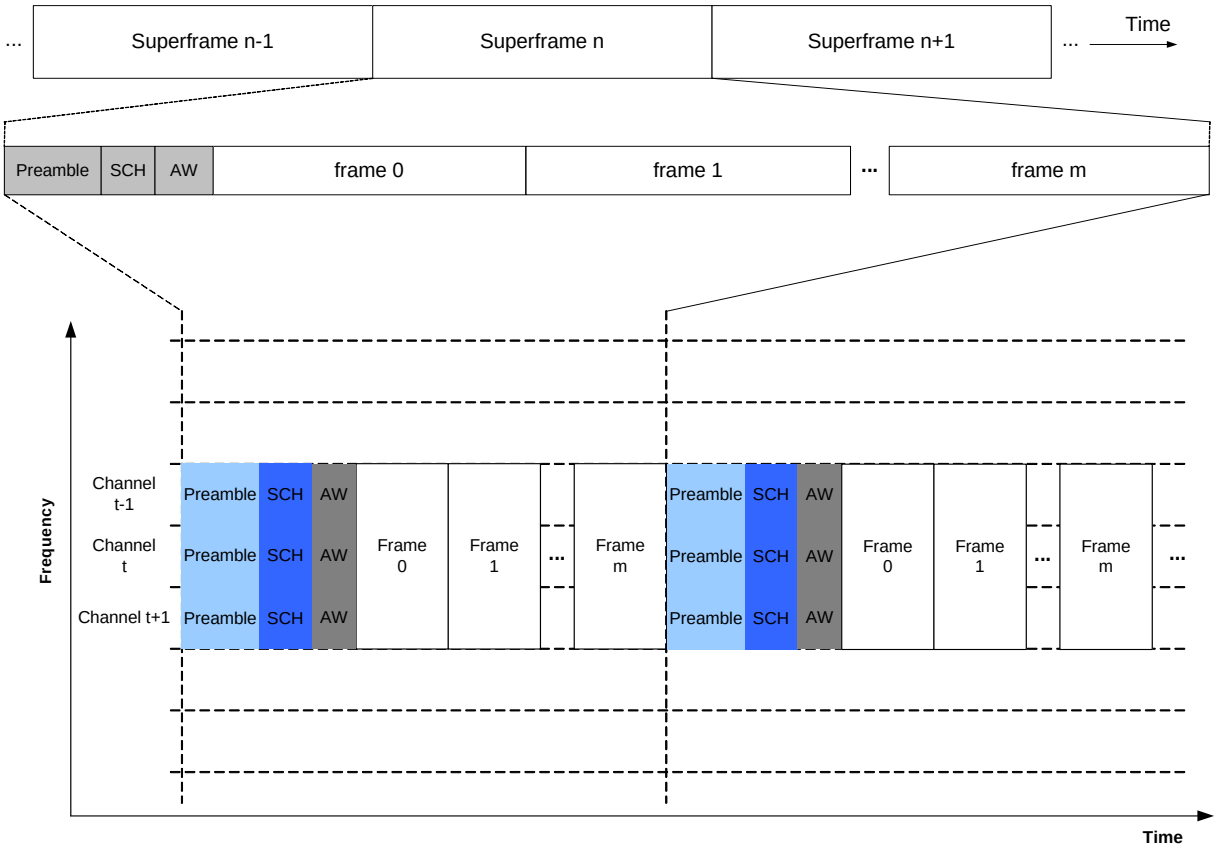

Figure 21 – Alerting the BS about the presence of an AAS-enabled CPE

13.6.4 TDD Support

Adaptive Arrays use channel state information in the PHY at both downstream and upstream. When channel state of the downstream is required at the BS, there are two ways to obtain it:

- By relying on reciprocity, thus using the upstream channel state estimation as the downstream channel state.
- By using feedback, thus transmitting the estimated channel state from the CPE to BS.

The first method is simpler and is well suited for TDD systems. However, given that this first method may be inaccurate, the second method is also supported in CMAC.

Two mechanisms are available in CMAC to obtain channel state information. The first relies on two MAC control messages: AAS-CFB-REQ and AAS-CFB-RSP (see 8.29). The second method is based on piggybacking the channel state measurement request and report together with the incumbent measurement procedures, and therefore improving efficiency (see 8.22). Channel state information requests instruct the CPE to measure, the results of which shall be returned in the response after the measurement period has ended. The BS shall provide an upstream allocation to enable the CPE to transmit this response.

13.6.5 Requesting Bandwidth

AAS CPEs might not be able to request bandwidth using the usual contention mechanism. This happens because the adaptive array may not have a beam directed at the CPE when it is requesting bandwidth, and the Bandwidth Request will be lost. In order to avoid this situation, an AAS CPE is directed by the BS as to whether or not it may use broadcast allocations for requesting bandwidth. The BS may change its direction dynamically using the AAS broadcast permission field, which is carried by the RNG-RSP message. The CPE shall signify by using the AAS broadcast capability field in the RNG-REQ message whether or not it can receive the broadcast messages.

When a CPE is directed not to use the broadcast CID to request bandwidth, it is the responsibility of the BS to provide a polling mechanism to learn about the CPE bandwidth requirements.

14. Contention Resolution

The BS controls assignments on the upstream channel through the US-MAP messages and determines which minislots are subject to collisions. Collisions may occur during Initial Ranging, Request, UCS Notification, and Coexistence intervals defined by their respective IEs. The potential occurrence of collisions in Request Intervals is dependent on the CID in the respective IE. Here we discuss the decision a CPE has to make in order to resolve collision in the upstream direction for both Request and Initial Ranging. Since in the case of UCS Notification and Self-Coexistence (CBP packet transmission) no explicit feedback is expected to be received from the BS, collision resolution does not apply.

Since a CPE can have multiple upstream service flows (each with its own CID), it makes these decisions on a per CID or per service QoS basis. The mandatory method of contention resolution that shall be supported is based on a truncated binary exponential backoff, with the initial backoff window and the maximum backoff window controlled by the BS. The values are specified as part of the UCD message and represent a power-of-two value. For example, a value of 4 indicates a window between 0 and 15; a value of 10 indicates a window between 0 and 1023. When a CPE has information to send and wants to enter the contention resolution process, it sets its internal backoff window equal to the Request (or Ranging for initial ranging) Backoff Start defined in the UCD message referenced by the UCD Count in the US-MAP message currently in effect (the map currently in effect is the map whose allocation start time has occurred but which includes IEs that have not occurred).

The CPE shall randomly select a number within its backoff window. This random value indicates the number of contention transmission opportunities that the CPE shall defer before transmitting. A CPE shall consider only contention transmission opportunities for which this transmission would have been eligible. These are defined by Request IEs (or Initial Ranging IEs for initial ranging) in the US-MAP messages. Note that each IE may consist of multiple contention transmission opportunities.

Using bandwidth requests as an example, consider a CPE whose initial backoff window is 0 to 15 and assume it randomly selects the number 11. The CPE must defer a total of 11 contention transmission opportunities. If the first available Request IE is for 6 requests, the CPE does not use this and has 5 more opportunities to defer. If the next Request IE is for 2 requests, the CPE has 3 more to defer. If the third Request IE is for 8 requests, the CPE transmits on the fourth opportunity, after deferring for 3 more opportunities.

After a contention transmission, the CPE waits for a Data Grant Burst Type IE in a subsequent map (or waits for a RNG-RSP message for initial ranging). Once received, the contention resolution is complete. The CPE shall consider the contention transmission lost if no data grant has been given within T16 (or no response within T3 for initial ranging). The CPE shall now increase its backoff window by a factor of two, as long as it is less than the maximum backoff window. The CPE shall randomly select a number within its new backoff window and repeat the deferring process described above.

This retry process continues until the maximum number (i.e., Request Retries for bandwidth requests and Contention Ranging Retries for initial ranging) of retries has been reached. At this time, for bandwidth requests, the PDU shall be discarded. For initial ranging, proper actions are specified in 15.5. Note that the maximum number of retries is independent of the initial and maximum backoff windows that are defined by the BS. For bandwidth requests, if the CPE receives a unicast Request IE or Data Grant Burst Type IE at any time while deferring for this CID, it shall stop the contention resolution process and use the explicit transmission opportunity.

The BS has much flexibility in controlling the contention resolution. At one extreme, the BS may choose to set up the Request (or Ranging or Self-Coexistence or UCS Notification) Backoff Start and Request (or Ranging or Self-Coexistence or UCS Notification) Backoff End to emulate an Ethernet-style backoff with its associated simplicity and distributed nature as well as its fairness and efficiency issues. This would be done by setting Request (or Ranging or Self-Coexistence or UCS Notification) Backoff Start = 0 and Request (or Ranging or Coexistence or UCS Notification) Backoff End = 10 in the UCD message. At the other end, the BS may make the Request (or Ranging or Self-Coexistence or UCS Notification) Backoff Start and Request (or Ranging or Self-Coexistence or UCS Notification) Backoff End identical and frequently update these values in the UCD message so that all CPE are using the same, and hopefully optimal, backoff window.

14.1 Transmission Opportunities

A transmission opportunity is defined as an allocation provided in a US-MAP or part thereof intended for a group of CPEs authorized to transmit bandwidth requests or Initial Ranging requests. This group may include either all CPEs having an intention to join the cell or all registered CPEs or a multicast polling group. The number of transmission opportunities associated with a particular IE in a map is dependent on the total size of the allocation as well as the size of an individual transmission.

The size of an individual transmission opportunity for each type of contention IE shall be published in each transmitted UCD message. The BS shall always allocate bandwidth for contention IEs in integer multiples of these published values.

15. Network Entry and Initialization

Before a CPE can be serviced by a BS, it needs to enter the network and negotiate its capabilities with the BS. This may involve many tasks (e.g., sensing channels) and frame exchanges between the CPE and the BS, and this whole procedure is hereby referred to as network entry and initialization. More importantly, during this process the CPE needs to ensure that before it first transmits to the BS, its communication will not cause harmful interference with incumbents. In other words, the network entry and initialization process has to be designed to be what is hereby referred to as *incumbent safe*, which essentially means that incumbent system protection shall be guaranteed.

Figure 22 illustrates a scenario where the need for the definition of an incumbent safe bootstrap procedure can be easily seen. In this figure, consider that CPE 4 is powered down whereas the BS is transmitting in the cell which is under normal operation. Further, assume that the TV station in Figure 22 is powered up and starts transmitting in the same channel (i.e., channel #52 in this example) that is being used by the BS for its transmissions in the

cell. Assuming that the BS cannot detect the signal from the TV station in such low SNR values, it will continue to use the channel and hence may cause harmful interference to incumbent users. Therefore, whenever CPE 4 is powered up, it should be capable of detecting that the BS is operating in a channel that it is occupied by an incumbent service, and so the CPE shall not associate with this BS. If permitted, however, the CPE could send a very short notification to the BS indicating that the BS is using a channel occupied by an incumbent. As we can see, the definition of an incumbent safe bootstrap phase is critical for cognitive radio systems. CMAC incorporates algorithms to address this need.

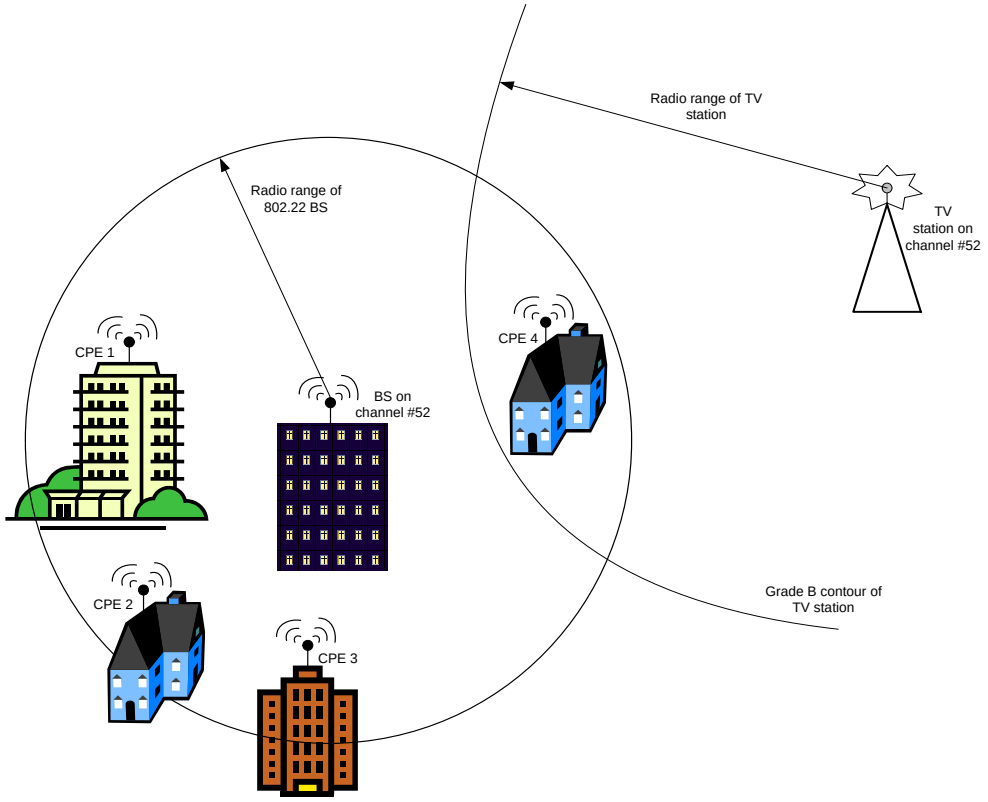

Figure 22 – Scenario where a safe bootstrap operation is required to protect incumbents

First and foremost, CMAC does not presuppose any pre-assigned channel where a CPE is able to look for a BS given the time-varying and unpredictable nature of channel occupancy. Hence, the first task a CPE must perform once it attempts to join a network is to scan the set of channels it is programmed to and capable of. Even though a BS within the coverage of the CPE may be grouping multiple channels together, it shall periodically send a SCH in each channel (see 5.1) which allows the CPE to recognize and, if appropriate, proceed with the network entry and initialization procedure with the corresponding BS.

The procedure carried out by the CPE to perform network entry and initialization is as follows:

1. Scan channels searching for a BS[8].
2. Once SCH is received, ascertain that the use of the channel(s) is permitted (i.e., does not interfere with incumbents)
3. Synchronize to the BS.
4. Obtain the transmit parameters from the BS, which are contained in the UCD message.
5. Perform ranging and Negotiate basic capabilities.
6. Authorize CPE and Perform key exchange.

[8] To reduce the time taken to lock to a BS, algorithms such as Bandwidth-Greedy and Reference Channel could be implemented.

7. Perform registration.
8. If indicated as desired by the CPE during registration (REG-REQ message), perform other optional initialization procedures such as establish IP connectivity, establish time of day, and transfer operational parameters.
9. Set up connections.

Figure 23 summarizes the network entry and initialization procedure carried out by CPEs. Note that each these steps taken by the CPE consist of a set of actions and error verification. In the following subsections, we provide a more detailed view of these steps and their individual responsibilities.

15.1 BS Initialization

WRAN systems may have to reinitialize for example if there is a situation that no TV channel is found empty during WRAN operation due to which it has to shut down. The process of initialization is not straightforward in the case of WRANs due to their unlicensed nature. The solution involves the following steps. The WRAN BS starts by consulting the TV usage database and the regional WRAN information base to find potentially empty channels. To ensure these channels are indeed empty, it performs sensing to find one or more empty channels. The WRAN BS begins its service on channels found vacant.

15.2 Scanning Downstream Channels

On initialization or after signal loss, the CPE shall acquire a downstream channel. The CPE shall have non-volatile storage in which the last operational parameters are stored and shall first try to reacquire this downstream channel. If this fails, it shall begin to continuously scan the possible channels of the downstream frequency band of operation until it finds a valid downstream signal.

Once the PHY has achieved synchronization, as given by a PHY Indication, the MAC shall attempt to acquire the channel control parameters for the downstream and then the upstream.

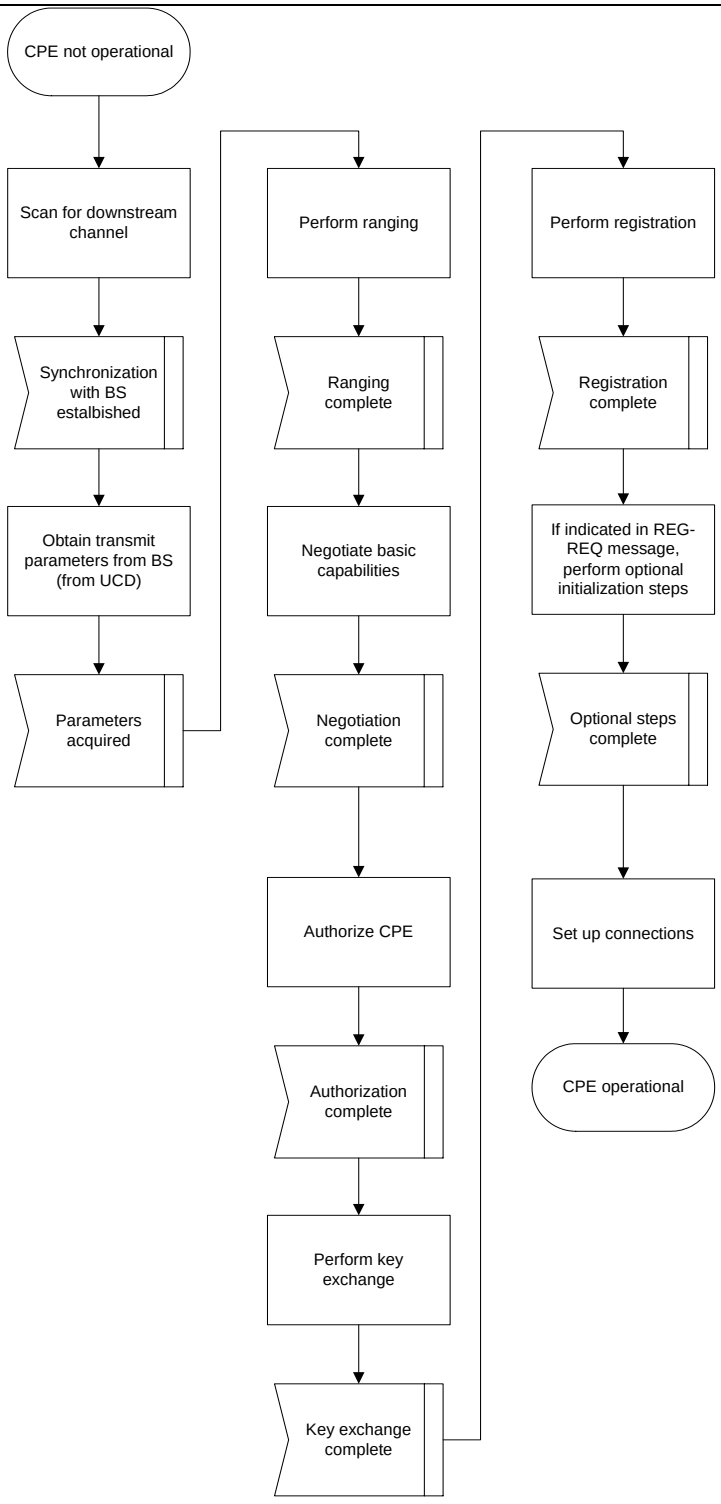

Figure 23 – CPE network entry and initialization procedure

15.3 Obtaining Downstream Parameters

The MAC shall search for the SCH message from the BS, which indicates the beginning of the superframe. To improve the joining latency in case a long superframe is in use by the BS, the CPE shall use energy detection to

help ascertain about the presence/absence of an 802.22 BS in a particular channel. If the energy detected is below the detection threshold, the CPE can safely move to the next channel.

After having received an SCH in a channel, the CPE shall perform sensing not only in the set of channels indicated in the SCH, but also in all other affected channels (e.g., N-t through N+t, where $t \leq 15$). During this sensing, the CPE shall attempt to identify incumbent operation. If incumbents are detected, the MAC should disregard the channel and, if permitted (e.g., by the DFS model parameters), send a short control message to the BS indicating that is using a channel occupied by an incumbent[9]. In case the BS receives such notification, it may take numerous actions as described in 21.1.

Provided no incumbents are found, the CPE may proceed to the next step. Here, the MAC shall search for the DS-MAP MAC management messages. The CPE achieves MAC synchronization once it has received at least one DS-MAP message. An CPE MAC remains in synchronization as long as it continues to successfully receive the SCH, DS-MAP and DCD messages for its channel(s). If the Lost DS-MAP Interval (Table 226) has elapsed without a valid DS-MAP message or the T1 interval (Table 226) has elapsed without a valid DCD message or Lost SCH counts of SCH are missed, an CPE shall try to re-establish synchronization. The process of acquiring synchronization is illustrated in Figure 24. The process of maintaining synchronization is illustrated in Figure 25.

15.4 Obtaining Upstream Parameters

After synchronization, the CPE shall wait for a UCD message from the BS in order to retrieve a set of transmission parameters for a possible upstream channel. These messages are transmitted periodically from the BS for all available upstream channels and are addressed to the MAC broadcast address.

If no upstream channel can be found after a suitable timeout period, then the CPE shall continue scanning to find another downstream channel. The process of obtaining upstream parameters is illustrated in Figure 26.

The CPE shall determine from the channel description parameters whether it may use the upstream channel. If the channel is not suitable, then the CPE shall continue scanning to find another downstream channel. If the channel is suitable, the CPE shall extract the parameters for this upstream from the UCD. It then shall wait for the next DS-MAP message and extract the time synchronization from this message. Then, the CPE shall wait for a bandwidth allocation map for the selected channel. It may begin transmitting upstream in accordance with the MAC operation and the bandwidth allocation mechanism.

The CPE shall perform initial ranging at least once. If initial ranging is not successful, the procedure is restarted from scanning to find another downstream channel.

The CPE MAC is considered to have valid upstream parameters as long as it continues to successfully receive the SCH, US-MAP and UCD messages. If at least one of these messages is not received within the time intervals specified in Table 226, the CPE shall not use the upstream. This is illustrated in Figure 27.

[9] This message shall be the shortest MAC frame (i.e., the size of the general MAC header), and will possibly not overrun the DTV receiver's interleaver.

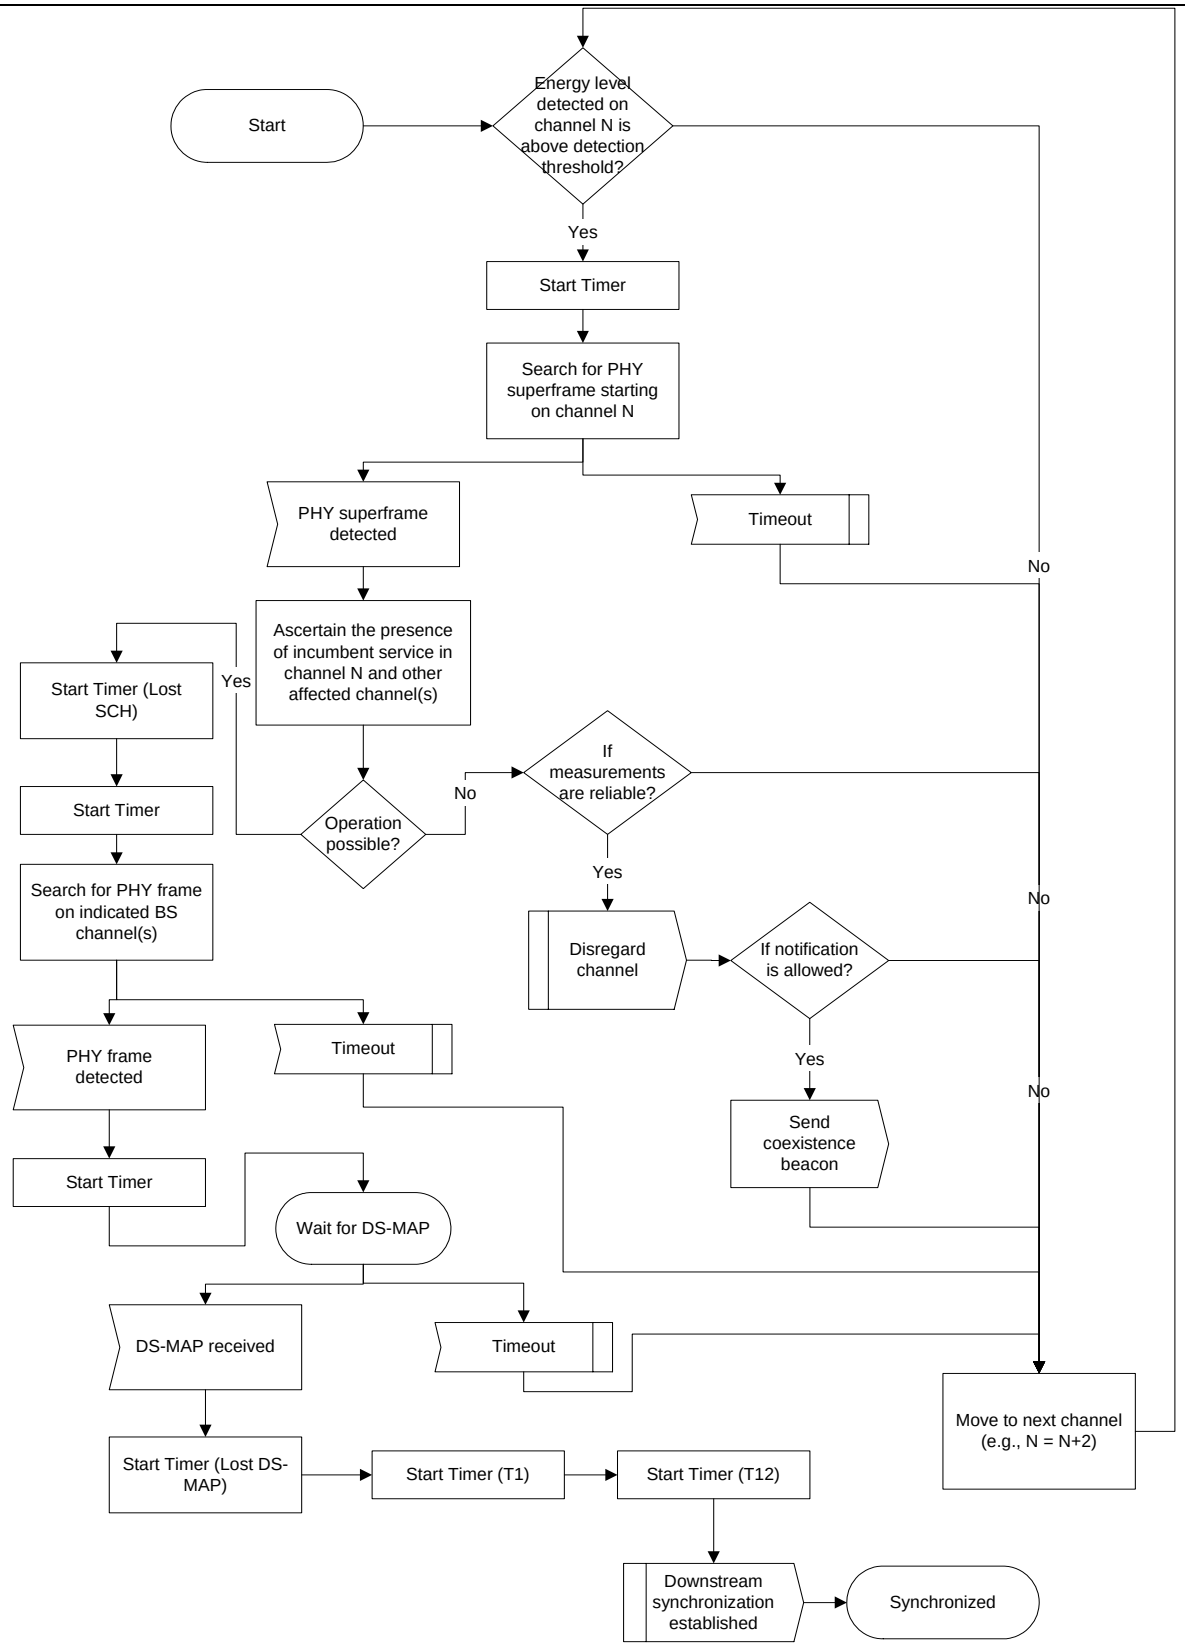

Figure 24 – Obtaining downstream parameters

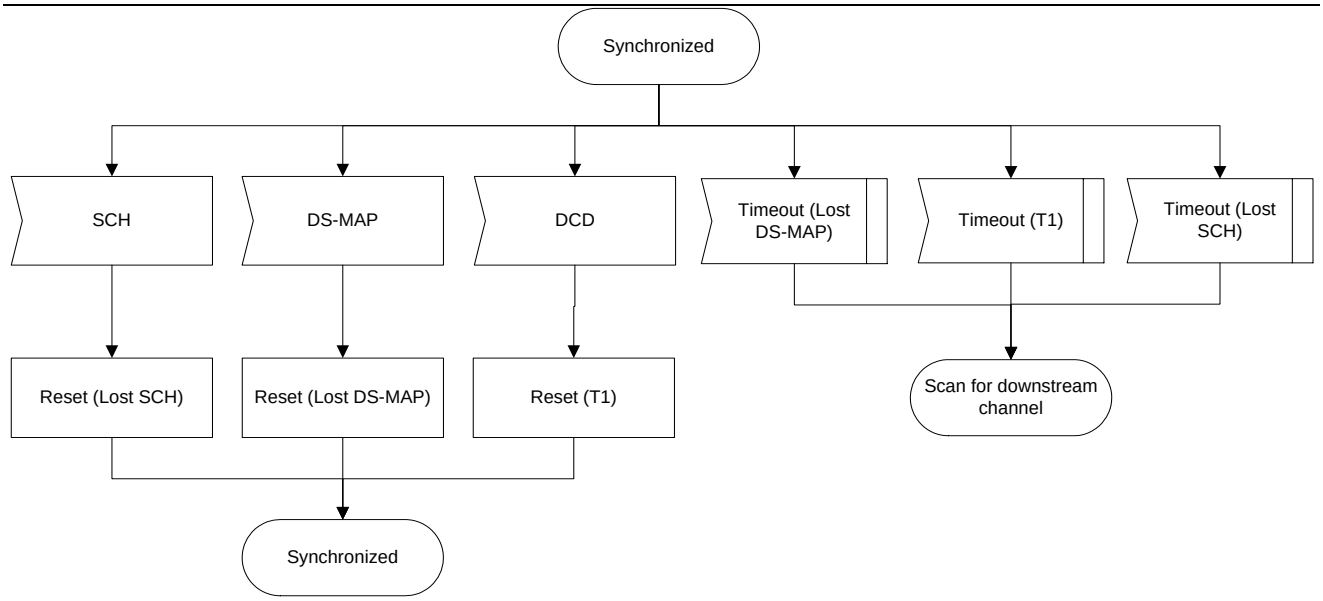

Figure 25 – Maintaining downstream parameters

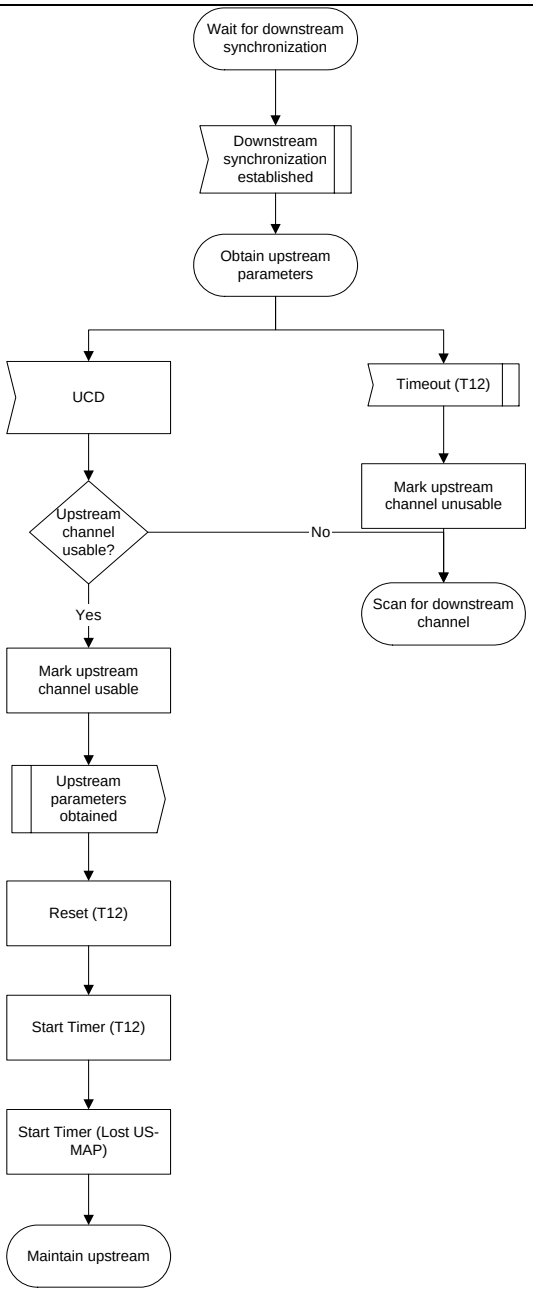

Figure 26 – Obtaining upstream parameters

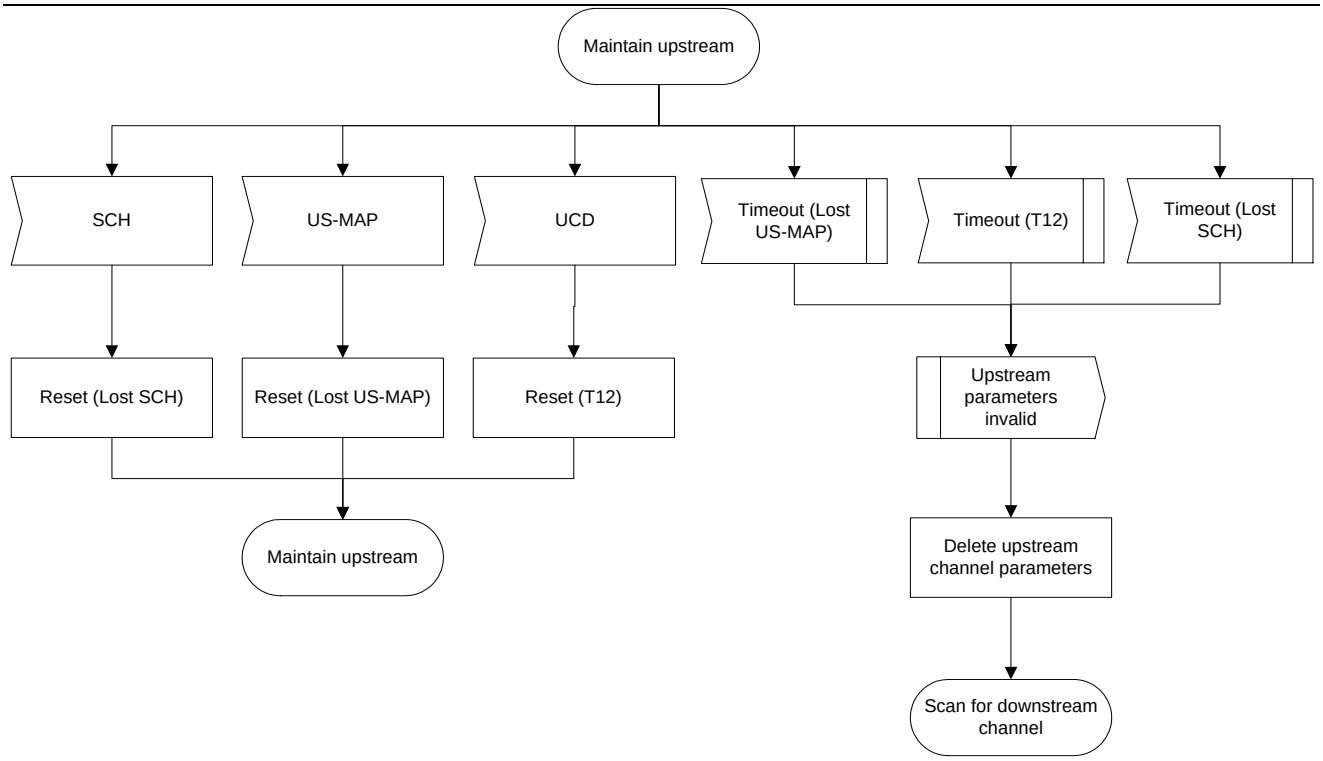

Figure 27 – Maintaining upstream parameters

15.5 Initial Ranging and Automatic Adjustments

Ranging is the process of acquiring the correct timing offset and power adjustments such that the CPE's transmissions are aligned with the BS receive frame, and received within the appropriate reception thresholds. The timing delays through the PHY shall be relatively constant. Any variation in the PHY delays shall be accounted for in the guard time of the upstream PHY overhead.

15.5.1 Contention-based Initial Ranging and Automatic Adjustments

First, a CPE shall synchronize to the downstream and learn the upstream channel characteristics through the UCD MAC management message. At this point, the CPE shall scan the US-MAP message to find an Initial Ranging Interval. The BS shall allocate an Initial Ranging Interval consisting of one or more transmission opportunities. The size of each transmission opportunity shall be as specified by the Ranging request opportunity size (see 8.3.1).

The CPE shall put together a RNG-REQ message to be sent in an Initial Ranging Interval. The CID field shall be set to the non-initialized CPE value (zero). Alternatively, the initial ranging process shall begin by sending initial-ranging CDMA codes on the US allocation dedicated for that purpose, in addition to RNG-REQ messages sent on contention slots.

Ranging adjusts each CPE's timing offset such that it appears to be co-located with the BS. The CPE shall set its initial timing offset to the amount of internal fixed delay equivalent to collocating the CPE next to the BS. This amount includes delays introduced through a particular implementation and shall include the downstream PHY interleaving latency, if any.

When the Initial Ranging transmission opportunity occurs, the CPE shall send the RNG-REQ message or a CDMA code. Thus, the CPE sends the message as if it were collocated with the BS.

The CPE shall calculate the maximum transmit signal strength for initial ranging, $P_{TX\_IR\_MAX}$, from the following equation:

$P_{TX\_IR\_MAX}$ = EIR x $P_{IR,max}$ + BS_EIRP – RSS

where the EIR x $P_{IR,max}$ and BS_EIRP are obtained from the DCD, and RSS is the measured RSSI, by the CPE, as described in the PHY.

In the case that the receive and transmit gain of the CPE antennae are substantially different, the CPE shall use the following equation:

$P_{TX\_IR\_MAX}$ = EIR x $P_{IR,max}$ + BS_EIRP – RSS + (G_{RX_CPE} – G_{TX_CPE})

where G_{RX_CPE} is the CPE receive antenna gain and G_{Tx_CPE} is the CPE transmit antenna gain.

In the case that the EIR x $P_{IR,max}$ and/or BS_EIRP are/is not known, the CPE shall start from the minimum transmit power level defined by the BS.

NOTE – The EIR x $P_{IR,max}$ is the maximum equivalent isotropic received power, which is computed for a simple single antenna receiver as $RSS_{IR,max}$ – GANT_BS_RX, where the $RSS_{IR,max}$ is the received signal strength at antenna output and GANT_BS_RX is the receive antenna gain. The BS_EIRP is the equivalent isotropic radiated power of the base station, which is computed for a simple single-antenna transmitter as P_{TX} + GANT_BS_TX, where P_{TX} is the transmit power and GANT_BS_TX is the transmit antenna gain.

In the case that the CPE uses RNG-REQ messages, the CPE shall send the RNG-REQ at a power level below $P_{TX\_IR\_MAX}$, measured at the antenna connector. If the CPE does not receive a response, the CPE shall resend the RNG-REQ at the next appropriate Initial Ranging transmission opportunity at one step higher power level. If the CPE receives a response containing the frame number in which the RNG-REQ was transmitted, it shall consider the transmission attempt unsuccessful but implement the corrections specified in the RNG-RSP and issue another RNG-REQ message after the appropriate backoff delay. If the CPE receives a response containing its MAC Address, it shall consider the RNG_RSP reception successful.

When a BS detects a transmission in the ranging slot that it is unable to decode, it may respond by transmitting a RNG-RSP that includes transmission parameters, but identifies the frame number and frame opportunity when the transmission was received instead of the MAC Address of the transmitting CPE.

In the case that the CPE uses CDMA, the CPE shall send a CDMA code at a power level below $P_{TX\_IR\_MAX}$, measured at the antenna connector. If the CPE does not receive a response, the CPE shall send a new CDMA code at the next appropriate Initial Ranging transmission opportunity at one step higher power level. If the CPE receives a RNG-RSP message containing the parameters of the code it has transmitted and status continue, it shall consider the transmission attempt unsuccessful but implement the corrections specified in the RNG-RSP and issue another CDMA code after the appropriate backoff delay. If the CPE receives an US-MAP containing a CDMA allocation IE with the parameters of the code it has transmitted, it shall consider the RNG-RSP reception successful, and proceed to send a unicast RNG-REQ on the allocated BW.

Once the BS has successfully received the RNG-REQ message, it shall return a RNG-RSP message using the initial ranging CID. Within the RNG-RSP message shall be the Basic and Primary Management CIDs assigned to this CPE. The message shall also contain information on RF power level adjustment and offset frequency adjustment as well as any timing offset corrections. At this point the BS shall start using invited Initial Ranging Intervals addressed to the CPE's Basic CID to complete the ranging process, unless the status of the RNG-RSP message is success, in which case the initial ranging procedure shall end.

If the status of the RNG-RSP message is continue, the CPE shall wait for an individual Initial Ranging interval assigned to its Basic CID. Using this interval, the CPE shall transmit another RNG-REQ message using the Basic CID along with any power level and timing offset corrections.

The BS shall return another RNG-RSP message to the CPE with any additional fine tuning required. The ranging request/response steps shall be repeated until the response contains a Ranging Successful notification or the BS aborts ranging. Once successfully ranged (RNG-REQ is within tolerance of the BS), the CPE shall join normal data traffic in the upstream. In particular, state machines and the applicability of retry counts and timer values for the ranging process are defined in Table 226.

NOTE – The burst profile to use for any upstream transmission is defined by the Upstream Interval Usage Code (UIUC). Each UIUC is mapped to a burst profile in the UCD message.

NOTES

1—The BS shall allow the CPE sufficient time to have processed the previous RNG-RSP (i.e., to modify the transmitter parameters) before sending the CPE a specific ranging opportunity. This is defined as CPE Ranging Response Processing Time in Table 226.

2—For multichannel support, the CPE shall attempt initial ranging on every suitable upstream channel before moving to the next available downstream channel.

On receiving a RNG-RSP instruction to move to a new downstream frequency and/or upstream channel ID, the CPE shall consider any previously assigned Basic, Primary Management, and Secondary Management CIDs to be deassigned, and shall obtain new Basic, Primary Management, and Secondary Management CIDs via initial ranging and registration.

It is possible that the RNG-RSP may be lost after transmission by the BS. The CPE shall recover by timing out and reissuing its Initial RNG-REQ. Since the CPE is uniquely identified by the source MAC address in the Ranging Request, the BS may immediately reuse the Basic, Primary Management, and Secondary Management CIDs previously assigned. If the BS assigns new Basic, Primary Management, and Secondary Management CIDs, it shall make some provision for aging out the old CIDs that went unused.

16. Multiple Channel Support

An important feature of CMAC is the capability to take advantage of the simultaneous availability of multiple vacant TV channels, be these contiguous or not. The simultaneous use of non-contiguous TV channels (also called channel aggregation) is made possible through the flexible architecture provided by CMAC (as discussed in 1.2), whereas the use of contiguous TV channels is done through the channel bonding mechanism (see Sections 3 and 5.1). In fact, multiple channel support is a critical feature of an 802.22 standard since, as reported in [6] after extensive measurements of real spectrum usage, there is a significant amount of spectrum available.

In this section, it is described the several features of CMAC when operating under multiple channels.

16.1 Operation under Multiple TV Channels

Whenever the incumbent detection procedure together with the distributed sensing capability conclude that it is safe to do it so, the BS may group multiple contiguous TV channels for the purpose of performance improvement, support of higher number of CPEs within a MAC frame, and achieving longer ranges. When in the multiple channel mode of operation, the BS shall transmit in each TV channel the SCH frame preceded by the superframe

preamble as shown in Figure 3. Within the SCH the BS shall indicate which TV channels are being grouped together, which will allow CPEs to detect the multiple channel mode of operation.

Once the superframe preamble and SCH are repeatedly transmitted in each channel, the BS shall immediately initiate the first frame by transmitting the MAC frame across all TV channels in use as per indicated in the SCH in effect. This is depicted in Figure 28, which present an exemplary time structure of a MAC frame where only key zones are illustrated. Within a multiple channel MAC frame, the BS has the freedom to schedule DS and US traffic that spans any number of TV channels.

16.2 Operation under Change in Number of TV Channels

Whenever the number of channels the BS is using changes, this will cause a change in the number of logical channels available at PHY and hence impact the scheduling function (MAC slot size remains the same). In such an event, it is implementation specific if any action whatsoever will be taken. However, irrespective of the actions to be taken, the MAC shall never change the MAC frame size. This is needed as to seek better self-coexistence with other overlapping 802.22 cells. Also, in the next frame after a change in the number of channels the MAC should transmit the DS-MAP and US-MAP messages to adjust the transmission schedule of CPEs to the new channel configuration.

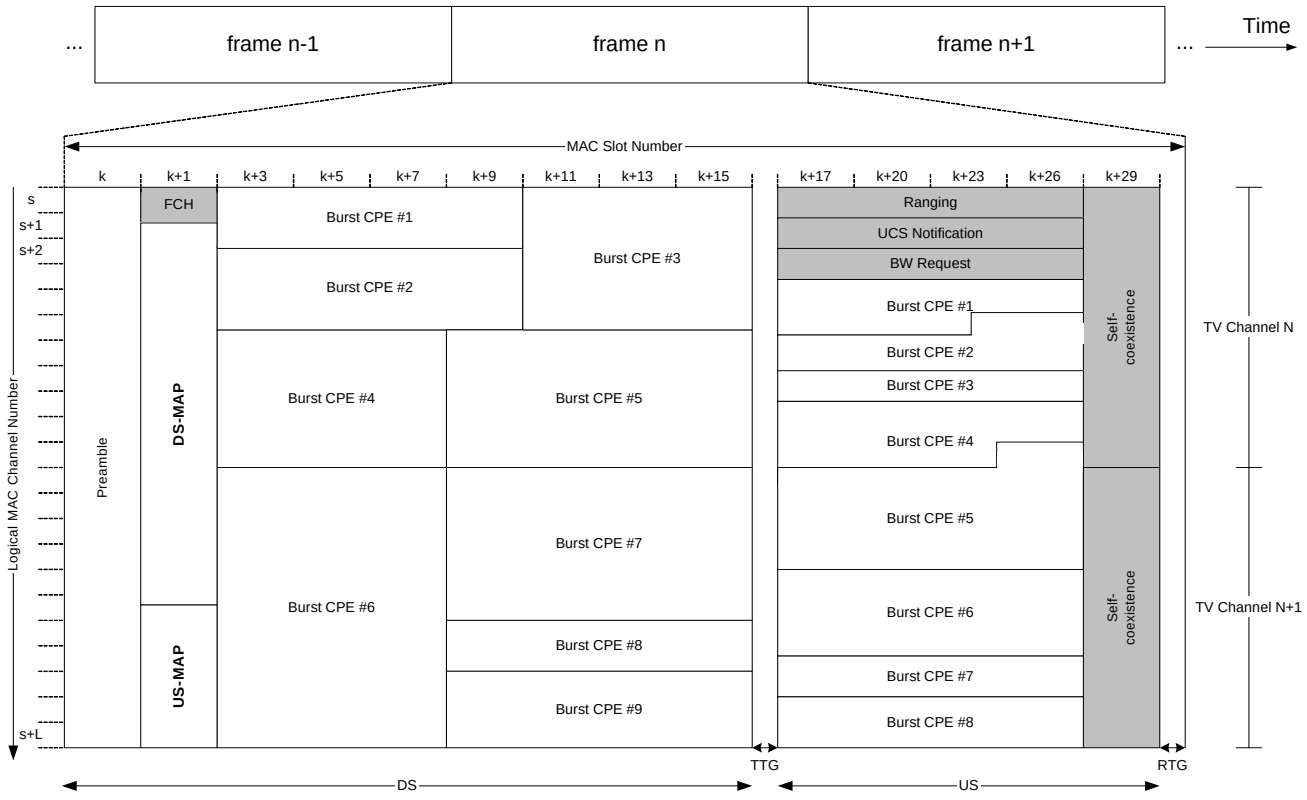

Figure 28 – Time structure of a MAC frame (with only key zones)

16.3 Channel Grouping and Matching

Channel grouping and matching is a procedure applied at the BS level and implemented within the Spectrum Manager. It is applicable only when there is more than one protocol stack available at the BS (see 1.2).

When multiple protocol stacks (and hence, multiple channels) are available, some form of radio resource management is useful. First of all, we note that some excessive MAP overhead is required when dealing with multiple channels in the system. As shown in Figure 29, the MAP must specify a list of the upstream and downstream channel pairs for all CPEs in the system. Such MAP overhead for specifying multi-channel allocation is very large.

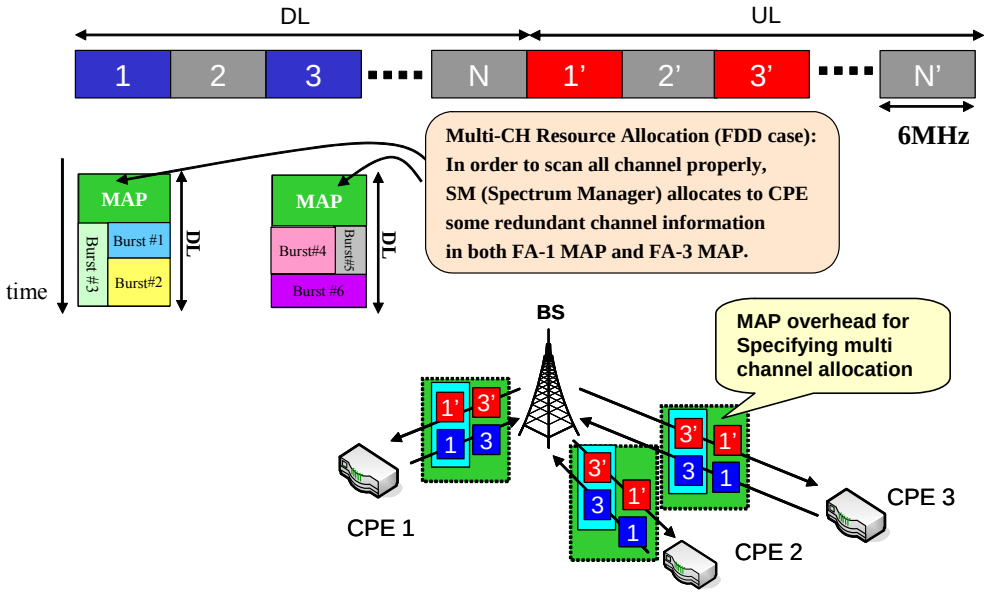

Figure 29 – Multi-channel resource allocation

Channel grouping is employed to select a group of CPEs that are assigned to the same channel. Furthermore, channel matching is to select active set 1 for individual CPE (See 21.6). In Figure 30, we find that CPE 2 and CPE 3 belong to the same channel group. We note that channel grouping is associated to channel matching in the sense that each channel group shares the same channel matching result. Once channel matching and grouping is performed, the different MAPs can be transmitted in the different bands via the different stacks, which reduces the MAP overhead to cover all CPEs in the system.

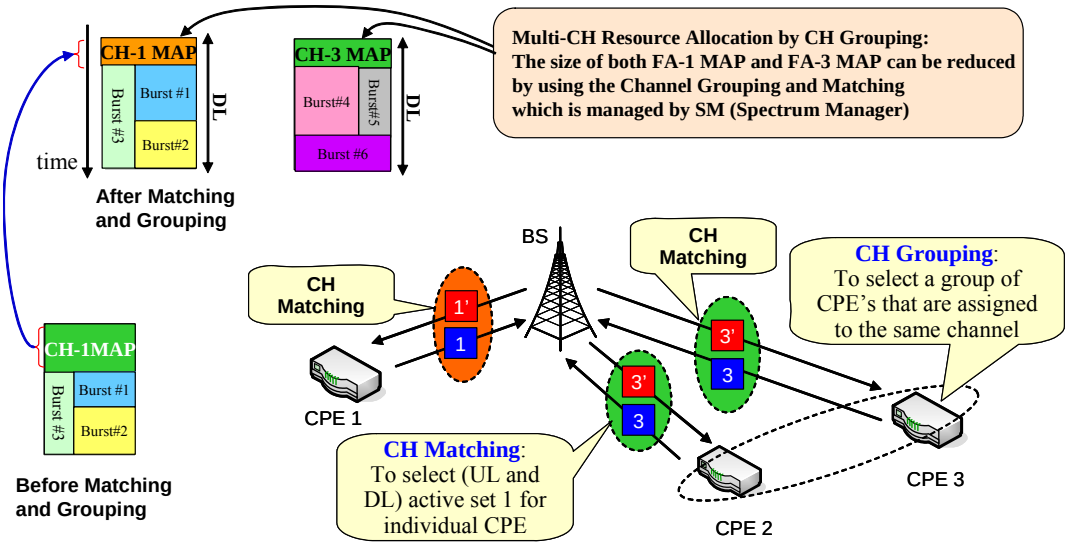

Figure 30 – Channel grouping and matching

Depending on the time-varying channel condition, it can be reported to BS via BLM-REP message so that its own active set 1 can be updated. Following the active set update, channels can be regrouped. The detail procedure is depicted in Figure 31. The results of channel grouping can be announced in the DS-MAP prefix. Figure 32 shows our new signal field for DS-MAP Prefix, which is appended right before the DS-MAP. It is used for announcing the result of channel regrouping as the active set 1 is updated as a part of radio resource management.

Another aim of channel grouping and matching is maximization of the bandwidth utilization. The proposed system provides some means of channel switching when a CPE find a channel with better quality, which can maximize the average system throughput. In other words, it is important to determine the channel that warrants a quality, as the system throughput is directly determined by the channel quality. To this end, we need to address the procedures of selecting the active set 1 for individual CPEs in the downstream and upstream ("Channel matching" procedure), and selecting a group of CPE's assigned to the same channel ("Channel grouping" procedure). Even though there is no incumbent user detection, channel matching and channel grouping procedures shall be performed whenever there is more than one stack, so as to maximize the system utilization (average throughput) while minimizing the system cost with the constraints on guard band for FDD operation, co-channel interference, cross-talk in transceiver, and so on.

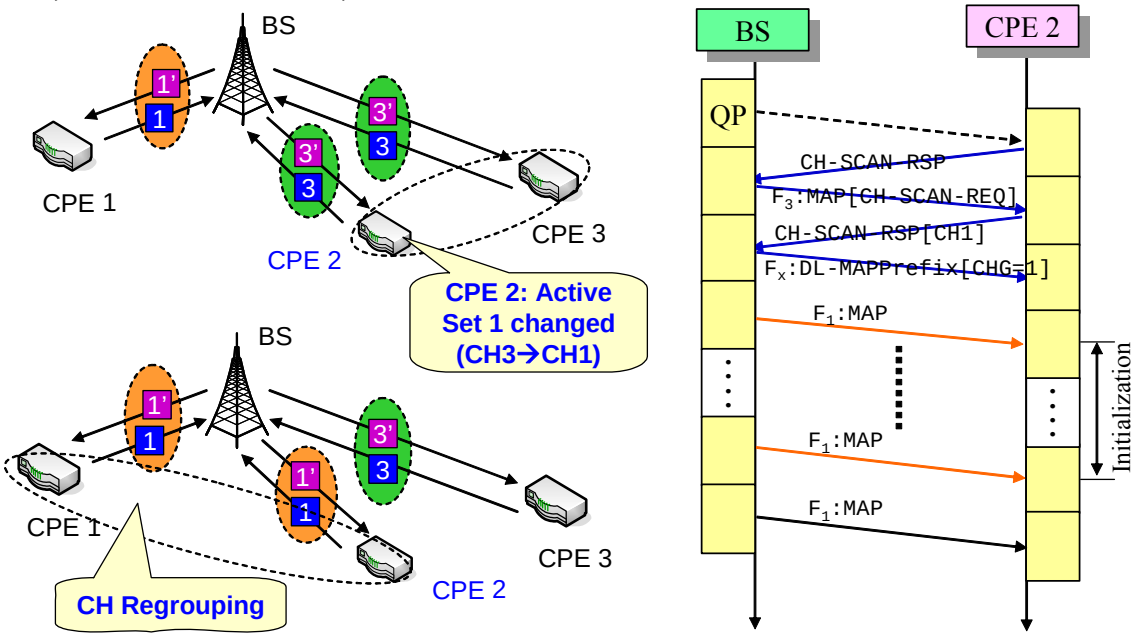

Figure 31 – Active set update and channel regrouping

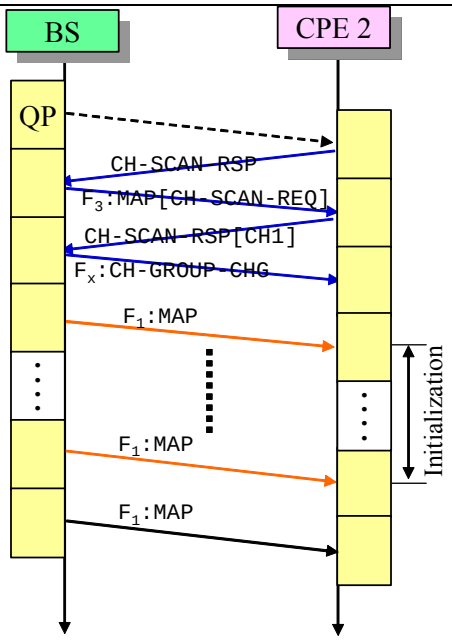

- DL-MAP Information Element

Field	Note
Extended DIUC	
Number of CPEs_to_update	
for (i=1; Number of CPEs_to_update) {	
CPE CID	Primary CID
Group ID	New group ID
DL channel ID	New group DL
UL channel ID	New group UL
}	

Figure 32 – CH-GROUP-CHG information element

16.4 Explicit Outband Signalling for Hidden Incumbent System Detection

16.4.1 Hidden Incumbent Systems

Assume that a WRAN system is on service in channel x. The BS sensed some channels and it recognized channel x was available, or BS just started the service based on its database information. Some CPEs inside the incumbent system radio area may not be able to decode the BS signal because of strong interference. So, the CPEs cannot report the existence of the incumbent system and current status to the BS. BS cannot recognize this situation because of no information. Also, some incumbent users have experienced interference from the WRAN system. This situation is depicted in Figure 33.

The hidden Incumbent case is occurred when a BS starts the service, the BS changes the service channel. Candidate channel broadcasting by BS and sensing reports for the candidate bands by CPEs can reduce the possibility of hidden incumbent system. But, BS may change its service channel without notification.

16.4.2 Explicit Outband Signalling for Hidden Incumbent Case Detection

To address the hidden incumbent system case, BS periodically broadcasts Out-band signal including the information on current channel in some of other unoccupied channels (e.g., candidate channels). The Out-band signal is control signal on the band other than current band. This broadcasting signal follows the same PHY and MAC frame architecture (not to necessitate additional protocol or PHY module). When some CPEs cannot decode the BS's current service channel, the CPEs try to sense other channels to locate the BS signal. If CPEs receive the explicit out-band broadcast signal, the CPEs recognize the current service channel id. If the current channel is already sensed and is found to be not decidable at the CPEs, then the CPE sends a report to the BS using the upstream in out-band. After receiving the report, BS changes its service channel to other available band because BS notices the existence of the hidden incumbent. Figure 34 shows the explicit outband signalling for hidden incumbent case detection.

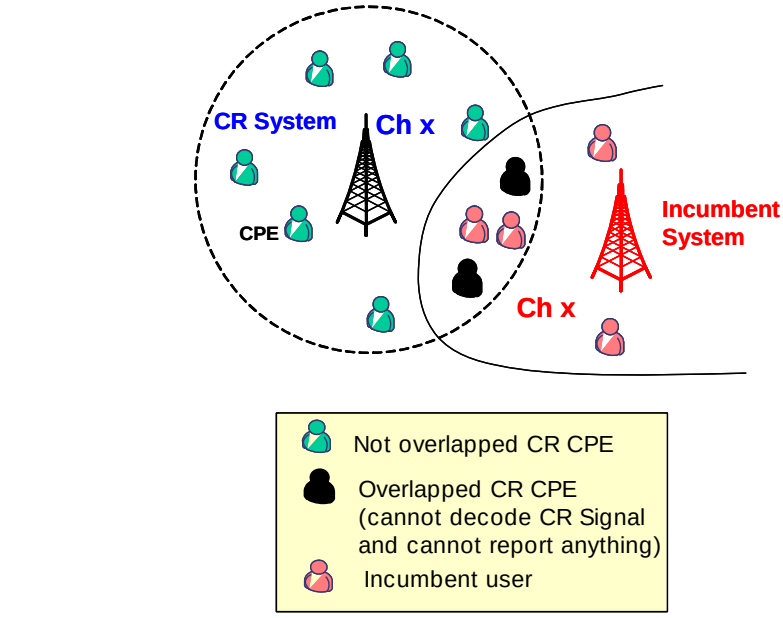

Figure 33 – Example of a hidden incumbent system

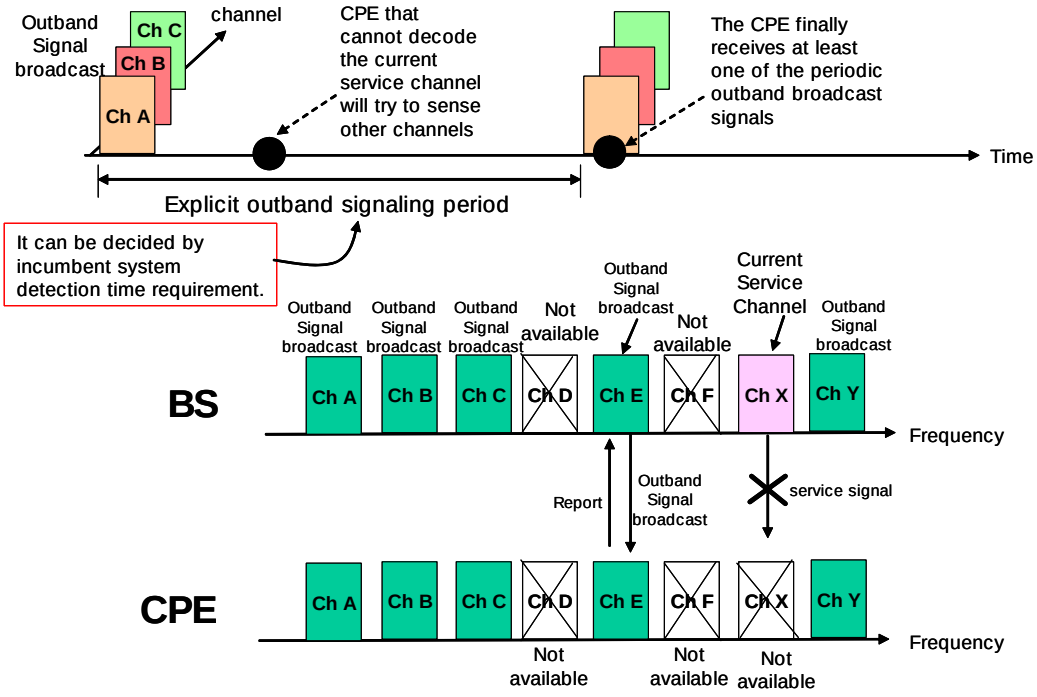

Figure 34 – Explicit outband signalling for hidden incumbent case detection

Explicit out-band broadcast signal follows the usual PHY and MAC frame architecture like in the service channel. Figure 35 shows the frame structure of explicit out-band broadcast signal. The SCH includes a flag to indicate that the current MAC frame is for regular service or for out-band broadcast signal. DS-Bust includes service channel information, such as service channel numbers and candidate channel numbers.

When a CPE receives a out-band signal and if the current service channel is not decodable by the CPE, then the CPE sends “Hidden incumbent report” to BS using the broadcasting US-Burst. Hidden incumbent report can

indicate “the current service channel x is not decodable” and the report can indicate “the current service channel x is used by an incumbent system” if the CPE can recognize incumbent signal. Also, the report can include sensing result for some other channels.

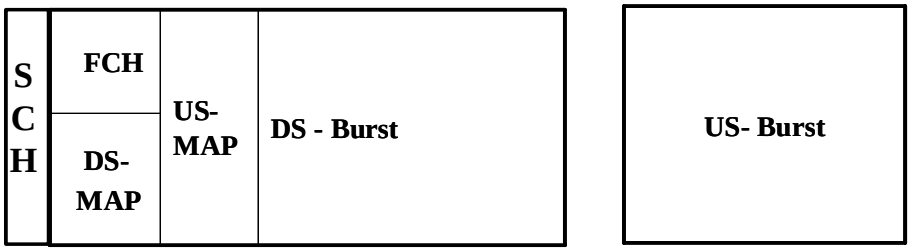

Figure 35 – The outband signalling uses the same frame structure, but with a slight change to the SCH

BS can allocate upstream resource for CPE’s “Hidden incumbent report” using one of the following two methods. First method is that BS allocates explicit resource to each CPE after CPE initialisation procedure. BS recognizes CPEs that want to send the reports and then allocates dedicated resource (upstream burst). This method requires additional overhead. Second method is that BS divides US resource into US-Burst slots for all unknown CPEs according to maximum hidden incumbent report size. Divided each US-Burst slot is indicated in US-MAP. When a CPE sends the hidden incumbent report, the CPE randomly selects a US-Burst, and then the CPE sends the report with the selected US-Burst.

16.5 Frequency Hopping

Based on the foundation provided in CMAC, it can also support frequency hopping which is treated here as an implementation issue. Frequency hopping may be useful when the number of vacant channels is greater than the number of 802.22 networks in an area. In this scenario, frequency hopping can be used to, for example, avoid quiet periods or provide better QoS to certain traffic types (e.g., voice).

For self-coexistence purposes, a WRAN shall not hop to a frequency channel that is occupied by another WRAN. WRANs that are operating in the same channel for a certain amount of time have knowledge of each other (e.g., through CBP). This serves to foster better self-coexistence and also incumbent protection (e.g., quiet period synchronization). Therefore, disallowing WRANs to hop into frequency channels that are already occupied by other WRANs aims at avoiding such scenarios. In fact, in these scenarios frequency hopping can be worse than staying in the current channel.

For the frequency hopping mechanism to be carried out, the WRAN must have accurate and up-to-date information on the spectrum occupancy. In particular, the WRAN shall be able to meet the DFS requirements for the protection of incumbents.

In summary, a WRAN network shall observe the following requirements before hopping to a new channel, say, the Backup Channel:

- The Backup Channel evaluation meets the Channel Availability Check Time requirement as specified in [3];
- The Backup Channel is not occupied by any incumbent;
- The Backup Channel is not occupied by another 802.22 network;
- Adjacent channels to the Backup Channel have also been checked for the presence of incumbents, so as to meet the EIRP profile.

Provided the aforementioned requirements are observed, frequency hopping can be implemented.

17. Ranging

To deal with the large propagation delays and varying RF signal quality between CPEs and the BS, CMAC incorporates a *ranging* procedure. Ranging is a collection of processes by which the CPE and BS maintain the quality of the RF communication link between them. Distinct processes are used for managing downstream and upstream.

17.1 Downstream Management

To maintain efficient operation between the BS and CPEs, the downstream burst profile is determined by the BS according to the quality of the signal that is received by each CPE. To reduce the volume of upstream traffic, the CPE monitors the CINR and compares the average value against the allowed range of operation. As shown in Figure 36, threshold levels bound this region. These thresholds parameters are specified in the DCD message, and shall be used by CPEs to determine their optimal burst profile. If the received CINR falls outside of the allowed operating region as determined by the threshold parameters, the CPE requests a change to a new burst profile using one of two methods:

1. If the CPE has been granted upstream bandwidth (a data grant allocation to the CPE's Basic CID), the CPE shall send a DBPC-REQ message in that allocation. The BS responds with a DBPC-RSP message.
2. If a grant is not available and the CPE requires a more robust burst profile on the downstream, the CPE shall send a RNG-REQ message in an Initial Ranging interval.

In either of these methods, the message is sent using the CPE's Basic CID. The coordination of message transmission and reception relative to actual change of modulation is different depending upon whether a CPE is transitioning to a more or less robust burst profile. Figure 37 shows the case where a CPE is transitioning to a more robust profile, while Figure 38 illustrates the transition to a less robust profile.

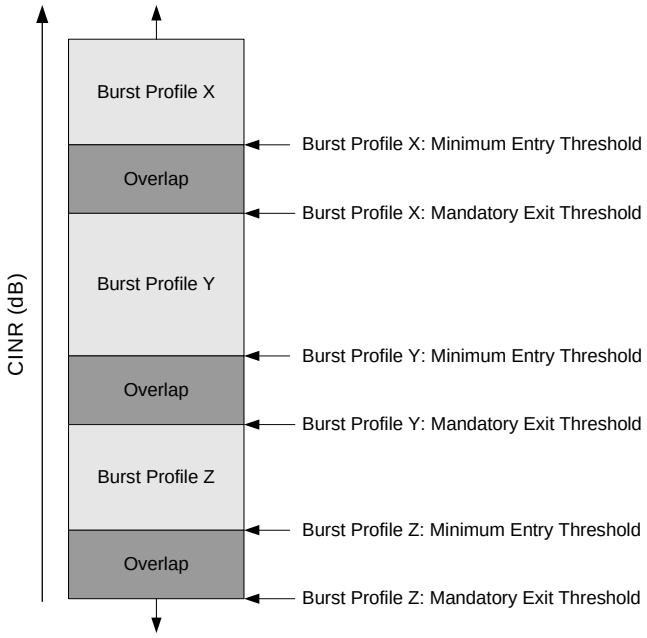

Figure 36 – Burst profiles and threshold utilization

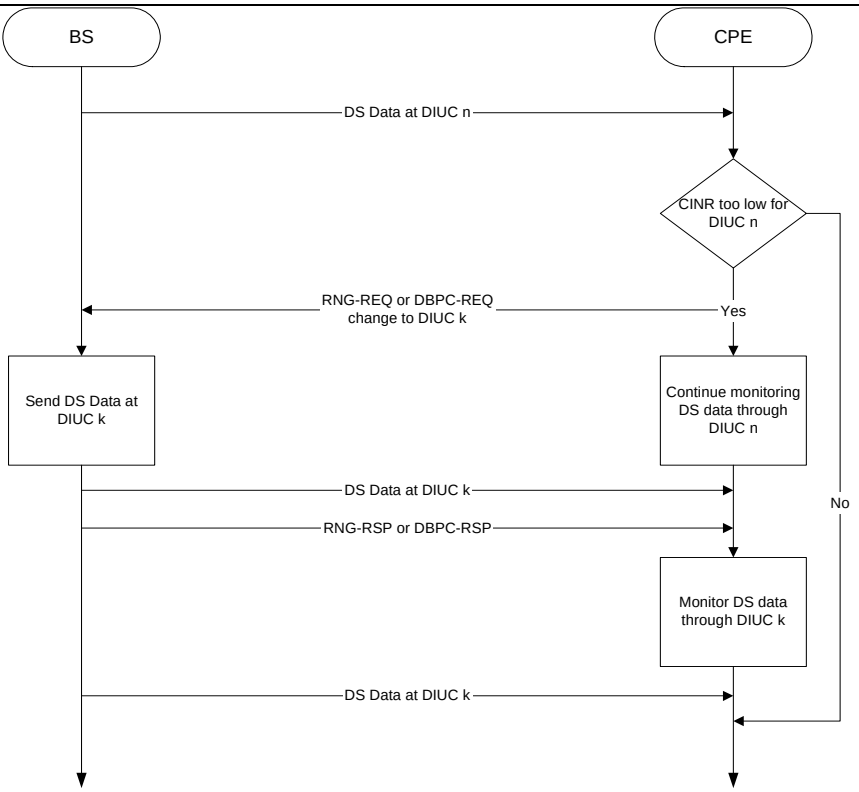

Figure 37 – Change to a more robust profile

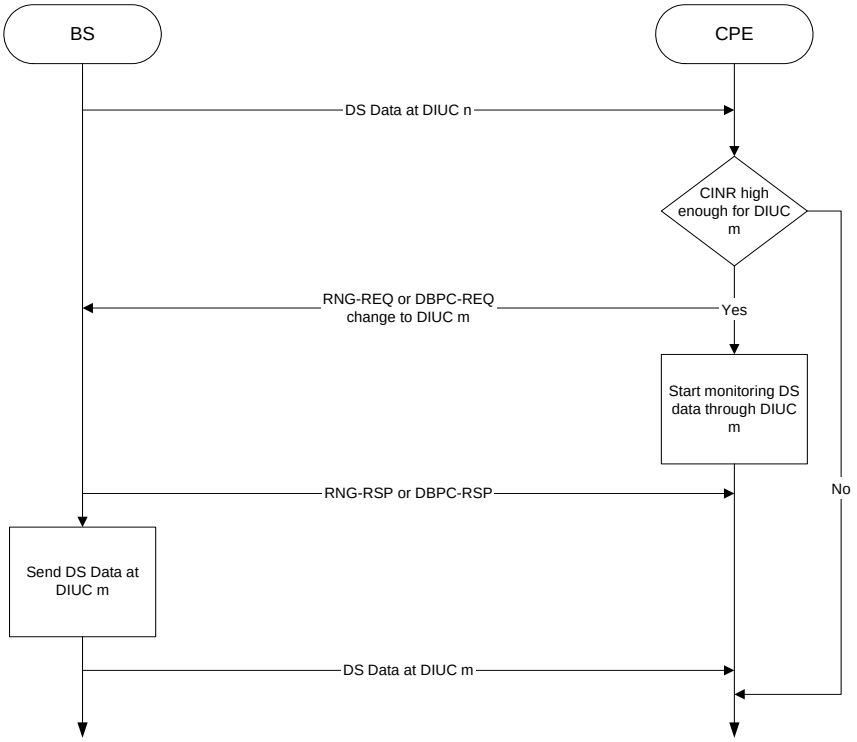

Figure 38 – Change to a less robust profile

17.2 Upstream Management

Upstream ranging management consists of two procedures: initial ranging and periodic ranging. Initial ranging (see 15) allows a CPE joining the network to acquire correct transmission parameters, such as time offset and Tx power level, so that the CPE can communicate with the BS. Following initial ranging, periodic ranging allows the CPE to adjust transmission parameters so that it can maintain upstream communications with the BS.

The following summarizes the general algorithm for periodic ranging.

1. For each CPE, the BS shall maintain a T27 timer (Table 226). At each expiration of the timer and provided the CPE is available (e.g., performing measurements), the BS shall grant bandwidth to the CPE for an upstream transmission. The timer is restarted each time a unicast grant is made to the CPE. As a result, as long as the CPE remains active, the BS does not specifically grant bandwidth to the CPE for a ranging opportunity.
2. Each CPE shall maintain a T4 timer (Table 226). The expiration of this timer indicates to the CPE that it has not been given the opportunity to transmit to the BS for an extended period of time. Operating on the assumption that its upstream transmission parameters are no longer usable, the CPE initiates a restart of its MAC operations.
3. For each unicast upstream burst grant, the BS determines whether or not a transmitted signal is present. If no signal is detected in a specified number of successive grants, the BS shall terminate link management for the associated CPE.
4. For each unicast upstream burst grant in which a signal is detected, the BS makes a determination as to the quality of the signal. If the signal is within acceptable limits and the data carried in the burst includes the RNG-REQ message, the RNG-RSP message shall be issued with a status of *success.* If the signal is not within acceptable limits, the RNG-RSP message shall be issued that includes the appropriate correction data and a status of *continue*. If a sufficient number of correction messages are issued without the CPE signal quality becoming acceptable, the BS shall send the RNG-RSP message with a status of *abort*, and terminate link management of the CPE.
5. The CPE shall process each RNG-RSP message it receives, implementing any PHY corrections that are specified (when the status is *continue)* or initiating a restart of MAC activities (when the status is *abort*).
6. The CPE should respond to each upstream bandwidth grant addressed to it, provided its transmission does not cause harmful interference to incumbents. When the status of the last RNG-RSP message received is *continue*, the RNG-REQ message shall be included in the transmitted burst. When the status of the last RNG-RSP message received is *success*, the CPE shall use the grant to service its pending upstream data queues. If no data is pending, the CPE shall keep quiet and does not transmit any data, as to keep interference at lower levels and hence improve coexistence.

18. Channel Descriptor Management

As previously presented, channel descriptor messages (i.e., DCD and UCD) are broadcast by the BS to all CPEs at periodic intervals. Among other things, these channel descriptors define burst profiles which are used by US-MAP and DS-MAP messages for allocating upstream and downstream transmissions, respectively. Once broadcast by the BS and received by associated CPEs, a given channel descriptor shall remain valid until a new channel descriptor message with a different value for the Configuration Change Count field, is again broadcast by the BS. When this happens, this new channel descriptor shall overwrite all the information of the previous descriptor.

Once channel descriptors are known to all CPEs in an 802.22 cell, the BS shall set the UCD/DCD Count value, contained in US-MAP and DS-MAP messages, equal to the Configuration Change Count of the desired channel descriptor. This way, a BS can easily indicate to the CPEs which burst profile is to be used for a given allocation, and hence provides high flexibility to the BS in controlling which burst profile to use at any given time by simply changing the UCD/DCD Count value.

Figure 39 describes the procedure to migrate from one upstream channel descriptor to the next, while Figure 40 focuses on the same procedure but for the downstream channel.

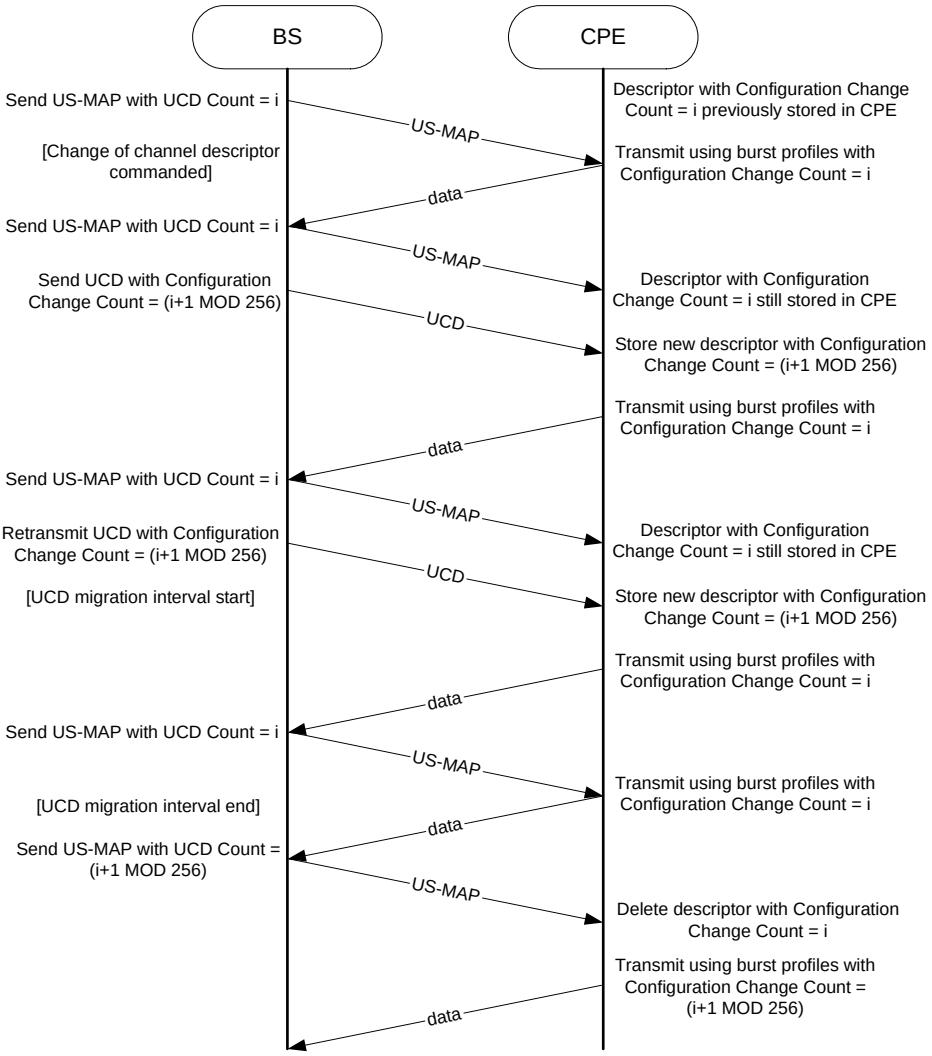

Figure 39 – UCD channel descriptor update

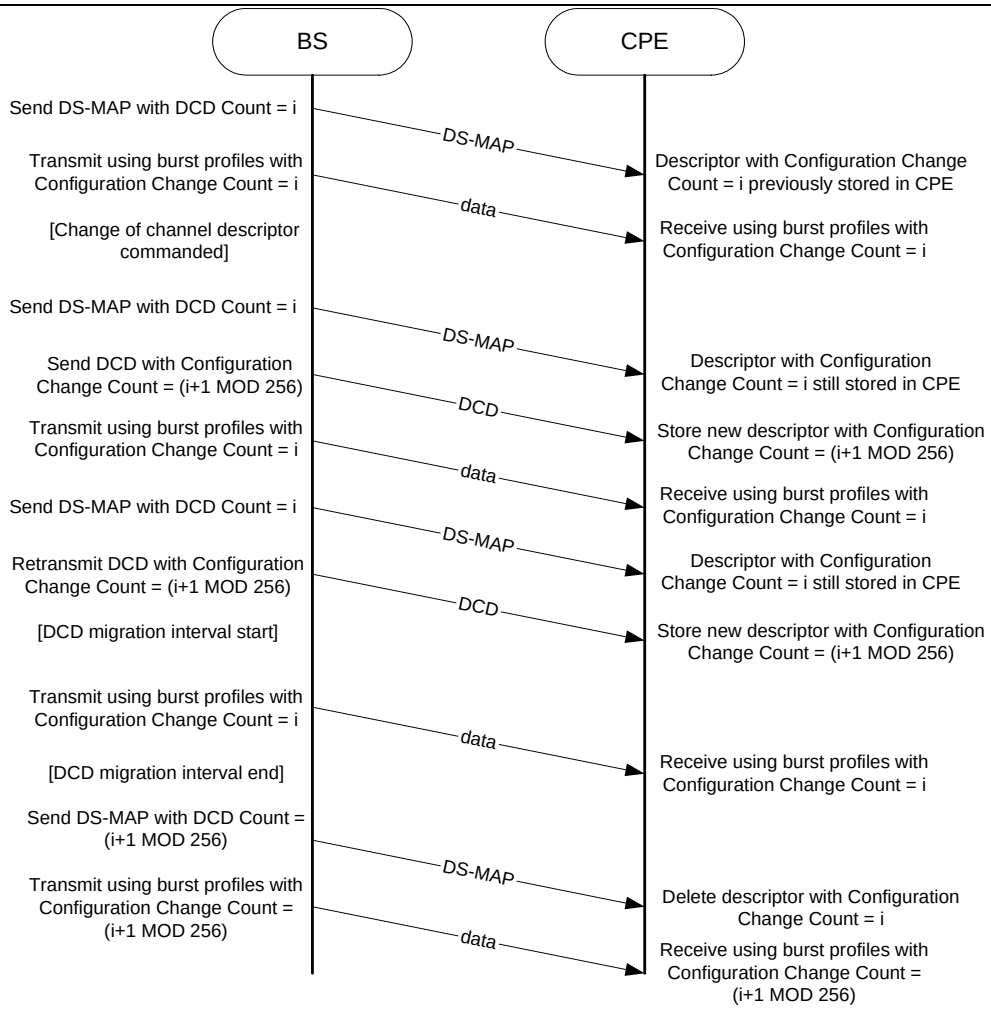

Figure 40 – DCD channel descriptor update

Finally, note that the Configuration Change Count shall be incremented by 1 modulo 256 for every new migration of channel descriptor. After issuing a DS-MAP or US-MAP message with the Configuration Change Count equal to that of the new generation, the old channel descriptor ceases to exist and the BS shall not refer to it anymore. When migrating from one generation to the next, the BS shall schedule the transmissions of the UCD and DCD messages in such a way that each CPE has the possibility to successfully hear it at least once.

19. Multicast Support

Multicast support is an important and integral part of CMAC. In CMAC, multicast groups are used not only for their traditional application of data delivery (e.g., streaming), but also for sending management commands to a set of CPEs. For example, the BS may wish to implement clustering algorithms for measurements and use the feature of multicast group to create such clusters. In this case, the BS could, for instance, simultaneously address a set of CPEs and share the load of measurements across clusters. That is, the BS could make certain clusters responsible for DTV measurements while other clusters would target Part 74 services. Another possible use of multicast connections is for CBP (see 21.2.1). In this case, the BS can maximize the use of the Coexistence IUC by properly selecting the CPEs who will transmit CBP packets.

In order to support multicast services with the purpose of management, CMAC defines a special type of multicast connection named multicast management CID. In this section, we describe the multicast feature of CMAC.

19.1 Group Management

The BS may add a CPE to a multicast group by sending an MCA-REQ message with the Join command. Upon receiving an MCA-REQ message, the CPE shall respond by sending an MCA-RSP message. A similar procedure is employed in the case of leaving a group. The protocol is shown in Figure 41 and Figure 42.

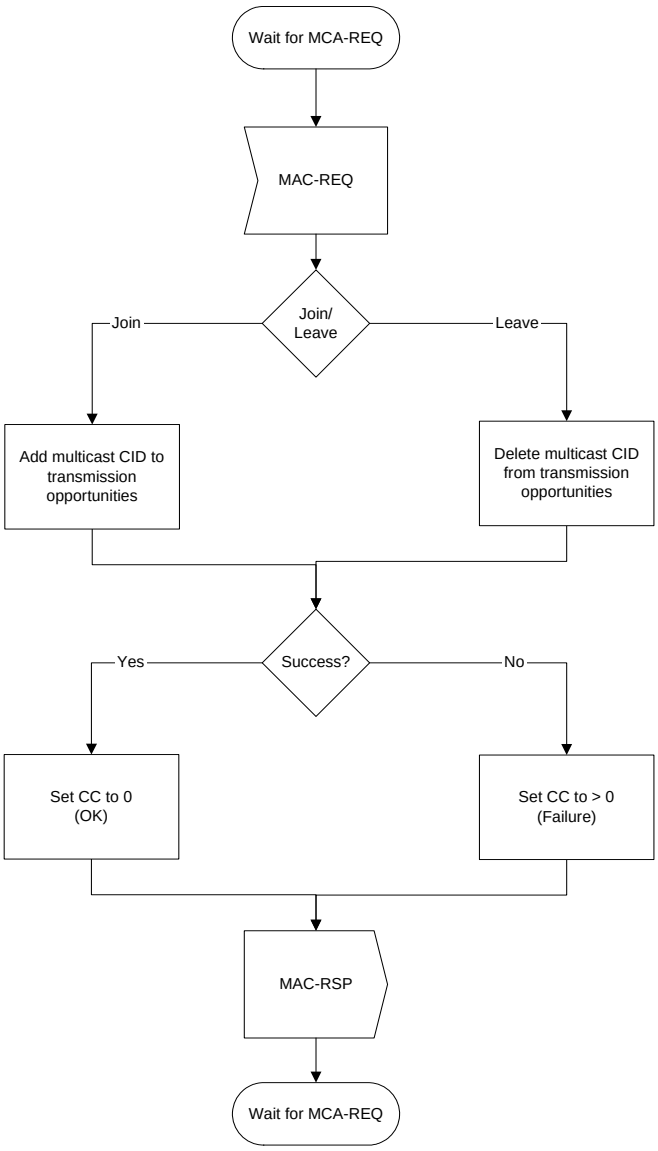

Figure 41 – Group management at CPE

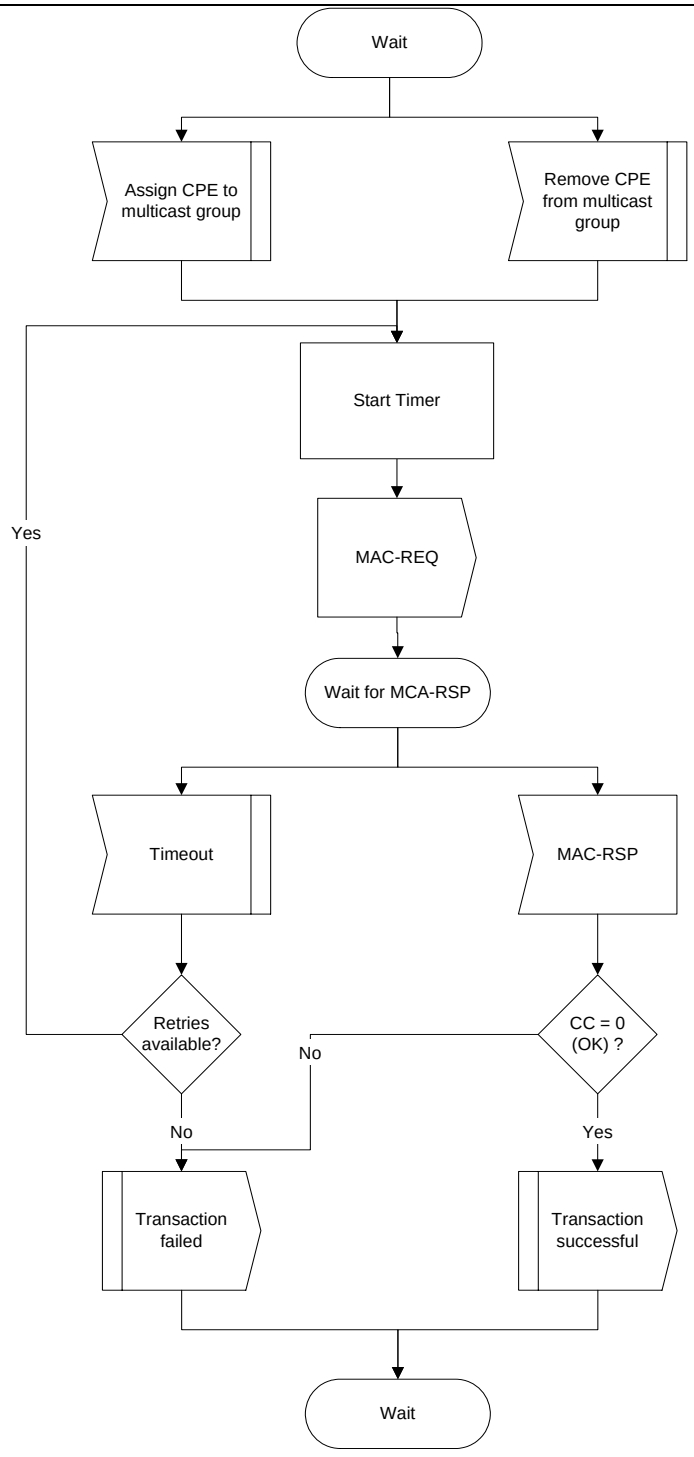

Figure 42 – Group management at BS

19.2 Multicast Connections

The BS may establish a downstream multicast service by creating a connection with each CPE to be associated with the service. Any available transport or multicast management CID value may be used for the service (i.e., there are no dedicated transport type of CIDs for multicast connections). To ensure proper multicast operation, the

CID used for the service is the same for all CPEs on the same channel that participate in the connection. Except for multicast management CIDs, the CPEs need not be aware that the connection is a multicast connection. The data transmitted on the connection with the given CID shall be received and processed by the MAC of each involved CPE. Since a multicast connection is associated with a service flow, it is associated with the QoS and traffic parameters for that service flow.

ARQ is not applicable to multicast connections.

If a downstream multicast connection is to be encrypted, each CPE participating in the connection shall have an additional security association (SA), allowing that connection to be encrypted using keys that are independent of those used for other encrypted transmissions between the CPEs and the BS.

20. QoS

CMAC adopts a similar QoS service model as specified in IEEE 802.16 (IEEE Std 802.16^{TM}-2004 [5]). It defines several QoS related concepts, which include the following:

a. Service Flow QoS Scheduling.
b. Dynamic Service Establishment.
c. Two-phase Activation Model.

20.1 Theory of Operation

The various protocol mechanisms described in this document may be used to support QoS for both upstream and downstream traffic through the CPE and the BS. This subclause provides an overview of the QoS protocol mechanisms and their part in providing end-to-end QoS.

The requirements for QoS include the following:

a. A configuration and registration function for preconfiguring CPE-based QoS service flows and traffic parameters.
b. A signaling function for dynamically establishing QoS-enabled service flows and traffic parameters.
c. Utilization of MAC scheduling and QoS traffic parameters for upstream service flows
d. Utilization of QoS traffic parameters for downstream service flows.
e. Grouping of service flow properties into named Service Classes, so upper-layer entities and external applications (at both the CPE and BS) may request service flows with desired QoS parameters in a globally consistent way.

The principal mechanism for providing QoS is to associate packets traversing the MAC interface into a service flow as identified by the CID. A service flow is a unidirectional flow of packets that is provided a particular QoS. The CPE and BS provide this QoS according to the QoS Parameter Set defined for the service flow.

The primary purpose of the QoS features defined here is to define transmission ordering and scheduling on the air interface. However, these features often need to work in conjunction with mechanisms beyond the air interface in order to provide end-to-end QoS or to police the behavior of CPEs.

Service flows exist in both the upstream and downstream direction and may exist without actually being activated to carry traffic. All service flows have a 32-bit SFID; admitted and active service flows also have a 16-bit CID.

20.2 Service Flows

A service flow is a MAC transport service that provides unidirectional transport of packets either to upstream packets transmitted by the CPE or to downstream packets transmitted by the BS[10]. A service flow is characterized by a set of QoS Parameters such as latency, jitter, and throughput assurances. In order to standardize operation between the CPE and BS, these attributes include details of how the CPE requests upstream bandwidth allocations and the expected behavior of the BS upstream scheduler.

A service flow is partially characterized by the following attributes[11]:

a. SFID: An SFID is assigned to each existing service flow. The SFID serves as the principal identifier for the service flow in the network. A service flow has at least an SFID and an associated direction.
b. CID: Mapping to an SFID that exists only when the connection has an admitted or active service flow.
c. ProvisionedQoSParamSet: A QoS parameter set provisioned via means outside of the scope of this standard, such as the network management system.
d. AdmittedQoSParamSet: Defines a set of QoS parameters for which the BS (and possibly the CPE) are reserving resources. The principal resource to be reserved is bandwidth, but this also includes any other memory or time-based resource required to subsequently activate the flow.
e. ActiveQoSParamSet: Defines a set of QoS parameters defining the service actually being provided to the service flow. Only an Active service flow may forward packets.
f. Authorization Module: A logical function within the BS that approves or denies every change to QoS Parameters and Classifiers associated with a service flow. As such, it defines an "envelope" that limits the possible values of the AdmittedQoSParamSet and ActiveQoSParamSet.

The relationship between the QoS Parameter Sets is as shown in Figure 43 and Figure 44. The ActiveQoSParamSet is always a subset[12] of the AdmittedQoSParamSet, which is always a subset of the authorized "envelope." In the dynamic authorization model, this envelope is determined by the Authorization Module (labeled as the AuthorizedQoSParamSet). In the provisioned authorization model, this envelope is determined by the ProvisionedQoSParamSet. It is useful to think of three types of service flows:

a. Provisioned: This type of service flow is known via provisioning by, for example, the network management system. Its AdmittedQoSParamSet and ActiveQoSParamSet are both null.
b. Admitted: This type of service flow has resources reserved by the BS for its AdmittedQoSParamSet, but these parameters are not active (i.e., its ActiveQoSParamSet is null). Admitted Service Flows may have been provisioned or may have been signaled by some other mechanism.
c. Active: This type of service flow has resources committed by the BS for its ActiveQoSParamSet, (e.g., is actively sending maps containing unsolicited grants for a UGSbased service flow). Its ActiveQoSParamSet is non-null.

[10] A service flow, as defined here, has no direct relationship to the concept of a "flow" as defined by the IETF Integrated Services (intserv) Working Group (IETF RFC 2212). An intserv flow is a collection of packets sharing transport-layer endpoints. Multiple intserv flows can be served by a single service flow.

[11] Some attributes are derived from the above attribute list. The Service Class Name is an attribute of the ProvisionedQoSParamSet. The activation state of the service flow is determined by the ActiveQoSParamSet. If the ActiveQoSParamSet is null, then the service flow is inactive.

[12] To say that QoS Parameter Set A is a subset of QoS Parameter Set B, the following shall be true for all QoS Parameters in A and B:
if (a smaller QoS parameter value indicates less resources, e.g., Maximum Traffic Rate)
A is a subset of B if the parameter in A is less than or equal to the same parameter in B
if (a larger QoS parameter value indicates less resources, e.g., Tolerated Grant Jitter)
A is a subset of B if the parameter in A is greater than or equal to the same parameter in B
if (the QoS parameter is not quantitative, e.g., Service Flow Scheduling Type)
A is a subset of B if the parameter in A is equal to the same parameter in B

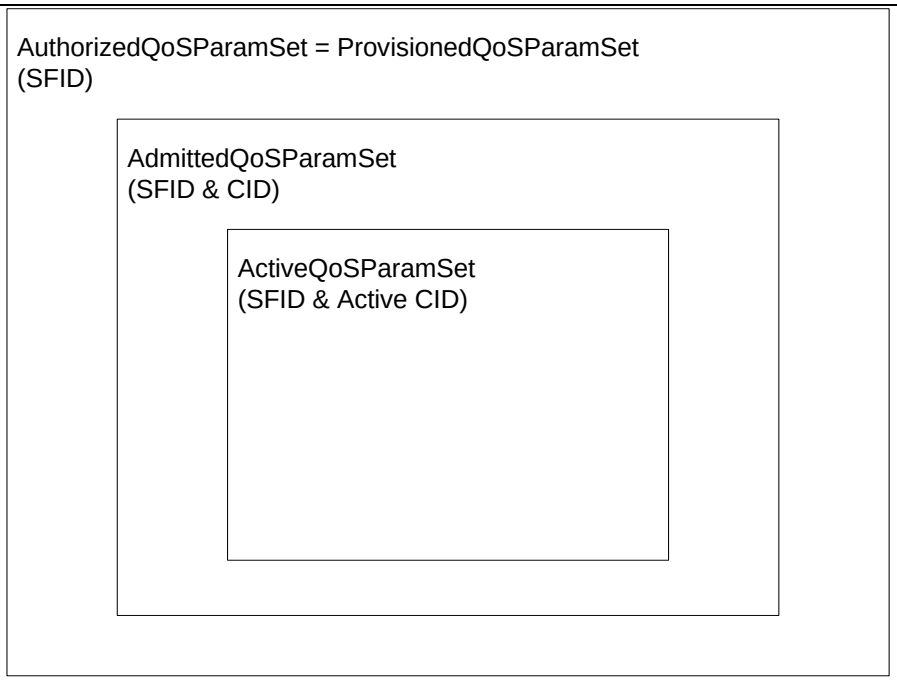

Figure 43 – Provisioned authorization model "envelopes"

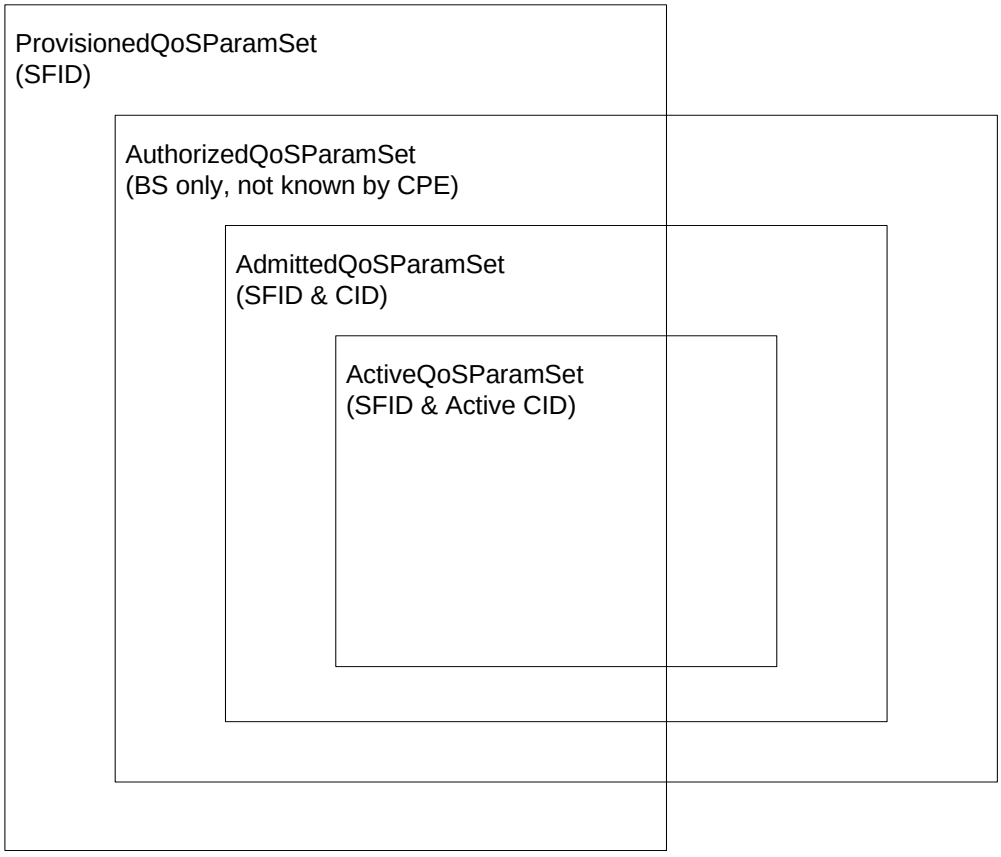

Figure 44 – Dynamic authorization model "envelopes"

20.3 Object Model

The major objects of the architecture are represented by named rectangles in Figure 45. Each object has a number of attributes; the attribute names that uniquely identify it are marked with an "*". Optional attributes are denoted with brackets. The relationship between the number of objects is marked at each end of the association line between the objects. For example, a service flow may be associated with from 0 to *N* (many) PDUs, but a PDU is associated with exactly one service flow. The service flow is the central concept of the MAC protocol. It is uniquely identified by a 32-bit (SFID). Service flows may be in either the upstream or downstream direction. Admitted and active service flows are mapped to a 16-bit CID.

Outgoing user data is submitted to the MAC SAP by a CS process for transmission on the MAC interface. The information delivered to the MAC SAP includes the CID identifying the connection across which the information is delivered. The service flow for the connection is mapped to MAC connection identified by the CID.

The Service Class is an optional object that may be implemented at the BS. It is referenced by an ASCII name, which is intended for provisioning purposes. A Service Class is defined in the BS to have a particular QoS Parameter Set. The QoS Parameter Sets of a service flow may contain a reference to the Service Class Name as a "macro" that selects all of the QoS parameters of the Service Class. The service flow QoS Parameter Sets may augment and even override the QoS parameter settings of the Service Class, subject to authorization by the BS.

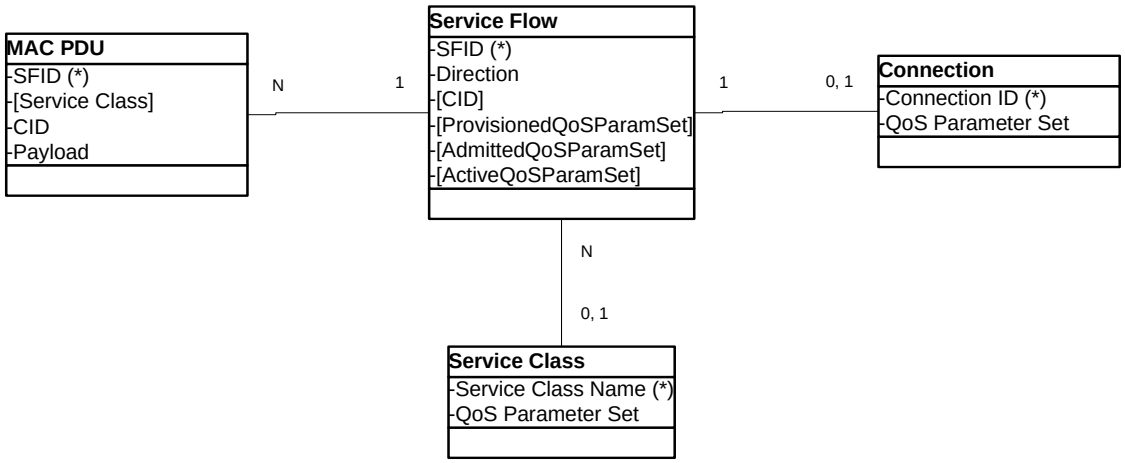

Figure 45 – Object model of the QoS service

20.4 Service Classes

The Service Class serves the following purposes[13]:

a. It allows operators, who so wish, to move the burden of configuring service flows from the provisioning server to the BS. Operators provision the CPEs with the Service Class Name; the implementation of the name is configured at the BS. This allows operators to modify the implementation of a given service to local circumstances without changing CPE provisioning. For example, some scheduling parameters may need to be tweaked differently for two different BSs to provide the same service. As another example, service profiles could be changed by time of day.

13 Service classes are merely identifiers for a specific set of QoS parameter set values. Hence, the use of service classes is optional. A service identified by a service class is treated no differently, once established, than a service that has the same QoS parameter set explicitly specified.

b. It allows higher-layer protocols to create a service flow by its Service Class Name. For example, telephony signaling may direct the CPE to instantiate any available provisioned service flow of class "G711."

Any service flow may have its QoS Parameter Set specified in any of three ways:

- By explicitly including all traffic parameters.
- By indirectly referring to a set of traffic parameters by specifying a Service Class Name.
- By specifying a Service Class Name along with modifying parameters.

The Service Class Name is "expanded" to its defined set of parameters at the time the BS successfully admits the service flow. The Service Class expansion can be contained in the following BS-originated messages: DSA-REQ, DSC-REQ, DSA-RSP, and DSC-RSP. In all of these cases, the BS shall include a service flow encoding that includes the Service Class Name and the QoS Parameter Set of the Service Class. If an CPE-initiated request contained any supplemental or overriding service flow parameters, a successful response shall also include these parameters.

When a Service Class name is given in an admission or activation request, it is possible that the returned QoS Parameter Set may change from activation to activation. This can happen because of administrative changes to the Service Class's QoS Parameter Set at the BS. If the definition of a Service Class Name is changed at the BS (e.g., its associated QoS Parameter Set is modified), it has no effect on the QoS Parameters of existing service flows associated with that Service Class. A BS may initiate DSC transactions to existing service flows that reference the Service Class Name to affect the changed Service Class definition.

When an CPE uses the Service Class Name to specify the Admitted QoS Parameter Set, the expanded set of TLV encodings of the service flow shall be returned to the CPE in the response message (DSA-RSP or DSC-RSP). Use of the Service Class Name later in the activation request may fail if the definition of the Service Class Name has changed and the new required resources are not available. Thus, the CPE should explicitly request the expanded set of TLVs from the response message in its later activation request.

20.5 Authorization

An authorization module shall approve every change to the service flow QoS Parameters. This includes every DSA-REQ message to create a new service flow and every DSC-REQ message to change a QoS Parameter Set of an existing service flow. Such changes include requesting an admission control decision (e.g., setting the AdmittedQoSParamSet) and requesting activation of a service flow (e.g., setting the ActiveQoSParamSet). The authorization module also checks reduction requests regarding the resources to be admitted or activated.

In the static authorization model, the authorization module stores the provisioned status of all "deferred" service flows. Admission and activation requests for these provisioned service flows shall be permitted, as long as the Admitted QoS Parameter Set is a subset of the Provisioned QoS Parameter Set, and the Active QoS Parameter Set is a subset of the Admitted QoS Parameter Set. Requests to change the Provisioned QoS Parameter Set shall be refused, as shall requests to create new dynamic service flows. This defines a static system where all possible services are defined in the initial configuration of each CPE.

In the dynamic authorization model, the authorization module also communicates through a separate interface to an independent policy server. This policy server may provide the authorization module with advance notice of upcoming admission and activation requests, and it specifies the proper authorization action to be taken on those requests. Admission and activation requests from an CPE are then checked by the Authorization Module to ensure that the ActiveQoSParamSet being requested is a subset of the set provided by the policy server. Admission and activation requests from an CPE that are signaled in advance by the external policy server are permitted.

Admission and activation requests from an CPE that are not pre-signaled by the external policy server may result in a real-time query to the policy server or may be refused.

Prior to initial connection setup, the BS shall retrieve the Provisioned QoS Set for an CPE. This is handed to the Authorization Module within the BS. The BS shall be capable of caching the Provisioned QoS Parameter Set and shall be able to use this information to authorize dynamic flows that are a subset of the Provisioned QoS Parameter Set. The BS should implement mechanisms for overriding this automated approval process (such as described in the dynamic authorization model). For example it could:

a. Deny all requests whether or not they have been preprovisioned.
b. Define an internal table with a richer policy mechanism but seeded by the Provisioned QoS Set.
c. Refer all requests to an external policy server.

20.6 Types of Service Flows

It is useful to think about three basic types of service flows. This subclause describes these three types of service flows in more detail. However, it is important to note that there are more than just these three basic types (see 8.8.10.4).

20.6.1 Provisioned

A service flow may be provisioned but not immediately activated (sometimes called “deferred”). That is, the description of any such service flow contains an attribute that provisions but defers activation and admission (see 8.8.10.4). The network assigns a SFID for such a service flow. The BS may also require an exchange with a policy module prior to admission.

As a result of external action beyond the scope of this specification, the CPE may choose to activate a provisioned service flow by passing the SFID and the associated QoS Parameter Sets to the BS in the DSC-REQ message. If authorized and resources are available, the BS shall respond by mapping the service flow to a CID.

As a result of external action beyond the scope of this specification, the BS may choose to activate a service flow by passing the SFID as well as the CID and the associated QoS Parameter Sets to the CPE in the DSCREQ message. Such a provisioned service flow may be activated and deactivated many times (through DSC exchanges). In all cases, the original SFID shall be used when reactivating the service flow.

20.6.2 Admitted

This protocol supports a two-phase activation model that is often utilized in telephony applications. In the two-phase activation model, the resources for a “call” are first “admitted,” and then once the end-to-end negotiation is completed (e.g., called party’s gateway generates an “off-hook” event), the resources are “activated.” The two-phase model serves the following purposes:

a. Conserving network resources until a complete end-to-end connection has been established;
b. Performing policy checks and admission control on resources as quickly as possible, and in particular, before informing the far end of a connection request; and
c. Preventing several potential theft-of-service scenarios.

For example, if an upper-layer service were using UGS, and the addition of upper-layer flows could be adequately provided by increasing the Maximum Sustained Traffic Rate QoS parameter, then the following procedure might be used. When the first higher-layer flow is pending, the CPE issues a DSA-REQ with the admitted Maximum

Sustained Traffic Rate parameter equal to that required for one higher-layer flow, and the active Maximum Sustained Traffic Rate parameter equal to zero. Later when the higher-layer flow becomes active, it issues a DSC-REQ with the instance of the active Maximum Sustained Traffic Rate parameter equal to that required for one higher-layer flow. Admission control was performed at the time of the reservation, so the later DSC-REQ, having the active parameters within the range of the previous reservation, is guaranteed to succeed. Subsequent higher-layer flows would be handled in the same way. If there were three higher-layer flows establishing connections, with one flow already active, the service flow would have admitted Maximum Sustained Traffic Rate equal to that required for four higher-layer flows, and active Maximum Sustained Traffic Rate equal to that required for one higher-layer flow.

An activation request of a service flow where the new ActiveQoSParamSet is a subset of the AdmittedQoSParamSet shall be allowed, except in the case of catastrophic failure. An admission request where the AdmittedQoSParamSet is a subset of the previous AdmittedQoSParamSet, so long as the ActiveQoSParamSet remains a subset of the AdmittedQoSParamSet, shall succeed.

A service flow that has resources assigned to its AdmittedQoSParamSet, but whose resources are not yet completely activated, is in a transient state. It is possible in some applications that a long-term reservation of resources is necessary or desirable. For example, placing a telephone call on hold should allow any resources in use for the call to be temporarily allocated to other purposes, but these resources shall be available for resumption of the call later. The AdmittedQoSParamSet is maintained as "soft state" in the BS; this state shall be maintained without releasing the nonactivated resources. Changes may be signaled with a DSC-REQ message.

20.6.3 Active

A service flow that has a non-NULL ActiveQoSParamSet is said to be an active service flow. It is requesting (according to its Request/Transmission Policy, as in 8.8.10.12) and being granted bandwidth for transport of data packets. An admitted service flow may be activated by providing an ActiveQoSParamSet, signaling the resources actually desired at the current time. This completes the second stage of the two-phase activation model (see 20.6.2).

A service flow may be provisioned and immediately activated. Alternatively, a service flow may be created dynamically and immediately activated. In this case, two-phase activation is skipped and the service flow is available for immediate use upon authorization.

20.7 Service Flow Creation

The provisioning of service flows is done via means outside of the scope of this standard, such as the network management system. During provisioning, a service flow is instantiated, gets a service flow ID and a "provisioned" type. For some service flows it may be specified that DSA procedure must be activated by Network Entry procedure. Enabling service flows follows the transfer of the operational parameters (see Figure 23). In this case, the service flow type may change to "admitted" or to "active;" in the latter case, the Service Flow is mapped onto a certain connection.

Service flow encodings contain either a full definition of service attributes (omitting defaultable items if desired) or a service class name. A service class name is an ASCII string, which is known at the BS and which indirectly specifies a set of QoS Parameters.

Triggers, other than network entry, also may cause creation, admission, or activation of service flows. Such triggers lay outside the scope of the standard.

Capability of handling each specific Service Flow parameter is optional.

20.7.1 Dynamic Service Flow Creation

19.7.1.1 CPE-Initiated

Creation of service flows may be initiated by either BS (mandatory capability) or by CPE (optional capability).

The CPE-initiated protocol is illustrated in Figure 46 and described in detail in 20.9. A DSA-REQ from an CPE contains a service flow reference and QoS Parameter set (marked either for admission-only or for admission and activation). BS responds with DSA-RSP indicating acceptance or rejection. In the case when rejection was caused by presence of non-supported parameter of non-supported value, specific parameter may be included into DSA-RSP.

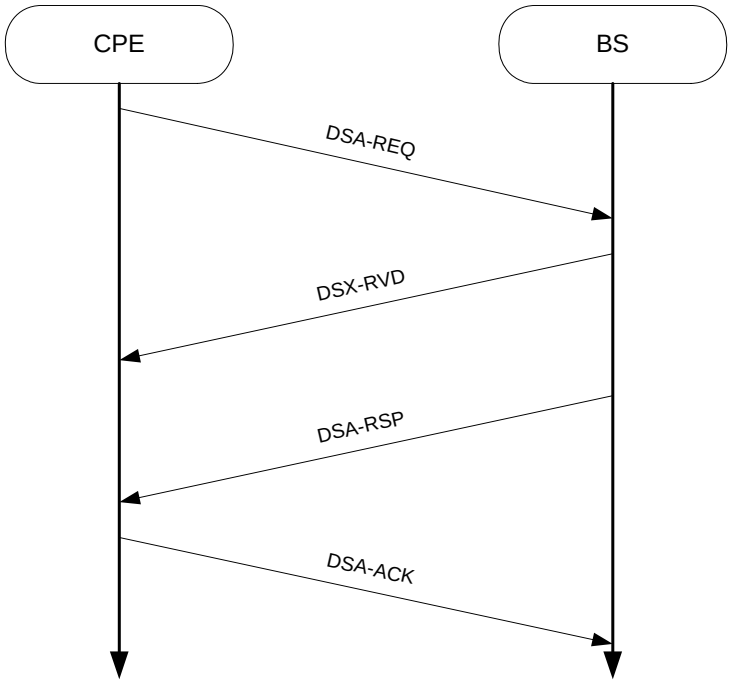

Figure 46 – DSA message flow (CPE-initiated)

19.7.1.2 BS-Initiated

A DSA-REQ from a BS contains an SFID for either one upstream or one downstream Service flow, possibly its associated CID, and a set of active or admitted QoS Parameters. The protocol is illustrated in Figure 47 and is described in detail in 20.9. The CPE responds with DSA-RSP indicating acceptance or rejection. In the case when rejection was caused by presence of non-supported parameter of non-supported value, specific parameter may be included into DSA-RSP.

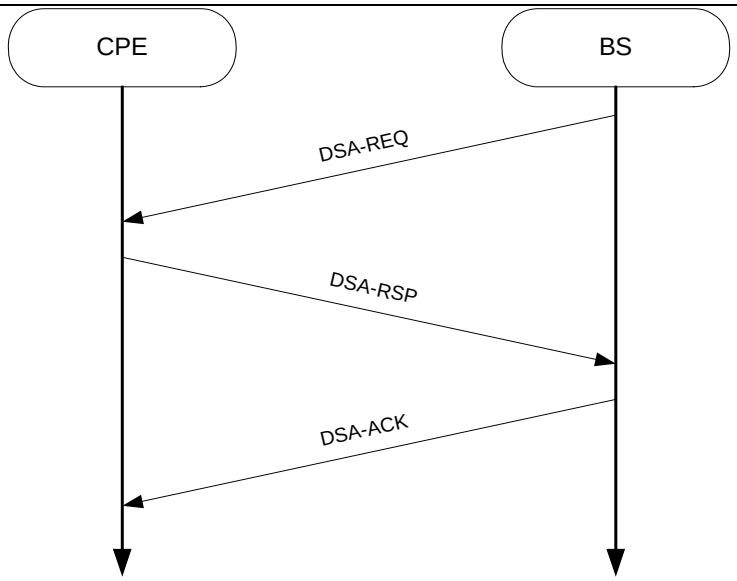

Figure 47 – DSA message flow (BS-initiated)

20.8 Dynamic Service Flow Modification and Deletion

In addition to the methods presented in 20.7.1 for creating service flows, protocols are defined for modifying and deleting service flows (see 20.9).

Both provisioned and dynamically created service flows are modified with the DSC message, which can change the Admitted and Active QoS Parameter sets of the flow. A successful DSC transaction changes a service flow's QoS parameters by replacing both the Admitted and Active QoS parameter sets. If the message contains only the Admitted set, the Active set is set to null and the flow is deactivated. If the message contains neither set ("000" value used for QoS Parameter Set type – see 8.8.10.4), then both sets are set to null and the flow is de-admitted. When the message contains both QoS parameter sets, the Admitted set is checked first, and if admission control succeeds, the Active set in the message is checked against the Admitted set in the message to ensure that it is a subset. If all checks are successful, the QoS parameter sets in the message become the new Admitted and Active QoS parameter sets for the service flow. If either of the checks fails, the DSC transaction fails and the service flow QoS parameter sets remain unchanged.

20.9 Service Flow Management

The service flow management of CMAC is similar to that of IEEE 802.16 (IEEE Std 802.16^{TM}-2004). Please refer to [5] for further information.

21. Coexistence

Coexistence is critical for the 802.22 air interface, which is required to include incumbent detection and protection mechanisms as well as self-coexistence measures in the very conception of the standard. To address incumbents, CRs techniques are incorporated into CMAC by means of distributed spectrum sensing, measurements, detection algorithms, and spectrum management. With regards to self-coexistence, the CBP is proposed which relies on

beacons to achieve efficient coexistence amongst overlapping 802.22 cells. The combination of these mechanisms forms a MAC layer that is highly flexible and adaptive to the environment, and can react to sudden changes in it[14].

Therefore, in this section we discuss the coexistence aspects of CMAC in order to protect incumbents, and also to address self-coexistence (i.e., coexistence with itself). In addition to this, CMAC supports advanced clustering schemes to be implemented in order to optimise the distributed measurement activity, and so clustering support is also covered.

21.1 Incumbents

CMAC provides all the capabilities for the effective detection and protection of incumbent services. A comprehensive set of measurement and spectrum management commands is available, which gives the BS the necessary flexibility to manage CPEs and obtain a reliable spectrum occupancy map of its cell and, if needed, change its operating parameters.

CPEs also have available plenty of ways to report measured information to the BS. In addition to a vast pool of MAC management frames, urgent coexistence situations can also be reported either through fields located in the general MAC header itself, or through the UCS. The simplified lifecycle of a single measurement activity is depicted in Figure 48. Each of the phases of this lifecycle requires special handling, and specific protocols and algorithms are proposed in this section to address their requirements. In the following subsections, we provide a detailed overview of the mechanisms available in CMAC for the management of incumbent measurements throughout all its lifecycle.

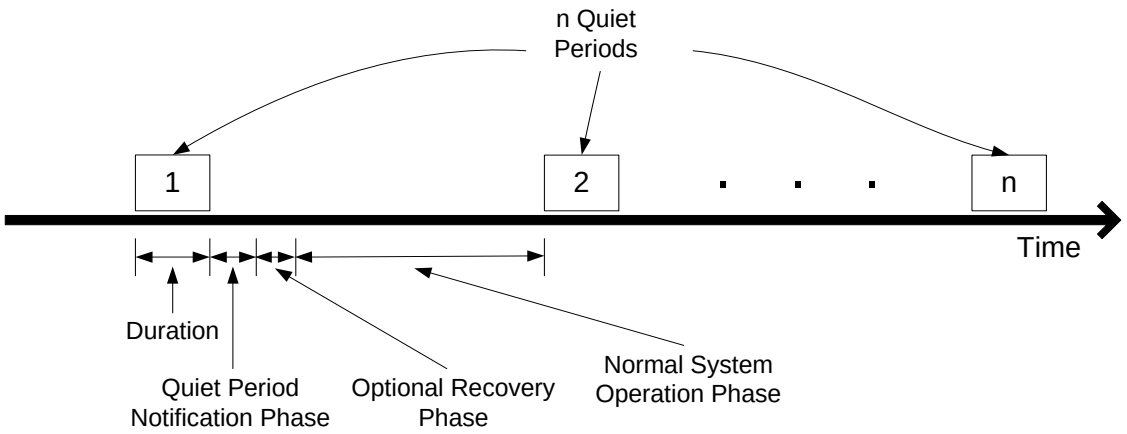

Figure 48 – Lifecycle of a measurement activity

21.1.1 Measurements Classification

Measurements in CMAC can be of types: in-band and out-of-band. In-band measurements refer to the case when the stations sense the same channels used for normal cell operation. For example if a BS uses channel N to communication with its CPEs, in-band measurement refers to the incumbent sensing activity which is performed in those channels where 802.22 transmissions have a direct effect (e.g., N-t through N+t, where, say, $t \leq 1$), since these are directly affected by the 802.22 operation in channel N. Similarly, out-of-band measurements refer to the case when the incumbent sensing activity is carried out in those channels other than N-t through N+t.

[14] Note that all this is in addition to TPC, which was discussed earlier and is also supported by CMAC.

It is important to note, however, that in-band and out-of-band measurements have a different meaning when used in the context of CBP measurements. For beacon measurements, all channels other than channel N are classified as being out-of-band rather than in-band since operation in these channels is not prohibited, as it is the case with incumbent protection.

Finally, note that for in-band measurements quiet periods shall be required, which is not the case for out-of-band measurements.

21.1.2 Measurements Management

CMAC supports a hierarchical measurement philosophy implemented by three management messages, namely, BLM-REQ, BLM-RSP, BLM-REP, and BLM-ACK (see Table 24 and 8.22)[15]. These management messages are used between the BS and CPEs to perform a wide range of measurement activities, either related to incumbents or to self-coexistence. With these messages, both in-band and out-of-band measurements can be performed.

In a single BLM-REQ command, the BS may simultaneously request CPEs to perform several types of measurements in a number of channels. Thus, a BLM-REQ is formed by a collection of single measurement requests. Each single measurement request specifies several parameters such as the frequency with which the BS wants CPEs to report back to it or if reports are to be autonomous. Furthermore, single measurement requests also define timing parameters, as illustrated in Figure 8. Upon receiving BLM-REQ message, the CPE shall examine this message's header and determine whether it is required to respond back with a BLM-RSP message. In all cases, the CPE shall carry out all the measurements as requested by the BS, if these are supported. CPEs shall report back to the BS with a BLM-REP message which contains measurements of what has been requested by the BS in the corresponding BLM-REQ message. These reports shall be sent with the periodicity specified by the BS in the corresponding BLM-REQ message. Once the measurement report message is successfully received at the BS, the BS shall respond back to the CPE with a BLM-ACK message to acknowledge its reception. In case the CPE does not hear the BLM-ACK message from the BS after some pre-specified timeout T28 (see Table 226), it shall assume that its BLM-REP message was lost and shall initiate retransmission of the BLM-REP message. The CPE shall attempt retransmission or measurement report messages up to BLM-REP retries (see Table 226). Once the BLM-ACK message is successfully received at the CPE, it shall then clear its local statistics to prepare for future measurements. Figure 49 illustrates the measurement message flow between the BS and a CPE.

[15] Since the correct reception of these messages is critical for coexistence purposes, they should be transmitted with a more robust modulation/coding.

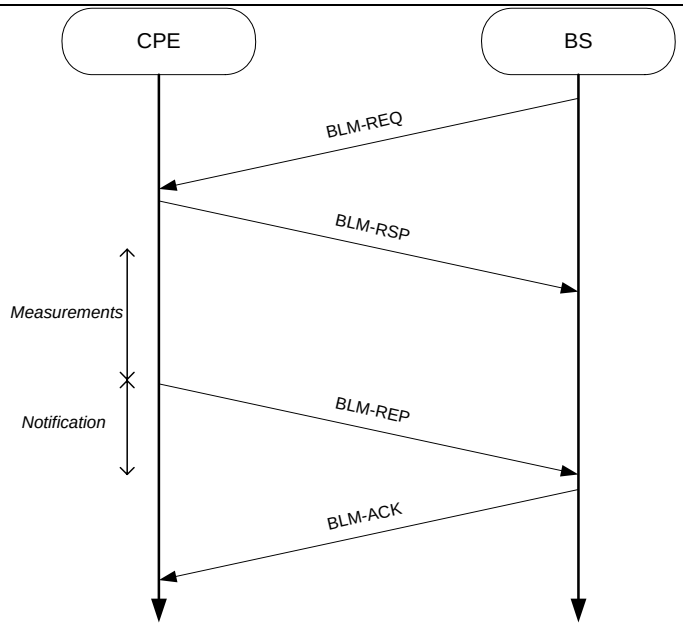

Figure 49 – Measurement message flow between BS and CPE

The nature of the reports received by the BS can be essentially of two types: regular or urgent. Regular reports refer to the cases where the BS has explicitly requested CPEs to report back to it with a certain periodicity (and so the BS can allocate sufficient upstream resources beforehand), and also when CPEs are allowed to report autonomously, for example, whenever enough data has been collected (in this case, CPEs may have to request for upstream resource allocation). Urgent reports are those that take place whenever an incumbent is detected in a channel in use by the 802.22 cell, and need to be reported back to the BS immediately (see 21.1.4). In this case, the BS should provide regular upstream UCS contention periods where CPEs can notify the BS about potential interference and any critical measurement results, or else the CPE can use the UCS and CN fields in the MAC frame header (6.1.1) to notify the BS about an urgent situation.

Once the BS analyses the reports from its various CPEs, it may wish to take steps to resolve any potential coexistence situation (either with incumbents of self-coexistence). To this end, CMAC supports a rich set of channel management messages (see Table 24 and 8.21) that enables the BS to act promptly and effectively as to resolve the coexistence situation. The IDRP protocol is also part of the detection recovery procedure (see 21.1.5). Transmission power control can also be employed for improving coexistence. In case of self-coexistence, other mechanisms available are "interference-free" scheduling and traffic constraints, and their use is discussed in 21.2.

21.1.3 Incumbent Detection

CMAC is able to fully manage incumbent detection for both TV signals and Part 74 services[16]. The DFS model is implemented as per defined in the functional requirement document.

21.1.4 Measurement Report and Notification

Channel occupancy by incumbents changes over time, and this is the reason why CPEs and BSs shall periodically sense the medium as to determine the presence or absence of incumbents. In a situation where an incumbent is operating in-band with the 802.22 cell, certain CPEs (or even the BS itself) will likely detect the incumbent activity through the distributed sensing technique. Whenever this happens, the CPE shall immediately notify and

[16] Details of the incumbent detection algorithms are provided in the corresponding PHY proposal to this WG.

report this situation to the BS. As described in 21.5, in-band incumbent detection can take place during two phases: quiet periods or normal system operation.

Regardless of the detection phase, during this time both BS and CPEs shall execute algorithms that allow the reliable detection of incumbent signals. On the other hand, the way CMAC responds in these two phases is quite different. More specifically, the BS and CPEs shall take different actions depending upon if the incumbent notification to be made was as a result of a quiet period or of detection during normal system operation. This is shown in Figure 50, and is based on the observation that the quiet period notification phase is likely going to be much more demanding in terms of measurement reports than the normal system operation notification phase. Hence, proper measures have to be enforced by the MAC protocol to accommodate for these two different phases of notification as they have different natures. The following subsections describe the behaviour of CMAC during these two phases.

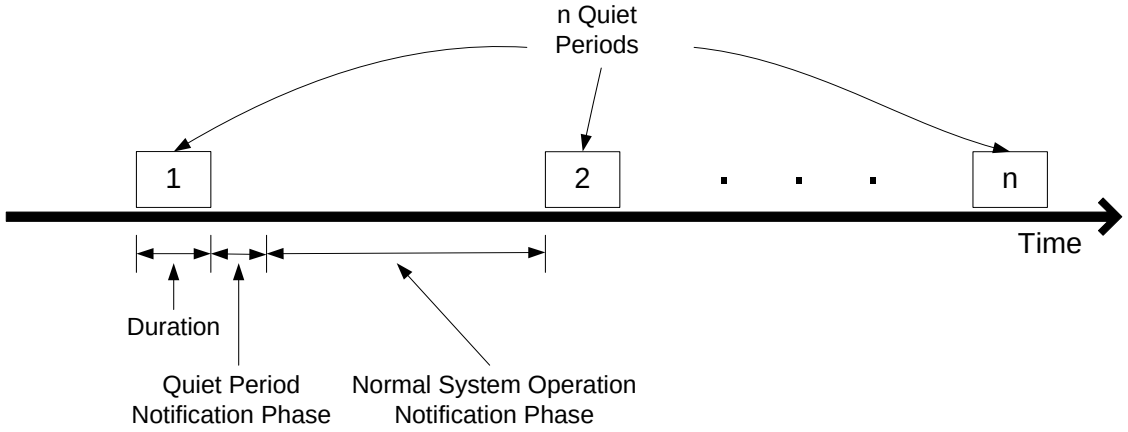

Figure 50 – Incumbent notification phases

21.1.4.1 During Quiet Period Notification Phase

After a quiet period, both the BS and CPEs have performed incumbent measurements. If the BS itself detected the presence of an in-band incumbent, it can proceed in different ways as discussed in 21.1.2. Regardless of that, in the next frame (and optionally in subsequent frames)[17] right after the end of the quiet period the MAC at the BS shall limit its downstream transmissions to the minimum necessary, and devote most of its frame allocation for upstream traffic. Not only that, to guarantee that most CPEs get a chance to reliably contact the BS with a measurement report, the BS shall divide the entire upstream bandwidth allocation into at the most two parts (not necessarily of equal size): dedicated per CPE upstream allocation and UCS notification slots (see 4 and 805512728.1332372.5.37573021.1.4.2.2)[18].

CPEs that are allocated upstream bandwidth shall use it to send to the BS a very brief report on its overall measurement outcome (i.e., incumbent detected or not, and in which channel). The way the BS indicates to the CPE that the allocation is primarily for measurement report is done through the MDP field located in the US-MAP message (see 8.4.1). Upon receiving measurement reports, the BS may decide to provide, in the next frame, for more upstream bandwidth allocation for those CPEs who indicated the presence of incumbents, and hence obtain a more comprehensive report. Therefore, only the minimum necessary upstream bandwidth is used for the first report stage, while in the second report stage more bandwidth can be allocated. We note that a CPE which did not detect any incumbent but which was allocated by the BS upstream bandwidth with the MDP field set (see

[17] The number of frames will highly depend on the number of CPEs.

[18] Note that, if necessary, the BS may guarantee the full reliability of the notification procedure (i.e., ensure that all CPEs are heard) by allocating upstream bandwidth for each and every CPE in its cell. In this case, the BS should not allocate time for the UCS Notification Slots. In some situations, however, this may lead waste of bandwidth especially when no incumbents were detected. There is a clear trade-off between notification reliability and data efficiency.

8.4.1), can choose to use this allocation for sending any other data other than measurement data. This will provide better use of resources and also serve to indicate to the BS that this particular CPE did not detect the presence of any incumbent in the previous measurement period.

In the event that a CPE was allocated upstream bandwidth but did not respond back to the BS, the BS shall assume the worst, that is, that the lack of response from the CPE was due to the fact that the CPE is already under interference from the incumbent and which, as a result, is causing collisions with the signals originated at the BS. In this case, the BS shall take measures to overcome the UCS as discussed in 21.1.2.

Those CPEs who have not been allocated dedicated upstream bandwidth, but who have detected the presence of an incumbent during the quiet period, shall use the UCS notification slots for the purpose of notification (the procedure to be used by CPEs in this case is the same as the one used during normal system operation, and can be found in section 805512728.1332372.5.37746721.1.4.2.2). If no such UCS notification slots are available, the CPE shall wait for subsequent frames where the BS will either allocate upstream bandwidth for this particular CPE or schedule UCS notification slots.

It is important to note that only those CPEs who have not been allocated upstream bandwidth in a frame are allowed to use the UCS notification slots in the same frame. Those CPEs having upstream bandwidth allocation with the BS shall not use the UCS notification slots (see 805512728.1332372.5.37785521.1.4.2.2 for further details).

To improve the reliability and performance of the system, two types of UCS Notification windows are possible (see 805512728.1332372.5.37840221.1.4.2.2). In the case of a contention-based CDMA UCS Notification, the CPE shall transmit the corresponding Incumbent Code and wait for a CDMA_Allocation_IE from the BS. This will allow the CPE the opportunity to report on any UCS. On the other hand, in the case of a contention-based UCS Notification, the BS shall determine the size of a dedicated per CPE upstream bandwidth allocation and UCS notification slot to be big enough to fit only the general MAC header (which is the smallest unit of information for incumbent notification purposes – see 6.1.1). The use of both of these types of notification schemes will provide a quick and reliable report from the CPEs to the BS to be made in the first stage, and allow the BS to dedicate, in a second stage, more upstream bandwidth resources for a full report only to those CPEs who have detected the presence of incumbents.

21.1.4.2 During Normal System Operation Notification Phase

The quiet period notification phase ends whenever the BS has acquired a reliable picture of the measurement outcome in its cell. This may mean, for example, that the BS has concluded that an incumbent has occupied a channel or that sufficient information has been obtained that indicates no incumbent activity. In CMAC, the end of a quiet period notification phase is indicated through the MDP field contained in the US-MAP message (see 8.4.1).

Once the quiet period notification phase is over, the normal system operation notification phase begins. In this phase, the BS shall allocate only the UCS notification slots for the specific purpose of incumbent notification given the lower expected demand during this phase. If the detection of an incumbent by a CPE takes place during this phase (as discussed in 21.5), the CPE can notify the BS in either of two ways depending upon whether or not the CPE has been granted upstream bandwidth allocation in the frame where the notification is to be made.

21.1.4.2.1 CPEs with Upstream Bandwidth Allocation

In case the CPE has sufficient upstream bandwidth allocation to send the BLM-REP back to the BS, it shall do so in the first available opportunity. Thus, BLM-REP messages take precedence over all other messages in what regards incumbent protection. Once the BS receives the BLM-REP message, it proceeds as outlined in 21.1.2 and takes any necessary steps to resolve the coexistence situation.

On the other hand, if not enough upstream bandwidth allocation is available to the CPE to fit a BLM-REP message, or not enough time is available for the creation of a report packet before the next upstream bandwidth allocation, or the CPE needs to transmit a small amount of sensitive traffic (e.g., voice) in order to meet QoS guarantees, an alternate method exists wherein the CPE sets the UCS and Channel Number fields in the general MAC header (see 6.1.1). By setting these fields, it will indicate to the BS about the urgent coexistence situation. Here, once the BS is notified through the UCS and Channel Number fields in the MAC header, it shall proceed in one of the following ways.

First, once it receives the first UCS notification it may allocate more upstream resources in the next frame to the reporting CPE so that this CPE can send the full report via a BLM-REP. Alternatively, the BS may send a BLM-REQ message to the CPE in question specifically requesting from this CPE more information that shall be sent in the next upstream allocation of this CPE. Another possibility is for the BS to play safe and immediately issue channel management messages in order to resolve the situation. Finally, one last option could be for the BS to delay taking any immediate action and wait for indications from other CPEs. In case no other CPEs report the same UCS with incumbents, the BS may conclude that a measurement report by a single CPE is not reliable and may disregard it. On the other hand, if multiple CPEs report the same coexistence situation in the same Channel Number, then the BS shall take one of the measures discussed above in order to resolve the issue.

21.1.4.2.2 CPEs without Upstream Bandwidth Allocation

Even if the CPE does not have any upstream bandwidth allocation with its BS, it still needs to report to the BS about the UCS with incumbents. In this case, the CPE shall use the upstream UCS notification slots (see 4) in order to reach the BS and indicate the UCS with incumbents. The UCS notification slots shall always be allocated by the BS in the same time/frequency region across frames. This will allow even those CPEs who have suddenly started to experience harmful interference[19] from an incumbent service to reliably notify the BS about the presence of the incumbent service. The CPE shall notify the BS in the UCS notification slot immediately after the incumbent service is detected, while it is still synchronized to the BS.

It is important to highlight that the only situation when CPEs are allowed to use the UCS notification slots is when they do not have upstream bandwidth allocation but yet need to reach the BS and report a UCS with incumbents. Out of these CPEs, only those CPEs who detected the presence of an incumbent are entitled to use the UCS notification slots. In other words, CPEs who have not been allocated upstream bandwidth in a given frame and who also did not detect an incumbent, shall not use the UCS notification slots but rather shall wait for the BS to take any action in this respect in future frames. If upstream bandwidth allocation is available, the CPE shall proceed as indicated in 805512728.1332372.5.38311821.1.4.2.1.

There are two possible ways a BS can allocate UCS notification windows, and hence the CPE shall use these slots in accordance to its type. They are: Contention-based UCS Notification and Contention-based CDMA UCS Notification.

Contention-based UCS Notification

In reporting a UCS through the use of the contention-based UCS notification slots, the CPE shall transmit only the general MAC header, typically without any payload. In this MAC header, the CPE shall set the UCS and Channel Number field accordingly so as to allow the BS to be notified of the UCS. Upon receiving the message from the CPE, the BS shall proceed as discussed in 21.1.2. If allowed by the BS and the CPE chooses to send a payload, this will consist of at most a single measurement report which indicates the type of incumbent signal detected and the power level (see 8.21.3.1.1).

Contention-based CDMA UCS Notification

[19] In this context, harmful interference is when the CPE is no longer able to decode packets coming from its BS.

In addition to the UCS notification window described above, the PHY also supports the use a CDMA mechanism for the purpose of UCS notification.

As specified in the PHY spec, the PHY has available a subset of Incumbent Codes that shall be used for contention-based CDMA UCS Notification. The CPE, upon needing to make a UCS Notification, shall select, with equal probability, an Incumbent Code from the code subset allocated to Incumbent Notification. This Incumbent Code shall be modulated onto an Incumbent Subchannel and transmitted during the appropriate UCS Notification window. The Incumbent Subchannel shall be selected by the PHY amongst the ones reserved by the MAC for UCS Notification.

Upon detection, the BS shall provide (an implementation dependent) upstream allocation for the CPE, but instead of indicating a Basic CID, the broadcast CID shall be sent in combination with a CDMA_Allocation_IE, which specifies the transmit region and Code that were used by the CPE. This allows a CPE to determine whether it has been given an allocation by matching these parameters with the parameters it used. The CPE shall use the allocation to transmit a MAC PDU with the UCS and CN fields in the MAC header properly set. In addition, if allowed by the BS, data could also be transmitted in this allocation (this is indicated by the Usage field – see Table 48).

If the BS does not issue the CDMA_Allocation_IE described above, the CPE shall assume that the Code transmission resulted in a collision and follow the contention resolution as specified in 14.

21.1.5 Incumbent Detection Recovery Protocol

The previous subsections discuss incumbent measurement control, how these measurements are reported to the BS, and how incumbents are detected. To recover from an UCS, protocols are needed that enable the network to efficiently and dynamically restore to its normal operation with minimal performance degradation. In other words, mechanisms are necessary that allow quick, reliable, and efficient recovery from an incumbent detection situation. This way, not only effective protection of primary systems can be guaranteed, but it would also potentially allow a minimum QoS level to be guaranteed and, most importantly, no apparent discontinuity of service despite the changing availability of radio spectrum resources. Therefore, this section addresses the incumbent detection recovery protocol (IDRP) that is employed in CMAC and which allows a quick service restoration under a UCS.

Clearly, depending upon many factors different actions are taken at the CPE and BS. Figure 51 and Figure 52 show in detail the procedures executed in IDRP in the cases of the BS and CPE, respectively. As we can see, the BS carries out a more simplified procedure than the CPE, since the latter is in charge of not only reporting any UCS, but also recovering from such states in a timely manner. In the case of a CPE, the first step in IDRP upon detection of a primary radio service is to notify the BS and await further instructions (this can be done a few times up to a pre-determined amount of time, as depicted in Figure 52), while for the BS it can immediately send spectrum management commands.

One of the key concepts of IDRP is the use of *backup channels*, which allow CMAC to quickly re-establish communication in the event of primary detection. Since CMAC provides for backup channel(s) (i.e., as discussed earlier, channel(s) which are kept in both the BS and CPEs and which can be resorted upon detection of an incumbent service), these are used by a CPE when looking for the BS. This way, the recovery procedure can be made significantly faster, as both the BS and CPE would know in advance where to restore the service should an incumbent initiate operation in their channel. To maximize the utility of the backup channel(s), this should be selected in a way to be completely disjoint from the current operational channel(s). This way, the likelihood that the backup channel is also affected when the incumbent service initiates operation in the operational channel can be significantly minimized. Figure 53 depicts the procedure by which CPEs keep track of backup channels in response to out-of-band measurements previously requested by the BS. Every CPE shall perform this procedure according to the DFS model defined in [3], which in its present version specifies a repetition period of execution equal to 6 seconds (i.e., 20% of the Channel Availability Check Time).

Finally, IDRP also incorporates a mechanism to overcome the situation that may occur and which leads to an erroneous perception from the BS that incumbents occupy all channels, which causes the interruption of all transmissions in a cell. In this case, a recovery method is necessary and here we propose to use the CPE to assist the BS in notifying about a new vacant channel. This procedure is shown in both Figure 51 and Figure 52, and is especially useful when the incumbent service has lower duty cycles (e.g., Part 74 services) and in mobile scenarios (although this is currently out of 802.22 scope). Typically, this would happen when only CPEs detect the presence of incumbent services for most of the channels. In these cases, the BS would have to rely upon CPEs to inform the status of these channels at a later point in time, otherwise the BS may be unable to reuse these channels indefinitely. In addition, a situation could arise wherein no vacant channels are left and hence the secondary system cannot operate. To overcome this (as shown in Figure 52), the CPE would periodically re-evaluate the status of a channel it has reported as occupied by incumbents. If this channel becomes again free in the future, the CPE would have the capability of sending a notification to the BS which would, in turn, periodically listen in this channel for any such incoming notification (as shown in Figure 51).

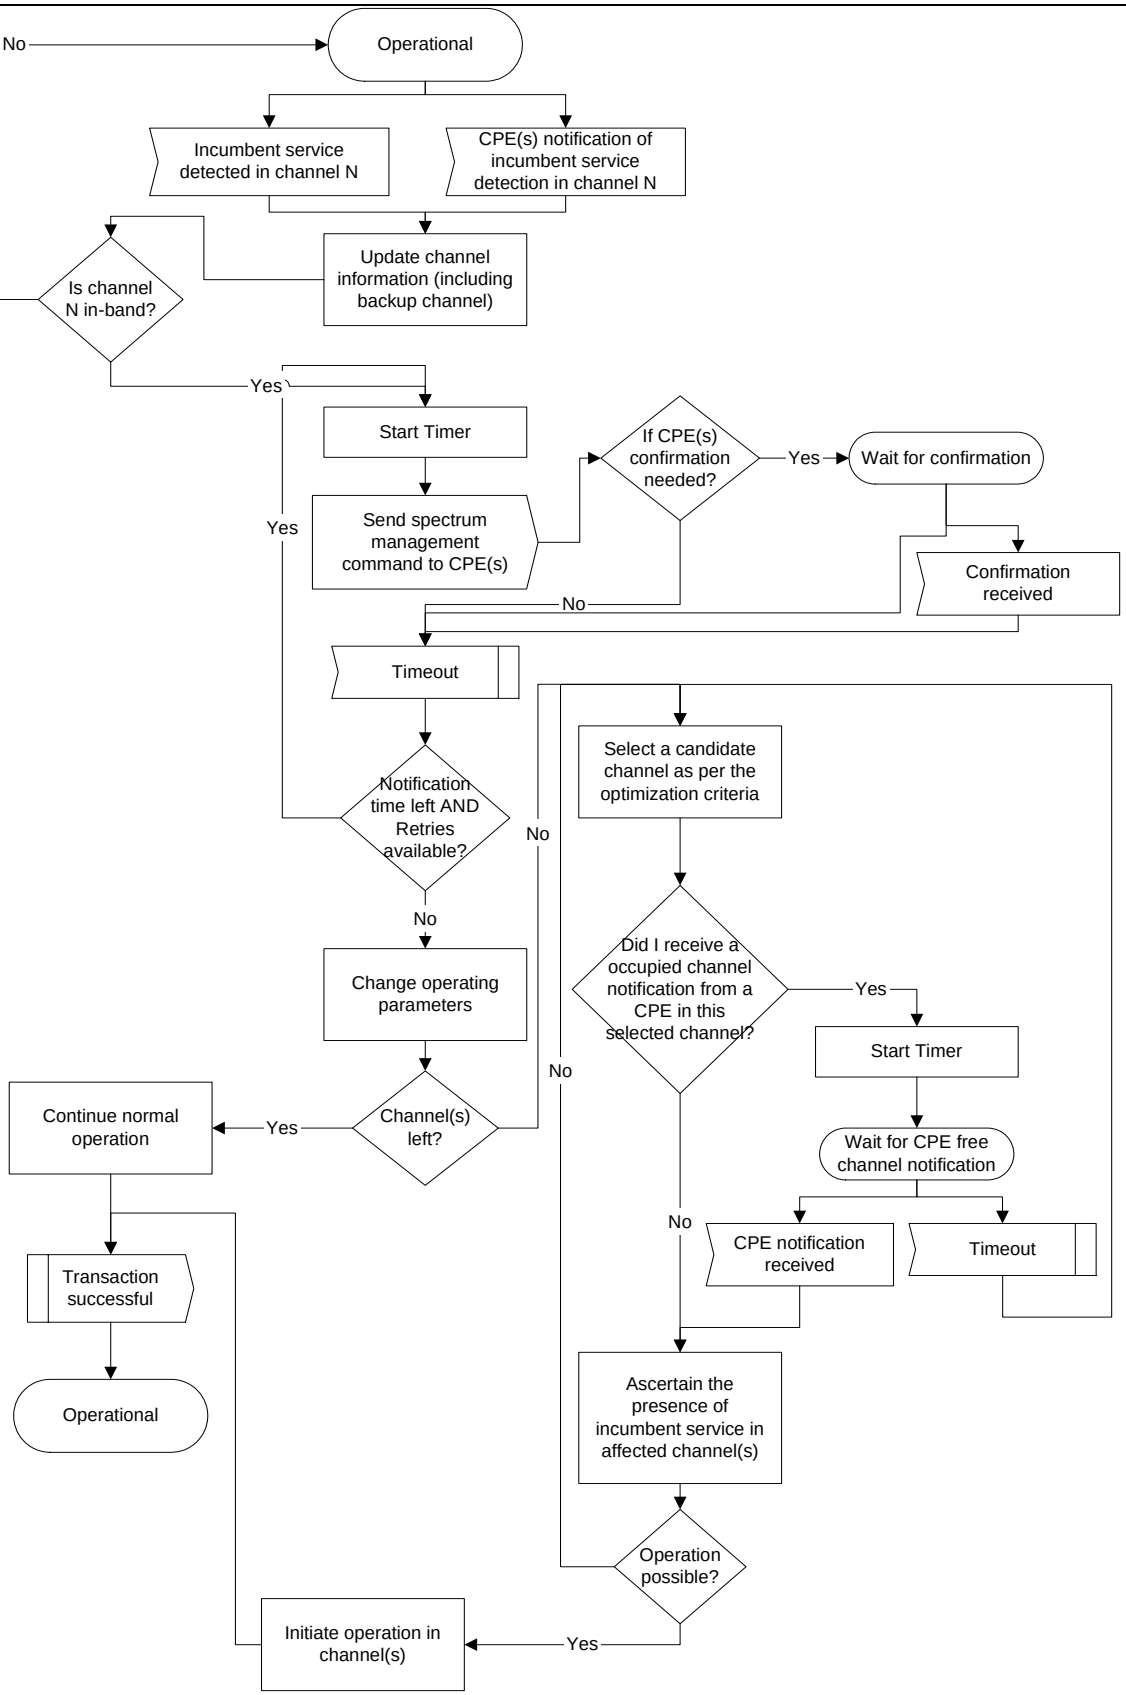

Figure 51 – IDRP at BS

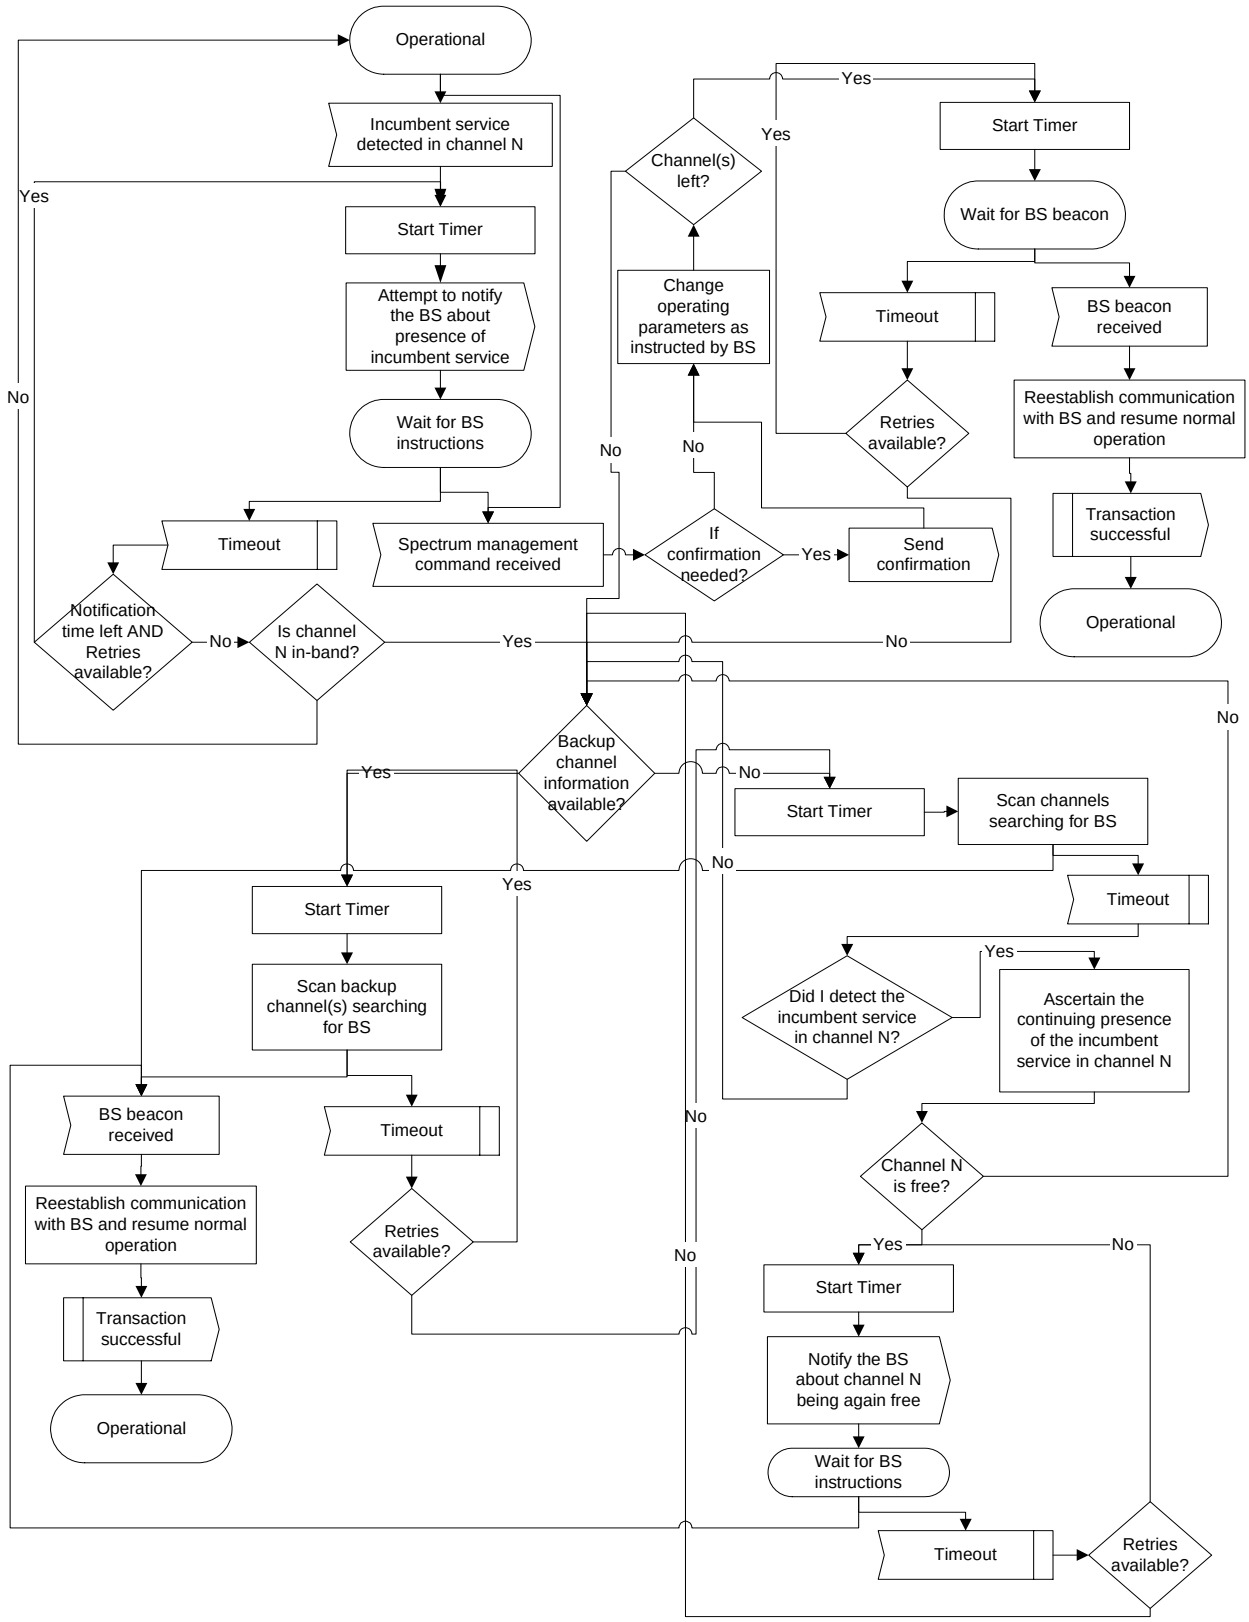

Figure 52 – IDRP at CPE

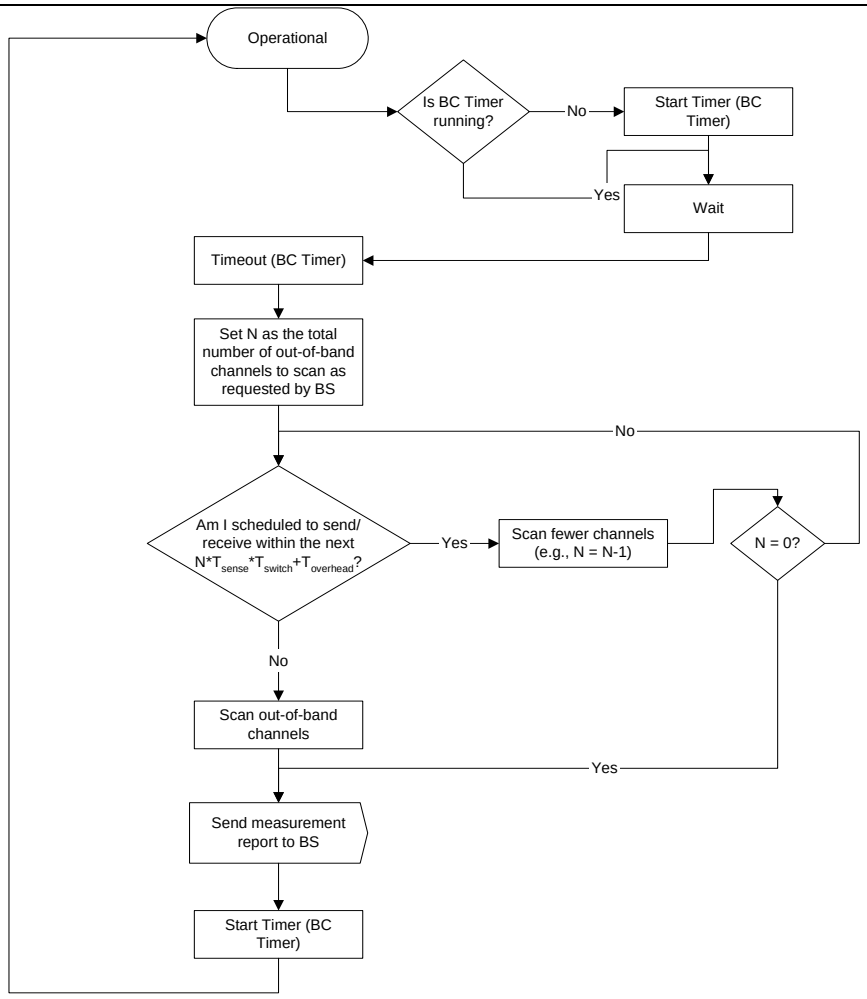

Figure 53 – Procedure to keep track of backup channels (and perform out-of-band measurements)

21.1.5.1 Incumbent Detection Recovery Protocol (to be merged with previous IDRP)

Herein, "Incumbent Detection Recovery" refers to the operation of selecting a new channel in the event of appearance of incumbent user (IU). There are three types of notification: Explicit, Implicit and Short implicit. These notification methods, a part of the recovery protocol, have different merits and demerits for incumbent appearance scenarios. First of all, explicit notification is the best way to notify to CPE for channel change (recovery). However, when CPE cannot hear the downstream, it is impossible to receive channel switch request (CHS-REQ) message. In such case, we can use implicit notification. But implicit recovery needs more time overhead comparing with explicit notification. Therefore, these notification methods can be specified by incumbent appearance specification. Figure 54 shows five possible the incumbent appearance scenarios:

- Case0: When IU detected by both BS and CPE
- Case1: When IU in upstream detected by BS
- Case2: When IU in downstream detected by BS
- Case3: When IU in upstream detected by CPE
- Case4: When IU in downstream detected by CPE.

In case 0, BS and CPE move their channel to predefined candidate channel immediately. Because both BS and CPE notice IU, they not need more time to notice each other. In other cases, we can choose one or more methods depending on IU appearance scenarios.

- **CASE 0: IU detected by both BS and CPE**

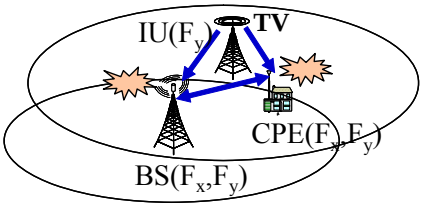

➔ Prompt recovery

CPE(F_x,F_y) ➔ CPE(F_x,F_z]
BS(F_x,F_y) ➔ BS(F_x,F_z)

- **CASE 1-4: What if IU can be detected by either BS or CPE only?**

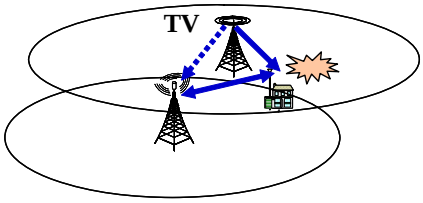

4 different cases		Interference detected in	
		IU = U/L	IU = D/L
IU detected by	BS	Case (i)	Case (ii)
	CPE	Case (iii)	Case (iv)

Figure 54 – Incumbent appearance scenarios

	IU appear in	Detected IU by	Implicit	Explicit	Short implicit
Case 1	UL	BS	o	o	o
Case 2	DL	BS	o	x	x
Case 3	UL	CPE	o	o	o
Case 4	DL	CPE	o	x	o

Figure 55 – Incumbent notification for each case

In FDD case, possible notification methods are summarized as shown in Figure 55. However, even if an incumbent user appears in the channel, BS can try to send the CHS-REQ message based on the response time [3]. In this reason, when an IU appears in active set, BS firstly tries explicit notification methods. However, FDD system has a critical problem to execute explicit notification. The problem is that the BS cannot make CHS-REQ message. The reason is that if BS fails receiving of BLM-REP, BS cannot sure what channel has a problem (DS, US or both?). Therefore the BS needs more information to decide the problem of DS and US channel.

To solve this problem we propose double checking response protocol (DCRP). Basically, in FDD mode, BLM-REP message are used not only acknowledgement of BLM-REQ but also transmission of sensing result itself. Figure 56 and Figure 57 show sensing result report message and corresponding CPE procedure.

	Response with	Convey information	Position
BLM_RSP	Automatically	Minimum notification and heavy coding	Self coexistence slot or UCS slot
BLM_REP	CH_SCAN_REQ_IE	Sensing result (in band, outs band)	Traffic slot

Figure 56 – Sensing result report messages

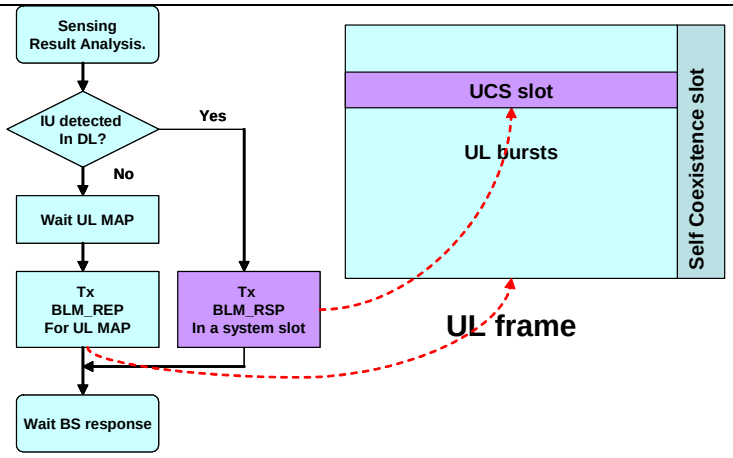

Figure 57 – Sensing result report procedure

Based on DCRP, the BS concludes what channel has a problem as shown in Figure 58. Finally, we can devise the corresponding recovery procedure depicted in Figure 59 and Figure 60 for FDD mode, and in Figure 61 and Figure 62 for TDD mode.

	Message receiving status		
CH_SCAN_RSP	O	X	X
CH_SCAN_REP	D.C	O	X
Channel Status(DL)	O	X	Unknown
Channel Status(UL)	O	O	X
BS action	Communication Continuous	CHG. DL	Send CHS_REQ msg For UL CHS After TL DL CHS

Figure 58 – BS sensing result report analysis

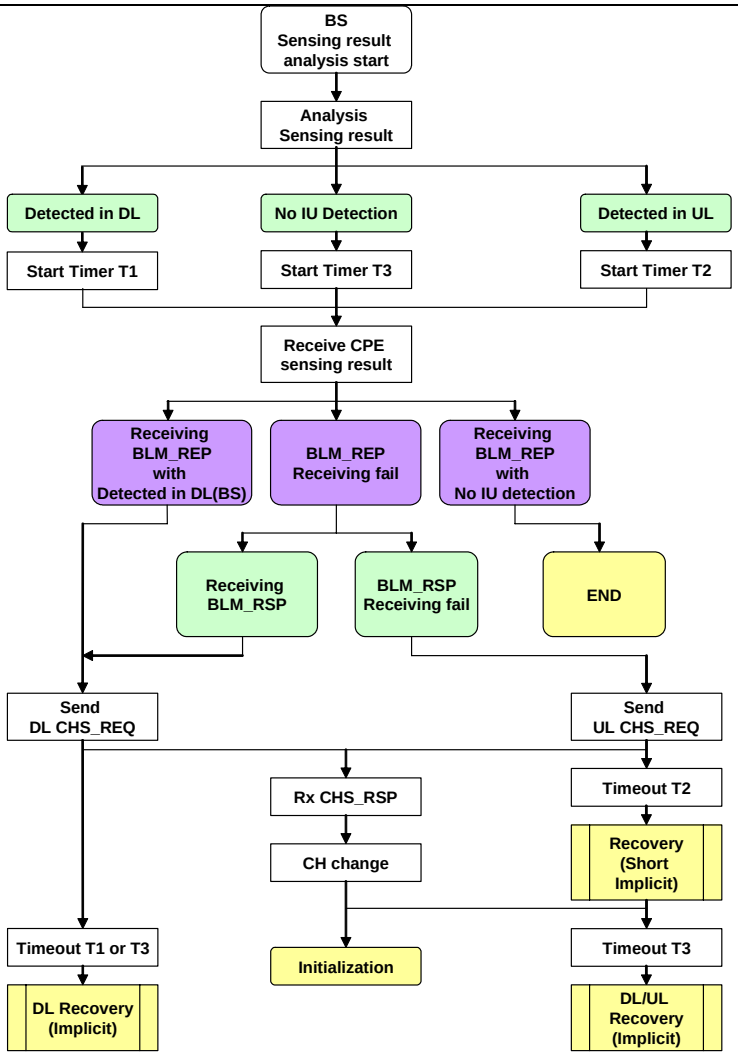

Figure 59 – BS recovery procedure in FDD mode

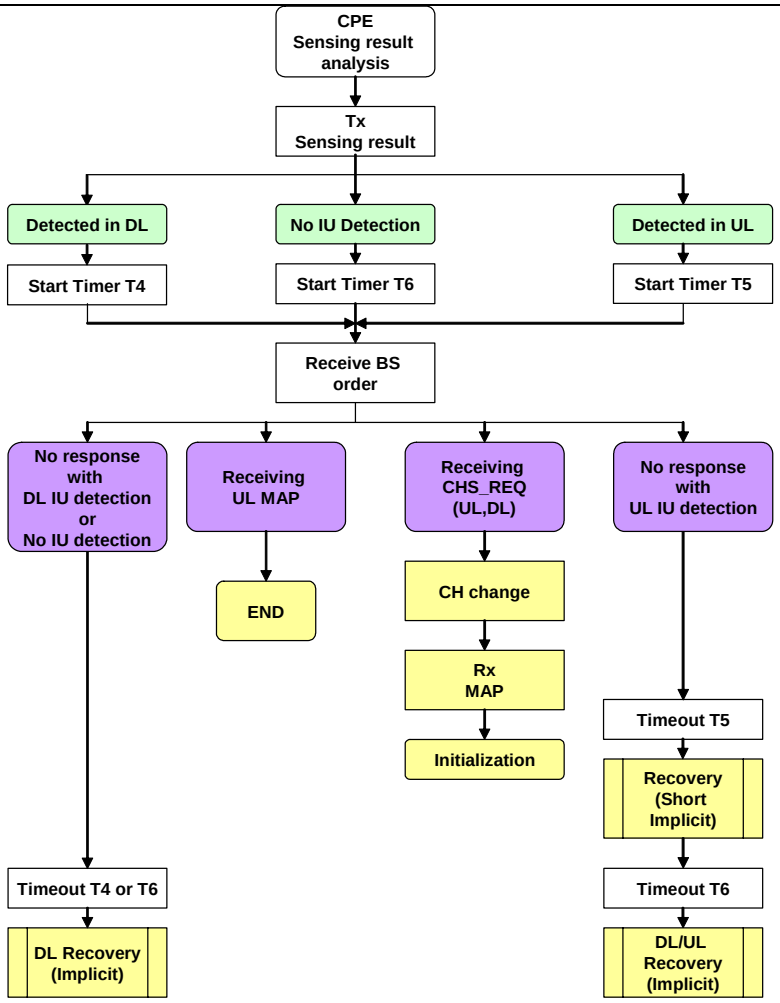

Figure 60 – CPE recovery procedure in FDD mode

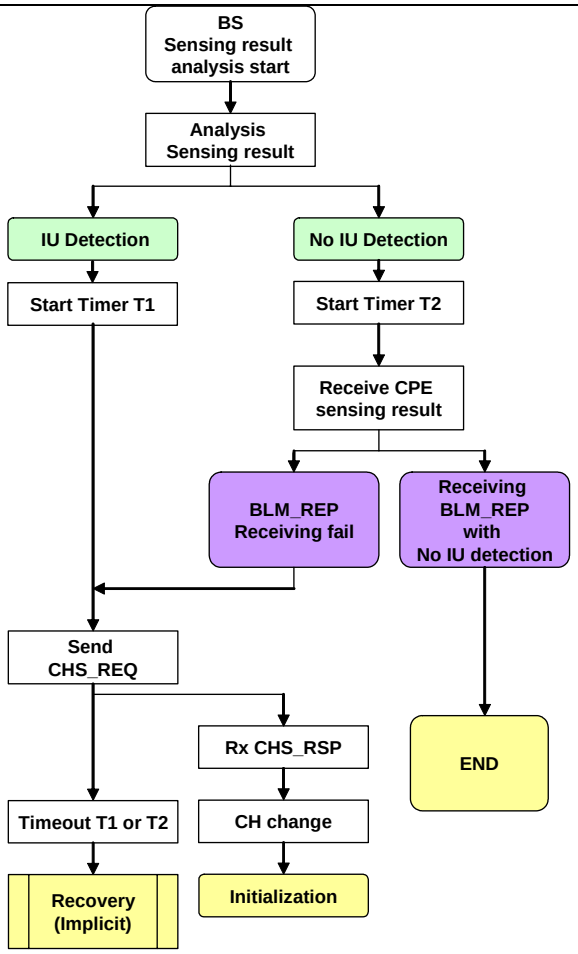

Figure 61 – BS recovery procedure in TDD mode

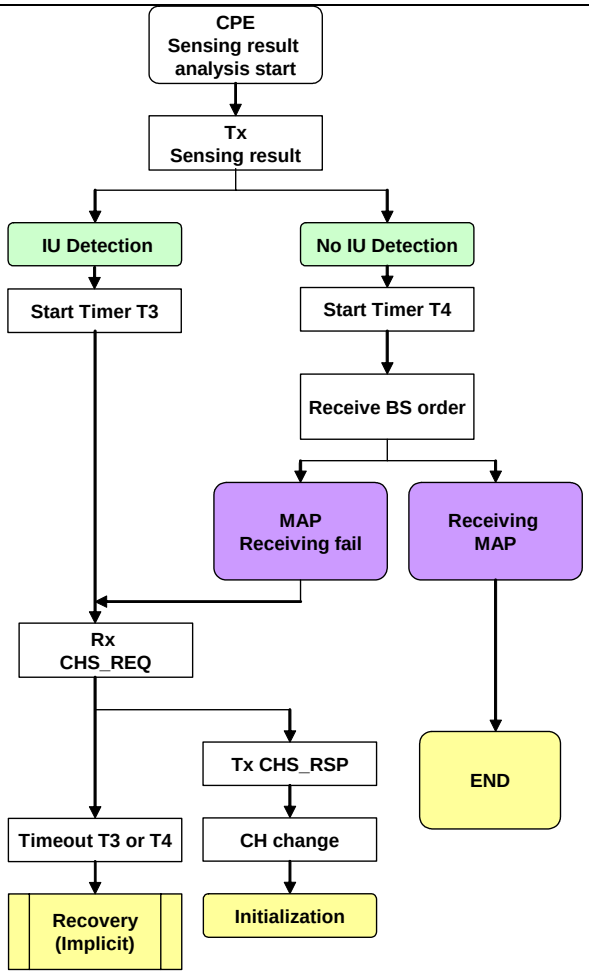

Figure 62 – CPE recovery procedure in TDD mode

21.1.6 DFS for Incumbent Protection

The 802.22 standard shall protect incumbent services operating in the TV broadcast bands. To this end, CMAC has been designed in such a way that the DFS model specified in [3] can be easily supported. As described earlier, this is done through measurement and spectrum management messages supported in CMAC (for example, see 8.21 and 8.22), which allow for the BS to control and implement the necessary behaviour to sense and protect the incumbent services.

21.1.7 Class B CPE for the Protection of Part 74 Services

The current 802.22 Study Group is considering some options for the protection of Part 74 devices (read, Wireless Microphones). As of this date, two major classes have been identified[20]. In the first class (class A), a separate beacon device (e.g., Zigbee like transmitter) is envisioned who will transmit short wireless microphone beacon (WMB) messages to notify collocated 802.22 system about the presence of a co-channel wireless microphone operation. In the second class (class B), the 802.22 system could support a special type of CPE (thus, an integral

[20] As discussed in the Study Group, each of these approaches has pros and cons (both technical and regulatory). Here, however, we focus only on the technical aspect.

part of the 802.22 standard) incorporating specific capabilities to inform collocated 802.22 systems about wireless microphone operation.

With respect to the final 802.22 standard, the class B approach is the one that requires the transmitter to be defined and included in the standard, whereas this is not the case for class A. For class A solutions, the 802.22 standard would only be required to understand a signal transmitted by an "alien" WMB device. We believe, however, that a single approach is not the best solution. Rather, we advocate that not only a separate beacon device be developed, but that the 802.22 standard also include built-in mechanism to allow the reliable protection of wireless microphone services.

Based on this observation, here we propose a solution to the class B approach. It is builds on the superframe foundation provided by CMAC, and can operate efficiently even under multiple overlapping 802.22 cells. In the following subsections we describe such class B proposed solution.

21.1.7.1 Scanning

Initially, the class B CPE (say, with special firmware and licensed under Part 74) shall be tuned to the desired frequency channel where the wireless microphone will operate. Upon initialization, this CPE shall scan the desired channel for a multiple number of the maximum superframe size in search for SCH packets transmitted by 802.22 BSs. In case no SCH is received for such duration of time, the CPE shall assume that no 802.22 system is operating in the channel. Despite of that, it shall continuously scan the desired channel in search for SCH transmissions as a BS may erroneously attempt to initiate transmission in this channel at any time.

If, on the other hand, one or more SCH packets are received, the CPE shall use the quiet period information contained in the SCH packet and build a *quiet period map* for all these on-channel BSs. An example of a quiet period map for an arbitrary TV channel is shown in Figure 63. As can be seen from this figure, it describes, for each BS operating in the considered TV channel, its schedule in terms of data transmission and quiet periods as per transmitted in the corresponding SCH packets. This map is then used by the class B CPE to determine the time periods which stand higher probability that a class B WMB transmission is received by the intended BS. These are, obviously, the quiet periods of each BS.

It may be possible that the class B CPE does not detect all 802.22 BS in operation in its vicinity (e.g., due to interference amongst multiple collocated 802.22 BSs), but as it is described in the next subsection, the notification scheme operates in stages with increasing reliability, which addresses this shortcoming.

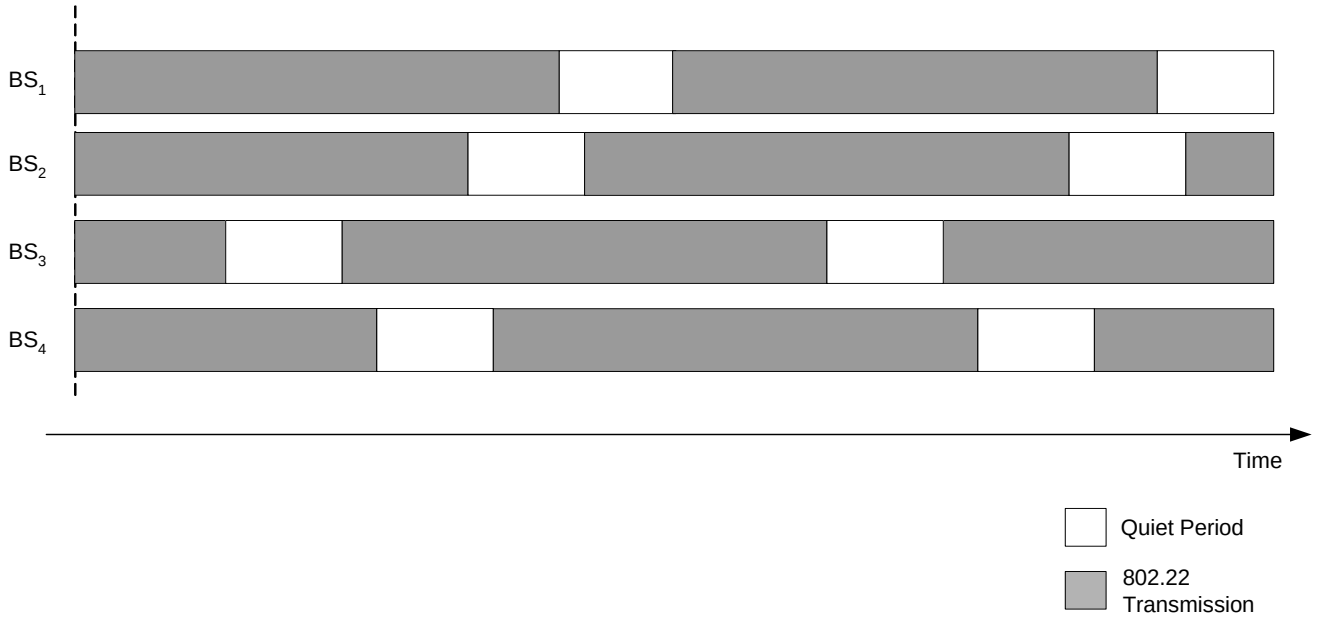

Figure 63 – Quiet period map for an arbitrary TV channel

21.1.7.2 Class B CPE Notification

Once the quiet period map is constructed, the class B CPE shall transmit WMB to notify all previously identified BSs about the scheduled wireless microphone operation. For this purpose, there is no need for a class B CPE to join a BS before sending WMBs. These WMB shall be transmitted during the quiet periods of the BSs (obtained from the quiet period map) to increase the probability of success. Furthermore, a class B CPE may transmit each WMB up to N_{Beacon} times to the same BS in order to further increase the transmission success probability.

The frame structure of a WMB is similar to that of BS beacons, with the key distinction that the System Type field in the SCH packet shall now be set to indicate that it is from a wireless microphone. Table 225 describes the structure of a WMB. In addition to the information about the scheduled use of the channel by the wireless microphone operator, the WMB shall also contain information that allow the 802.22 BS to authenticate the WMB, and hence prevent any misuse of this functionality. It is also possible that these devices possess pre-programmed security keys (e.g., that could be, in a secure way, established beforehand amongst wireless microphone operators and 802.22 service providers) that will allow enhanced security.

Table 225 – Structure of a wireless microphone beacon

Syntax	Size	Notes
Wireless_Microphone_Beacon_Format() {		Transmitted with well-known modulation/coding (e.g., QPSK rate ½)
ST	7 bits	System Type Indicates the type of the system using this band. According to Table 1, this field is to be set to 1
Tx ID	48 bits	Address that uniquely identifies the transmitter (in this case, the class B CPE)
Transaction ID	8 bits	The ID that uniquely identifies the transaction
Start Time	16 bits	Start time (in units of minutes) of the wireless microphone operation.
Duration	16 bits	Time (in units of minutes) during which the channel is going to be used by the wireless microphone service
Security Key	128 bits	
Reserved	1 bit	
Length	8 bits	
IEs	*Variable*	Information Elements • Common MAC IE (see Section 7) such as Transmit Power. • Location Information IE (see Table 169)
HCS	8 bits	Header Check Sequence See Table 6
}		

21.1.7.3 BS Acknowledgement and Dissemination

The BS shall acknowledge the reception of a WMB by including a Part 74 acknowledgement IE in the SCH (see 7.7). This is used by the BS to signal the class B CPE that the WMB was successfully received and that the BS will comply with the wireless microphone request. In case the class B CPE does not receive such acknowledgement, it shall repeat the notification procedure. This shall be done indefinitely until all BSs acknowledge receipt of the WMB.

Besides the BS, the reception of a WMB can also be done through any of the ordinary network CPEs. In this case, the CPE shall wait until the end of the quiet period of its cell and listen for the following SCH frame transmitted

by its BS. In case a Part 74 acknowledgement is not included in the SCH, the CPE shall assume that the BS did not receive the WMB and shall then take the responsibility to report to its BS about the presence of the wireless microphone service (see 8.22.3).

Once all previously identified BSs indicate the successful reception of the WMBs, the class B shall go back to the scanning phase and continue on the look out for BS beacons. As 802.22 BS move to other TV channels, the interference level in the desired wireless microphone channel decreases, and so increase the chances that other 802.22 BS which were not detected in the previous scanning phase now become detectable. If this procedure is repeated a few times, there is a high probability that the desired wireless microphone channel becomes free of any 802.22 cell operation.

Once BSs and CPEs switch to some other vacant channel as to free up the wireless microphone channel, the dissemination procedure shall start. In this phase, the BS shall continue to append the Part 74 acknowledgement IE to all of its SCH transmissions. This IE shall be appended for the duration of time that the wireless microphone is in operation in the channel where the WMB was transmitted. This is done to notify any other new collocated BSs, which do not know about the ongoing wireless microphone operation in the WMB channel, that they should not even try to transmit in the WMB channel. Therefore, a more effective protection of wireless microphone services can be accomplished.

21.2 Self-Coexistence

Contrary to other IEEE 802 standards where self-coexistence issues are only considered after the specification essentially is finalized, the IEEE 802.22 takes the proactive approach (as specified in its Requirements Document) and mandates that the MAC shall include self-coexistence protocols and algorithms as part of the initial standard conception and definition. As depicted in Figure 65, multiple 802.22 BSs and CPEs may operate in the same vicinity and provided appropriate measures are taken at the air interface level, self-interference may render the 802.22 system useless. Even if directional antennas are used at the CPEs (although this may be implementation dependent), self-coexistence issues are not at all overcome (see Figure 66). This is further aggravated by the fact that 802.22 coverage range can go up to 100 Km, and hence its interference range and impact on other collocated 802.22 cells is larger than in any other existing unlicensed technology.

CMAC addresses self-coexistence in two ways: through CBP and through inter-BS communication. Inter-BS communication is a passive method in the sense that it cannot be deliberately initiated, but depends on the periodic SCH packets transmitted by the BSs in the beginning of a superframe. CBP, however, behaves in both passive and active modes. In the following subsections, we discuss these two mechanisms which allow for appropriate self-coexistence amongst collocated 802.22 cells.

21.2.1 The Coexistence Beacon Protocol (CBP)

To cope up with the serious self-interference issues that may arise in a real deployment scenario, the CBP is proposed. The CBP is a best-effort protocol based on coexistence beacon transmissions. Since it follows the best-effort model, successful reception of coexistence beacons is not guaranteed (this behaviour is intentional). However, the mechanism for synchronization of overlapping BSs (see 21.3) and also the fact that CPEs do not continuously stay locked to a BS (see below), successful delivery of coexistence beacon transmissions has high probability.

21.2.1.1 CBP Packet Transmission

The structure and transmission of a CBP packet is shown Figure 64. It starts with a preamble which shall be common across all 802.22 networks (see PHY spec), and which is different from the superframe preamble. After the preamble, follows the SCH transmission. Within the SCH, the CT field (see 5.1) shall be set to 1 to indicate

that this is a CBP packet transmission, and hence that a MAC PDU with the beacon MAC header follows the SCH transmission. The use of the CT field provides an additional level of robustness to the preamble, so that receivers can exactly identify the type of content transmitted. By transmitting both the SCH (which contains information about the 802.22 cell) and the CBP MAC PDU (which contains information about the CPE reservations with its BS), the transmitting CPE conveys all necessary information to allow for better self-coexistence.

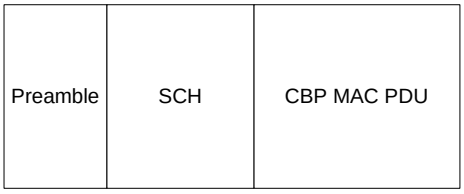

Figure 64 – The structure of a CBP packet

21.2.1.2 Description

In CBP, 802.22 entities (i.e., CPEs and BSs) are capable of transmitting beacons (see 6.1.2) which provide its recipients enough information to achieve satisfactory and good coexistence amongst overlapping 802.22 cells. These beacons are intended for inter-cell communication and carry specific information about a CPE's cell of attachment and downstream/upstream bandwidth allocations with the BS.

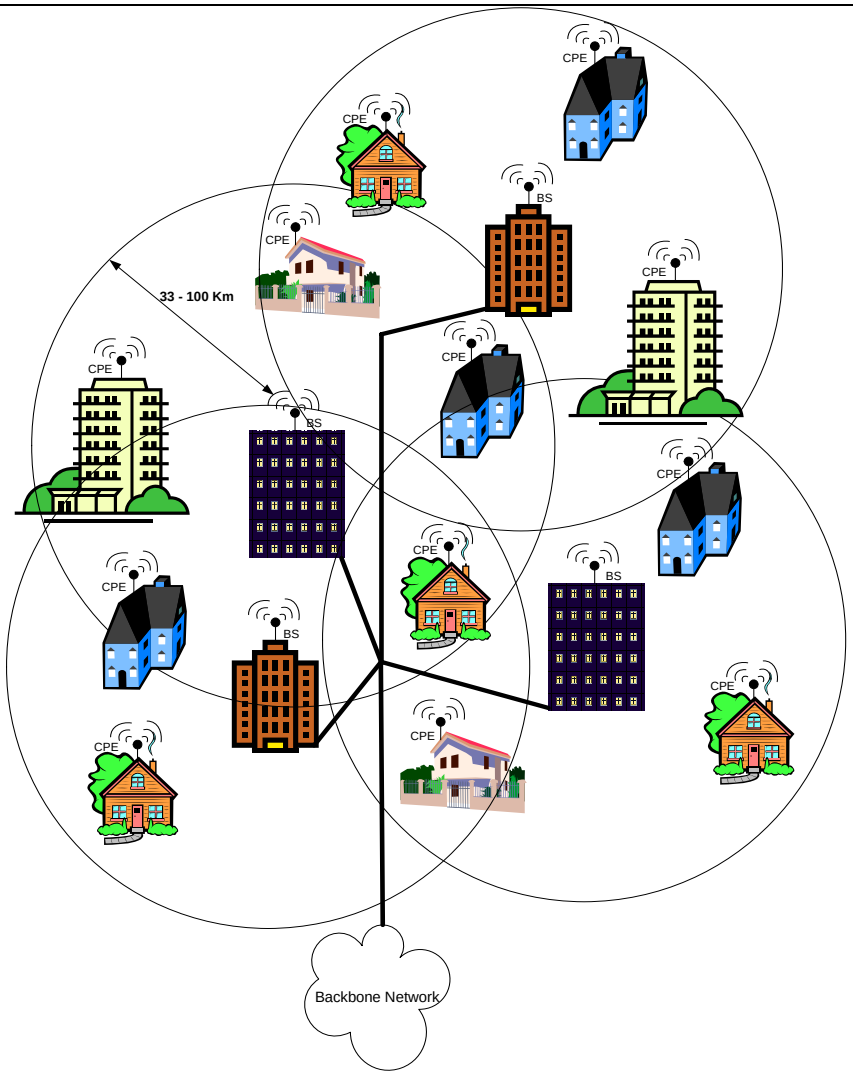

Figure 65 – Example of 802.22 deployment configuration

In CMAC, coexistence beacons are scheduled through the use of Coexistence IUC (both Passive and Active) which can be specified in US-MAP and DS-MAP messages. When scheduling a coexistence beacon, the connection ID contained in the MAP IE indicates which CPEs shall send the beacon within the specified scheduled time. This connection ID can be either unicast (e.g., a CPE's primary connection ID), multicast (i.e., a multicast management connection ID), or even the broadcast ID. In case of multicast, the BS can implement clustering algorithms that improve spectrum utilization and maximize the effectiveness of the coexistence beacons, as multiple CPEs would transmit a coexistence beacon during the same scheduled time (discussed in 21.4). Irrespective of the type of connection ID, the CPE shall always verify if the connection ID specified by the BS includes itself or not. This will determine the CPE's behaviour during the scheduled Coexistence IUC.

Another important remark to be made about the Coexistence IUC is that it defines a period of time where channel access is contention based. In other words, during this time CPEs shall use the contention access mechanism (see 14) to gain access to the medium and transmit the coexistence beacon. The reason why a contention-based access mechanism is preferred for sending coexistence beacons is that it maximizes the spectrum usage. In the majority of cases, it is anticipated that the BS will not schedule a single CPE to transmit beacons, but rather will attempt to improve coexistence by scheduling multiple CPEs to send beacons within the same time span (multicast management connections can be used for this purpose – see 8.10 and 19). Furthermore, when combined with the

clustering algorithms (see 21.4), the efficiency of this contention based is maximized because, for the same Coexistence IUC, the BS will only schedule CPEs not belonging to the same physical cluster (see 21.4).

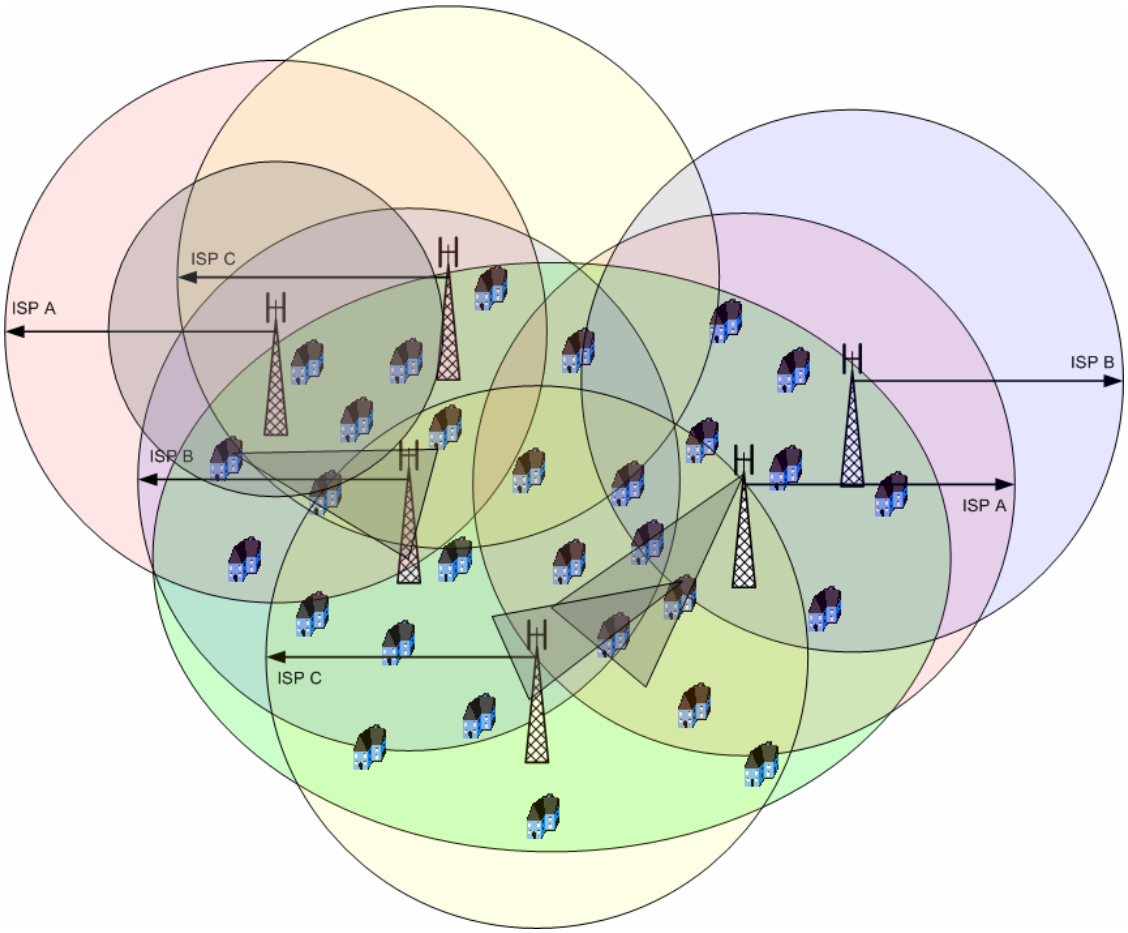

Figure 66 – The use of directional antennas at CPEs does not address self-coexistence issues

In order to maximize the probability that coexistence beacons are received from other collocated 802.22 cells, a CPE does not stay locked to the BS at all times during a frame. A CPE shall only be locked to the BS whenever it is scheduled to receive/send data from/to the BS (indicated through the US-MAP and DS-MAP messages). At all other times during the frame, the CPE shall be listening to the medium and searching for a coexistence beacon. Thus, the success probability and efficiency of the CBP is drastically increased. In case a CPE loses synchronization with BS while listening/receiving a coexistence beacon, it shall regain synchronization in the beginning following frame, which will cause few, if any, side effect.

Another mechanism that can be used by the BS to look out for CBP beacons is to schedule the Coexistence UIC in Passive Mode. Essentially, the Passive Mode defines a time where the CPEs shall not perform any transmission but simply listen to the medium on the look out for CBP packets and BS SCH beacons (since the same preamble is employed for both).

It is important to note that to increase the effectiveness of CBP, downstream/upstream bandwidth allocations made by BS to CPEs in a certain frame shall not change for a number of consecutive frames. This guarantees that the information carried in coexistence beacons is valid for at least a minimum duration of time, thus allowing enough time to the recipients of the coexistence beacon to implement self-interference mitigation mechanisms (discussed below). Further, at any time when a BS must allocate bandwidth to a CPE, it shall always seek to allocate this bandwidth based on the previous allocations (if any) to this CPE. That is, the BS shall always allocate bandwidth to a CPE using approximately the same combination of slot and logical channel. By doing so, this will

reduce the number of coexistence beacons that need to be transmitted by this CPE, since its neighbours would already have the information regarding the allocations as these have not been changed by the BS. Other optimisations are also possible to improve the efficiency of the CBP, and make the transmission of coexistence beacons less frequent.

Once a CPE receives coexistence beacons from other collocated CPEs belonging to different cells, it can use this information in many ways in order to improve coexistence. The first thing a CPE may want to do is to convey to the BS the received information. The BS, in turn, will implement so-called "interference-free" scheduling algorithm, which schedules the various upstream/downstream traffic from/to CPEs in such a way that these allocations do not intersect with the allocations of this CPE's interfering CPEs. Another use of this information is for bandwidth request purposes (see 6.1.3). In this case, the CPE may include constraint elements when requesting upstream bandwidth allocation to the BS, thus providing the information the BS would need to avoid allocating time for this CPE which interferes with other collocated CPEs. Yet another alternative is for the CPE not to send anything to the BS. Here, the BS would have to specifically send a TRC-REQ message (see 8.23) to the CPE requesting for any constraints it might have regarding allocation. Other uses besides these are also possible.

21.2.1.3 Trigger

As it is discussed later, CBP packets are used for multiple purposes in CMAC such as establishing and keeping synchronization, as well as for self-coexistence. Therefore, the process carried out at the BS to decide when and in which mode (i.e., passive or active) to schedule the Self-Coexistence IUC is mostly implementation dependent.

CMAC is capable, however, of providing the necessary information to assist the BS in this decision process. One example can be found in 8.21.3.1.3. In this case, it is recommended using the CPE statistics report as the basis for triggering the execution of CBP. For example, a decision criterion may be defined such that if the PER experienced by one or more CPEs (clustering can be used here) exceeds a predetermined threshold value PER_{CBP}, this would trigger the BS to schedule a Coexistence IUC for at least the corresponding CPEs.

Another simpler strategy for the BS is to implement a pseudo-random process wherein self-coexistence windows are statically scheduled with a certain frequency, but the mode (i.e., whether passive and active) is pseudo-random. This process is denominated pseudo-random in the sense that it can take into account other statistics in the decision process, such as traffic pattern.

21.2.2 Inter-BS Communication

CMAC incorporates inter-BS communication by allowing both BSs and CPEs to detect and receive SCH transmissions from other collocated BSs. In case of the BS, it may either periodically listen to or even schedule downstream/upstream per frame quiet periods (i.e., Coexistence in Passive Mode – see Table 33 and Table 44) with the goal of detecting SCH frames transmitted by other BSs within its transmission range. If SCH frames are received, the BS would have access to the channels other collocated BSs are employing for transmission, and this could be used for selecting disjoint channels and hence avoid self-interference. As for CPEs, they have the capability to receive SCH from collocated BSs and report back to its own BS (see 8.21.3.1.2).

Another possibility (as discussed in 21.2.1) is that a BS receives CBP packets (either during normal operation or during quiet periods). In this case, the BS would have not only information about channels being used, but also about specific time schedules. This would allow a finer control of self-coexistence, which may be desirable especially in the case where there are no other free channels where BSs can switch to.

21.2.3 Inter-BS Dynamic Resource Sharing

Finding unoccupied channels to operate in an WRAN system requires intensive computation. Spectrum resource sharing among neighbouring WRANs saves computation via cooperation among WRANs and facilitates the opportunistic usage of TV bands. WRAN systems may log the channel transactions and the settlement will be made afterward. Or, we may ask the WRAN systems to be supportive of each other and share the resources they have acquired for free.

Such resource sharing is facilitated by the use of the CBP protocol. As we have seen, CBP allows 802.22 stations (BS and CPEs) to transmit beacons that provide its recipients enough information to achieve better self-coexistence and resource sharing among overlapping 802.22 cells. These beacons are intended for inter-cell communication and carry specific information about a CPEs association and downstream/upstream bandwidth allocations with the BS, while keeping interference levels to a minimum.

21.2.3.1 Resource Renting among WRANs

Let the offerer stations be the ones assumed by the neighbouring base station to have available resource and the renter station the neighbour base station in need of spectrum. The renter station initiates broadcasting of the resource partition request through the CBP. Upon reception of the resource request, the neighbouring WRANs respond through CBP with their active and candidate sets. The union of candidate sets from the neighbour BSs forms the grand candidate set for the renter. The transaction is completed by sending the channel number chosen and the amount of renting time to the neighbour WRANs and receiving the acknowledgement from the offerer.

21.2.3.2 Dynamic Resource Offering among WRANs

In offering, the offerer initiates the intent to share its resource by broadcasting the resource advertisement message. The message contains the channel set information as described in the above sub-section. The renter in need of resource makes a renting request on a specific channel in the candidate set and the renting time to the WRAN BS, which has sent the advertisement message. Finally, the offerer finishes the procedure by sending acknowledgement to the renter BS.

21.2.3.3 Spectrum Etiquette

The channel selection at the renting station should obey the spectrum etiquette rule so that the chosen channel interferes with only a minimum number of channels being used in the neighbour WRANs. This is the reason why the offering BSs send not only the candidate sets but also the active sets of their own.

The number of channels which can be delivered in renting-offering processes should be optimized as the IEEE 802.22 standardization progresses so that high data rate efficiency is maintained while the spectrum etiquette rule is met as tightly as possible.

21.2.4 CBP Measurements

As presented in 8.22, CMAC also supports the feature wherein CPEs can receive and report information about overheard CBP packets. To this end, the BS shall include in the BLM-REQ message a Beacon Measurement Request element which requires the CPE to listen to CBP packets and report back to the BS. This way, the BS can obtain accurate information about other collocated 802.22 cells through its own CPEs, which is extremely useful in scenarios where BSs are not within radio range of each other and hence Inter-BS communication is not feasible. As discussed earlier, the CBP reports obtained by the BS allow it to take various actions such as change channels and employ interference-free scheduling.

To increase the probability for receiving CBP packets (and also BS beacons – see 21.2.2), the BS may periodically schedule short per frame quiet periods or self-coexistence quiet periods. These quiet periods are different than the incumbent quiet periods given that their time duration can be much shorter, and serves the

purpose of allowing primarily better self-coexistence. In addition, CMAC also includes a mechanism for synchronization of overlapping BSs (see 21.3), which significantly enhances the effectiveness of CBP.

21.3 Synchronization of Overlapping BSs

The key aspect that undermines possibly all existing solutions for coexistence in reservation-based wireless systems is the lack of synchronization amongst co-channel overlapping BSs. Synchronization is a hard problem to be solved, but the benefits gained from having it can be so significant that it is worth pursuing.

Traditionally, synchronization amongst overlapping BS has been a tackled through the backhaul. This simplifies both the PHY and MAC design, but it has as one of its major drawbacks the fact that it relies on third parties. In the particular case of 802.22, another critical drawback includes the fact that this technology is going to employ license-exempt operation, and hence the existence of a common backbone amongst competing operators serving a given location is very unlikely and cannot be assumed. This is further aggravated by the much longer coverage ranges that are expected from 802.22 cells.

Since coexistence is key in 802.22, synchronization becomes very critical in order to allow the 802.22 system to operate at its peak performance. In the case of 802.22, synchronization is beneficial both in the case of incumbent protection as well as for self-coexistence. In case of incumbents, synchronization is beneficial as it allows the quiet periods of overlapping BS to be synchronized (see 21.5.1). This will further enhance the incumbent detection probability, which can otherwise be compromised if it occurs randomly. In the case of self-coexistence, synchronization will make the self-coexistence mechanisms (see 21.2) to be much more effective, and hence provide efficient sharing of radio resources by overlapping 802.22 cells.

Based on the aforementioned facts, in this section we propose a scheme that allows overlapping BS to synchronize by aligning their frames in time. The scheme proposed addresses the problem with an over-the-air approach, in the sense that it does not rely on any sort of fixed backhaul infrastructure.

21.3.1 Assumptions

For any synchronization scheme to be mostly effective, some constraints should be imposed on the overall frame timings. In the specific case of CMAC, we propose that superframes should have the same and fixed length in terms of time, or at a minimum shall be an integral multiple of each other. Individual frames within a superframe should also have the same and fixed size, and at a minimum shall be an integral multiple of each other such as shown in Table 27. This will facilitate not only in establishing synchronization amongst overlapping cells, but, most importantly, in keeping it with very low overheads.

Last, but not the least, we assume that no GPS device is available to the 802.22 BSs. If such device is mandatory to be available, the task of synchronization becomes much simpler and can be accomplished by imposing an additional requirement that BSs shall only initiate superframes at specific points in time. Synchronization will, therefore, be guaranteed and may preclude the need for any real-time scheme.

21.3.2 Establishing Synchronization

An 802.22 cell shall actively seek for other overlapping 802.22 cells in order to establish synchronization, as well as provide means by which other collocated cells can find it. Besides the capability that a CPE shall have to be scanning for beacons whenever it is not transmitting or receiving (see 21.2.1), two other mechanisms shall also be used for this purpose which will considerably increase the probability of a successful synchronization: *self-coexistence quiet periods* and *self-coexistence windows* (in both passive and active modes).

Besides quiet periods for the detection of incumbents, the BS shall also schedule quiet periods for the purpose of self-coexistence (see 8.21.7 and 8.22.1), and these are hereby called as self-coexistence quiet periods. Typically, however, these quiet periods need not be as frequent as those for the detection of incumbents, although the BS has complete freedom to choose their occurrence[21]. During this time, both CPEs and the BS shall search for CBP or SCH packets transmitted by overlapping 802.22 terminals belonging to other 802.22 cells. Whenever a BS powers up, it shall never schedule self-coexistence windows in active mode (i.e., CBP transmissions) before at least one self-coexistence quiet period. This is done to ensure that, with high probability, a new 802.22 cell shall first synchronize with any other collocated 802.22 cell before announcing its existence through CBP packets.

Self-coexistence quiet periods should always be scheduled within the boundaries of a superframe, and shall be done in a random way to increase the probability that overlapping BSs successfully detect each other. The duration of a quiet period will typically be of one frame. The BS shall randomly pick the frame number between [0, FS-1], while the superframe number shall be derived from [0, NSIQP], where NSTQP is the Number of Superframes within an Incumbent Quiet Period. NSIQP can be easily derived from the TTQP field defined in Table 1. By doing this, we are enforcing that the frequency of self-coexistence quiet periods shall be the same as the incumbent quiet periods. Obviously, this can be dynamically changed by the BS if it can estimate an increment or decrement in the number of overlapping BSs (e.g., through PER statistics reported by CPEs or backhaul signalling).

Self-Coexistence windows (as depicted in Figure 28 and defined in 8.4.1) shall also be used for this purpose, and shall always be scheduled by the BS at the end of the US subframe (see Figure 28). The first key difference between self-coexistence quiet periods and self-coexistence window is the time granularity. While the former will typically take at least an entire frame, the latter happens within part of a frame. The second, and probably most significant, difference is that during the self-coexistence quiet periods CPEs and the BS do not perform any type of transmission, but only sense the channel. During self-coexistence windows, however, CPEs can transmit CBP packets if so scheduled by the BS (in the case of self-coexistence in active mode). Section 217.0.521.2.1.3 elaborates on the decision process at the BS used to determine whether to schedule a passive or active self-coexistence period.

For every CBP or SCH packet received, the BS and CPEs shall record the frame offset where they were received. Accuracy in this recording is critical for a successful synchronization. Figure 67 depicts the relationship between the Transmission Offset (see Table 8) and Reception Offset (see Table 164) fields for a frame of size F_S (in units of symbols). These fields are key for establishing synchronization between two overlapping 802.22 cells.

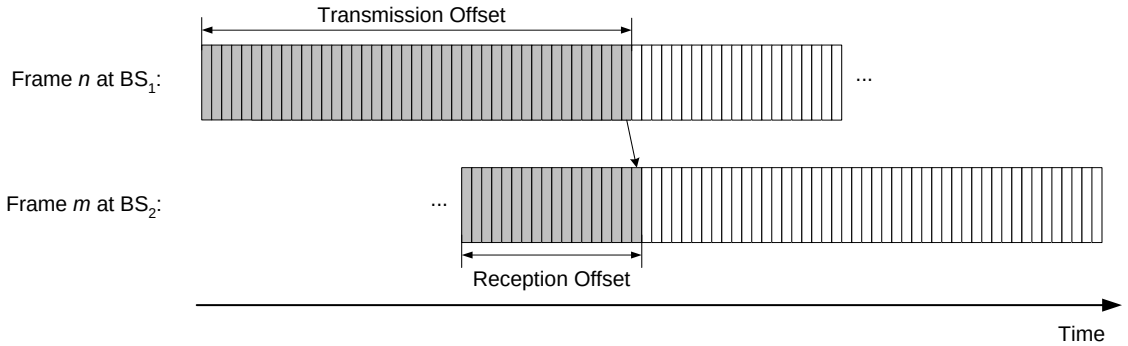

Figure 67 – Synchronization of overlapping BSs

In order for this fully distributed synchronization process to converge within an acceptable duration, a *Convergence Rule* shall be employed by the BS before any synchronization attempt. The correct application of this convergence rule shall guarantee network convergence in all scenarios. Mathematically speaking, BS_i,

[21] In practice, self-coexistence quiet periods could be scheduled during low peak hours, such as overnight, without causing any major impact on the system performance and responsiveness.

responsible for cell *i*, shall only attempt synchronization to BS_j, responsible for a neighbouring cell *j*, if and only if:

$$\left|\begin{array}{l}((FS_i - Frame\_Number_i - 1)\times FDC_i + (FDC_i - \text{Re}ception\_Offset)) - \\ ((FS_j - Frame\_Number_j - 1)\times FDC_j + (FDC_j - Transmission\_Offset))\end{array}\right| \le \frac{FS_i * FDC_i + 3 * SymbolSize}{2}$$

In this equation, $Frame\_Number_i$ is as specified in Table 164, and $Frame\_Number_j$ is as specified in Table 8 in the case of CBP and shall be –1 in case of a SCH. The other parameters to this equation are defined in Table 1 and in Figure 67. Given the requirement that $FS_i = FS_j = FS$ and $FDC_i = FDC_j = FDC$, we can further simplify this equation to:

$$\left|(Frame\_Number_j - Frame\_Number_i)\times FDC + Transmission\_Offset - \text{Re}ception\_Offset\right| \le \frac{FS * FDC + 3 * PHYSymbolSize}{2}$$

Therefore, BS_i shall apply this convergence rule to each and every synchronization alternative available. Only those that satisfy this rule can proceed to the next phase.

Even after the convergence rule is applied to all possible synchronization alternatives, although unlikely, multiple choices may still remain that satisfy the convergence rule. The BS shall attempt synchronization with each overlapping BS, one at a time. Unless the BS realizes that it is already synchronized with the overlapping network corresponding to the selected packet (i.e., Slide Amount equal to 0 or F_S – see below –, or through bookkeeping), the BS shall immediately construct and transmit a FSL-REQ message (see 8.25.1) as a broadcast to all CPEs in the cell, and shall not schedule any additional active mode self-coexistence interval until the scheduled time of the frame slide.

In constructing the FSL-REQ message, the Slide Count and Slide Amount fields shall be configured as depicted in Figure 68. Slide Count shall be equal to the number of frames right before the start of the superframe of the overlapping BSs, and Slide Amount shall be equal to the number of slots to the start of the overlapping BS' superframe. More specifically, the Slide Amount field and Direction fields in the FSL-REQ message shall be set according to the following rules:

$$Slide\ Amount = \begin{cases} FDC - Transmission\_Offset + \text{Re}ception\_Offset, & if\ (FDC - Transmission\_Offset + \text{Re}ception\_Offset \le \left\lceil \frac{FDC}{2} \right\rceil)\ \ (Case\ 1) \\ Transmission\_Offset - \text{Re}ception\_Offset, & otherwise\ (Case\ 2) \end{cases}$$

$$Direction = \begin{cases} 0\ (Right), & if\ (Case\ 1) \\ 1\ (Left), & otherwise\ (Case\ 2) \end{cases}$$

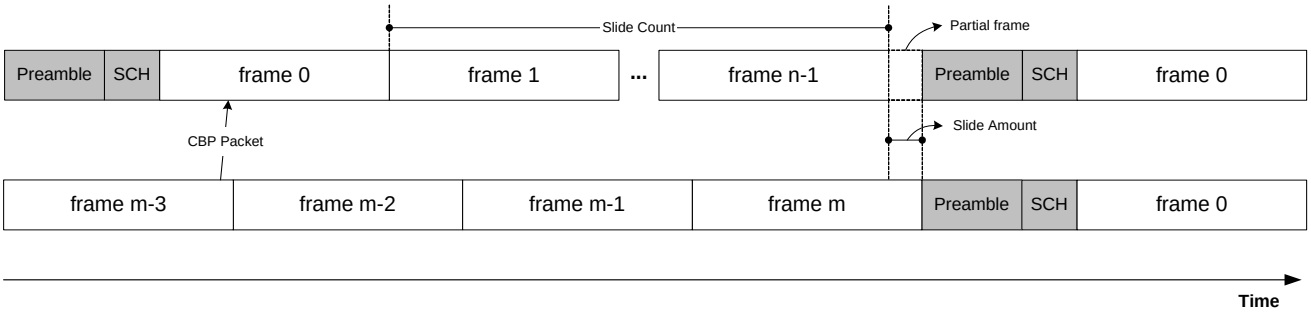

Figure 68 – Establishment of synchronization between overlapping BSs

On the CPE side, it shall report back to the BS on any received CBP or SCH packets, unless it receives a FSL-REQ message from the BS in the meantime in which case it shall terminate the self-coexistence procedure and return to normal operation. Regardless of whether it was the BS itself who detected the CBP or SCH packets, or if these were received as a result of a CPE report, the BS shall proceed in the same way as described above in attempting synchronization. The key difference at the BS is that the CPEs who received the same SCH and CBP packets will report different values for the Reception Offset due to the fact that they experience different propagation delays. In this case, it is up to the discretion of the BS to select one of the packets, as these refer to the same network. Section 21.3.3 describes the mechanism that the BS uses to cope up with the different propagation delays, which essentially consists of allowing guard bands.

Whenever sliding the frame to the left or to the right, the BS always shall adopt the same behaviour. That is, at the scheduled slide time the BS shall initiate transmission of a new superframe (see Figure 68). So that the frame slide does not disrupt any data communication, the BS scheduler shall take this slide into account when scheduling US and DS transmissions.

The frame slide operation may result in a partial frame as indicated in Figure 68. It is up to the BS to use this partial frame as it sees fit. For example, the BS could use this time as a quiet period. Another possibility is that the scheduler at the BS has the capability to schedule data transmissions during this time, and so airtime is not wasted. Yet another option would be to keep this partial frame as idle time, which may be a simple strategy whenever the partial frame size is only a few slots.

21.3.3 Confirmation and Maintenance of Synchronization

Once synchronization is accomplished, maintenance is a simpler process. Once the FSL-REQ message goes into effect and the frame is shifted, the BS shall schedule self-coexistence windows with a certain periodicity and always at the end of the US subframe (see Figure 28).

Confirmation and maintenance of synchronization is performed through periodic CBP packet transmissions and receptions during self-coexistence windows. Once the first CBP packet is successfully received from the overlapping cell, synchronization is completed and confirmed. At this point, the BS shall continue to schedule self-coexistence windows, but now with the main purpose of better self-coexistence through the exchange of traffic constraints. Of course, a positive side effect of the transmission and reception of CBP packets is that the overlapping cells can more easily remain synchronized. This will, in effect, provide the most performance gains for all synchronized 802.22 cells. Figure 69 shows an example of how two synchronized 802.22 cells communicate over the air through the self-coexistence window.

Given the large propagation delays in a 802.22 network, synchronization from the point of view of the BS does not mean exact synchronization for all CPEs. As discussed before, this is due to the different propagation delays experienced by different CPEs. To account for this disparity in propagation delays, and to accommodate the preamble transmission and the contention backoff interval, the BS shall schedule self-coexistence windows to be at least three slots wide. This way, the BS can provide a guard band and take into account the worst scenario for the transmission/reception of at least one CBP packet.

If a frame slide in a cell is the result of a report from one or more CPEs, together with the BS these CPEs shall be responsible for keeping track of the synchronization. During the maintenance phase, the CPEs shall periodically transmit/receive CBP packets to/from the synchronized cell in order to confirm continued synchronization. During this time, the CPE should not report to the BS on each and every CBP packet received. Rather, it shall restrict its control information exchange with the BS to only those that are needed for better self-coexistence and to implement “interference-free scheduling”. If, however, a CPE does not receive a CBP packet within some pre-determined period of time, it shall deem the synchronization with the corresponding cell to be lost and shall go back to proceed as discussed in 21.3.2.

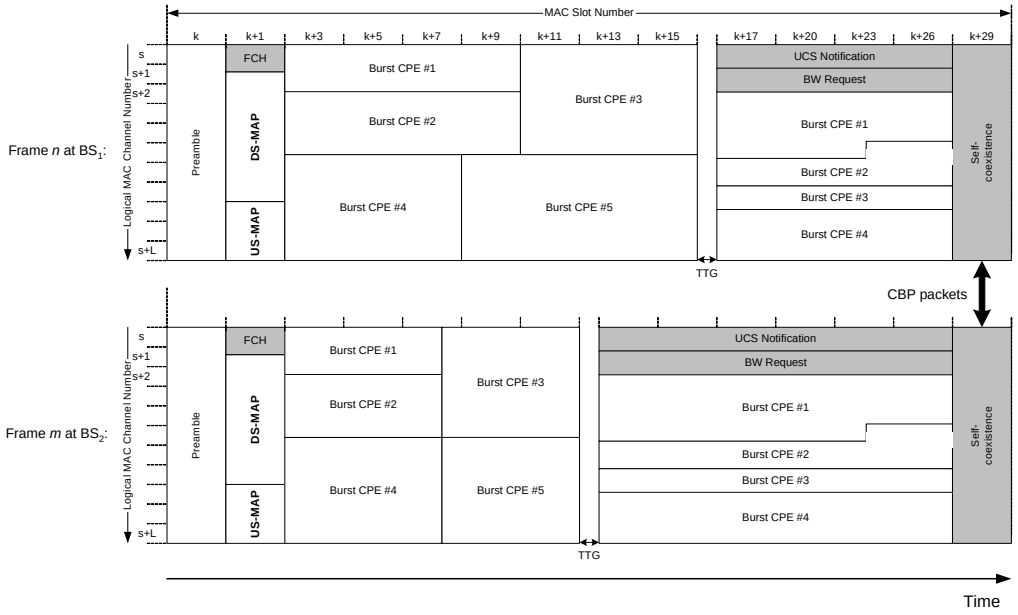

Figure 69 – Communication between two synchronized overlapping cells

21.4 Clustering

While absolutely necessary, the coexistence requirements present in 802.22 may have a significant impact on the overall cell performance. Here, we propose clustering as the solution approach to alleviate much of the overhead and redundancy involved in the execution of the coexistence mechanisms (whether with incumbents of self-coexistence) of an 802.22 cell. These clustering mechanisms would be coordinated by the BS, which would cluster CPEs together based on certain criteria (discussed later).

A clustering scheme for 802.22 is very suitable for many reasons. First, CPEs within an 802.22 cell are fixed, which allows cluster membership to remain nearly unaltered for a long period of time and hence reduces the control overhead. Secondly, in case of incumbents, the radio range of TV stations is very large which results in the sensing outcome of close-by CPEs to be *similar*[22]. So, in this case the BS is able to distribute the load of sensing across CPEs belonging to the same cluster, such that the sensing outcome of one CPE is also valid for all other CPEs belonging to the same cluster. Thirdly, in case of self-interference, clustering allows the BS to group together those CPEs that are less likely to contend for the medium simultaneously, hence decreasing both the access time by a CPE as well as the duration of time that has to be allocated by the BS for the CBP.

Figure 70 illustrates the overall concept of clustering. As it can be seen, the idea is that the BS would group CPEs together and would later utilize the CPEs belonging to these individual clusters to distribute the load of coexistence. For example, certain CPEs within a cluster could be responsible for sensing ATSC, while another CPE might be responsible for sensing DVB. The BS would then use the results of the sensing from a CPE within a cluster as a measure for the entire cluster. In the case of CBP, the main optimisation goal of the BS is to increase the probability that coexistence beacons be successfully transmitted. To this end, at a given Coexistence IUC, the BS shall only select a single CPE from each cluster (e.g., shown in Figure 70) to send coexistence beacons. By doing this, the BS aims at increasing the probability that these CPEs transmit beacons at nearly simultaneously, hence improving the spectrum utilization efficiency.

[22] The notion of "similar" sensing measurements can be based on thresholds. This issue is discussed in a later subsection.

It is important to note that the BS shall not use the clustering scheme for the protection of Part 74 services (this is the only exception). This is due to the very short transmission range of Part 74 devices as compared to TV stations. Hence, to provide effective protection, all CPEs shall perform sensing of Part 74 services.

21.4.1 Formation

The process by which a BS decides which CPEs to cluster together is very dependent on the criteria used. The criteria, in turn, depend upon what is the goal of the clustering (i.e., what will clusters be used for). Since in CMAC clustering is primarily used for improving the performance of coexistence mechanisms, the goal depends on the type of system to be protected: incumbents or self-coexistence. To address this issue, CMAC incorporates a two-stage process which are applied in any of the cases: first, the creation of Physical Clusters, and second, the creation of Logical Clusters.

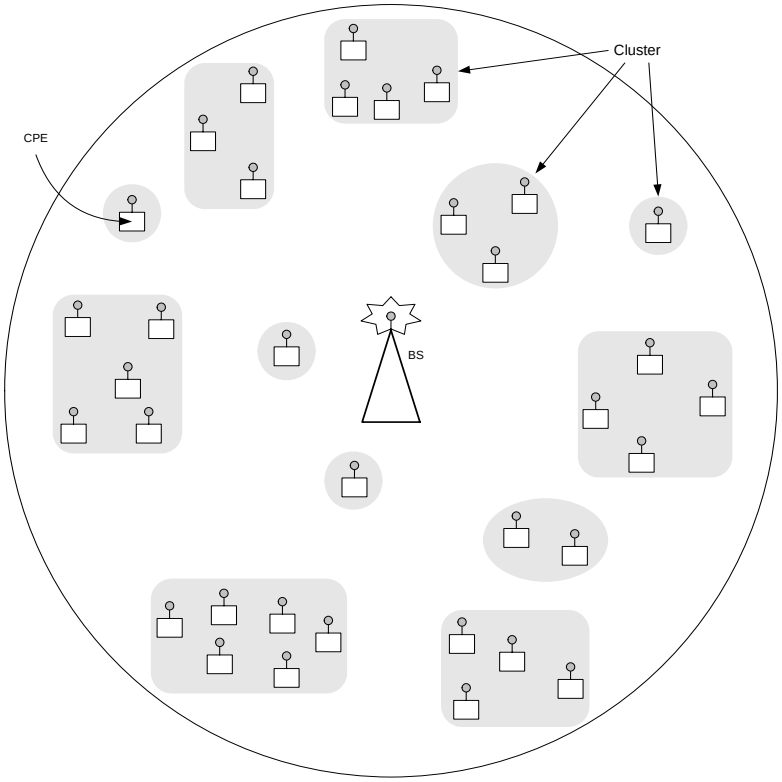

Figure 70 – Example of clustering. To improve performance, the BS groups CPEs into clusters based on a given selection criteria.

21.4.1.1 Physical Clusters

The first procedure a BS has to complete as to implement the clustering functionality is the creation of Physical Clusters (PCs). The process for the creation of these PCs is totally localized at the BS, and does not directly involve any CPE.

A PC can be defined in many ways. In CMAC, a PC is defined as a set of one or more CPEs whose incumbent sensing outcome for given channel(s) and for given RF profile(s), is similar. In other words, if the sensing outcome is similar, this means that the measurement outcome for, say, ATSC performed by a CPE C_1 within cluster G, can be generalized as valid for all $C_i \in G$, $i > 0$. Figure 71 depicts an example of PCs. Note that since the BS receives feedback from CPEs regarding their measurements outcome, the BS can easily implement an internal procedure to create physical clusters based on the measurements reported. Also, since PCs implicitly give

an indication of physical location, these can also be used for CBP clustering because devices within a PC are likely to contend with themselves for transmission of coexistence beacon.

Note that, as also shown in Figure 71, there may be cases where the number of CPEs within a PC is one (also referred to as a *unit cluster*). This may indicate two things. First, it may be that the BS was not able to find more than CPEs similar measurements outcomes. Secondly, the BS may disable clustering and hence each CPE would be considered as a PC.

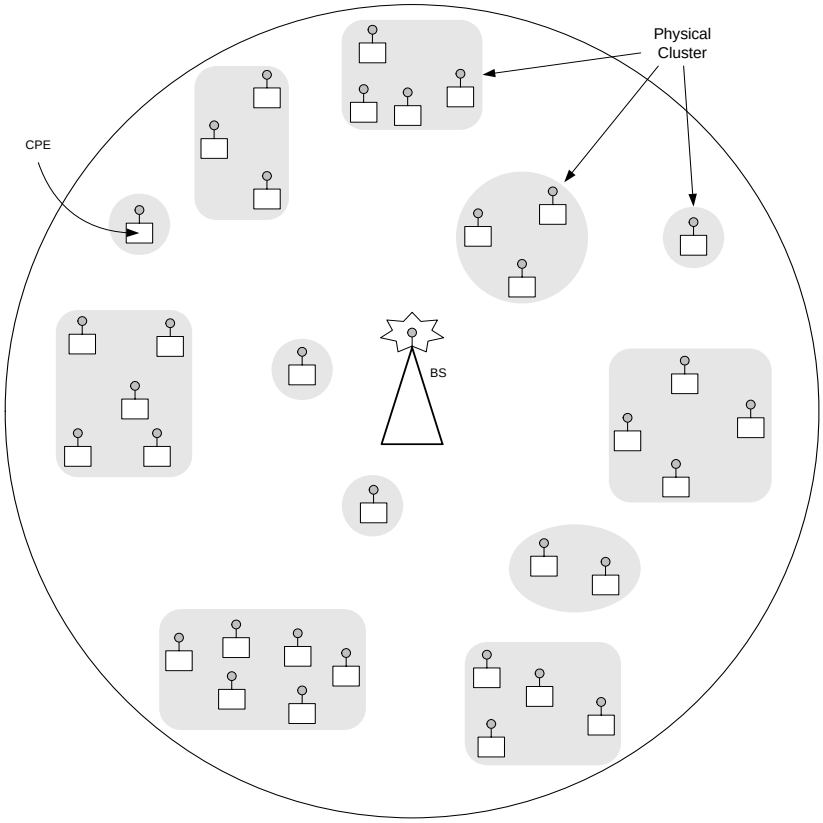

Figure 71 – Concept of physical clusters (created by the BS without any CPE involvement)

21.4.1.2 Logical Clusters

Once PCs have been internally created by the BS, the second step, namely, creation of Logical Clusters (LCs), can be started. The creation of LCs consists of the BS sending MCA-REQ to CPEs belonging to different PCs, wherein the same multicast management CID is used.

Figure 72 shows the creation of LCs, and how they relate to PCs. As we can see from this figure, CPEs belonging to different PCs are grouped into LCs, and each of these LCs can be simultaneously scheduled by the BS to perform a given coexistence task. For example, in Figure 72, it is indicated that unless a PC has a single CPE, different CPEs belonging to the same PC perform different type of RF measurements. Some may perform ATSC, while others may be responsible for NTSC, DVB, and so on. However, all CPEs participating in a LC perform the same type of incumbent measurement, and so the coexistence tasks associated with these LCs can be scheduled with a single management message (e.g., BLM-REQ) which considerably reduces overhead.

A similar procedure applies to scheduling Coexistence IUC in the case of CBP, as also indicated in Figure 72. CPEs belonging to a LC are less likely to contend for access, and hence the BS can optimise execution of the CBP by scheduling all CPEs belonging to a single LC to send coexistence beacons at the same scheduled time. This

improves performance as it decreases the likelihood of collisions at the receivers, as well as reduces the medium access time due to lower contention among CPEs within the same LC.

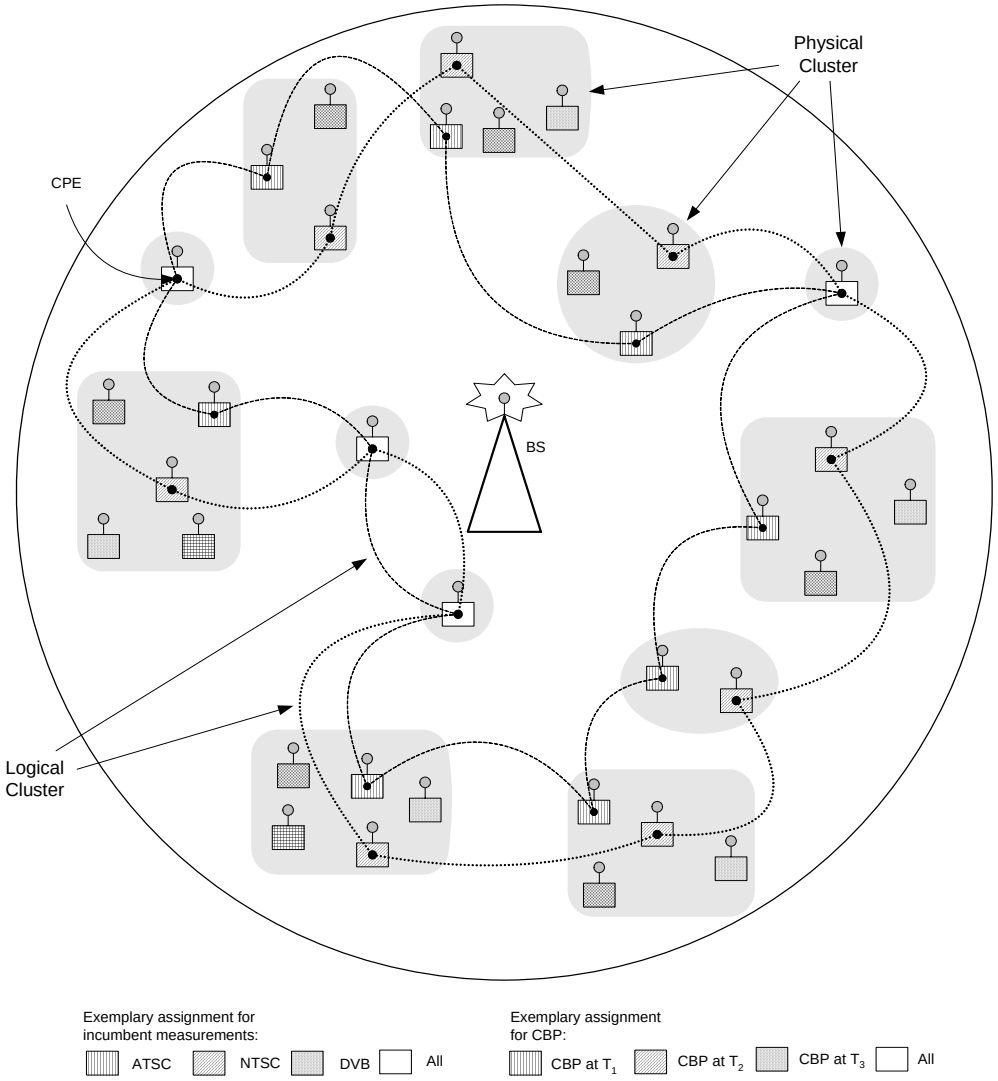

Figure 72 – Concept of logical clusters (CPEs across physical clusters are grouped and perform the same coexistence activity. BS and CPEs participate in this process.)

Note that, as also indicated in Figure 72, since unit clusters do not have more than one CPE to share the measurement load, these clusters shall be instructed by the BS to perform all type of measurements (possibly belonging to multiple LC, as shown in Figure 72). This is specially the case for incumbent measurements as they are more critical. In case of transmission of coexistence beacons though CBP, the BS does not need to include a unit clusters in a LC unless there is a need for self-coexistence.

21.4.2 Algorithm

We propose the use of the k-means clustering algorithm for implementing clustering in a 802.22 cell. This is a widely studied and understood algorithm that is a simple (unsupervised learning) and can be used to cluster data. Although this algorithm is guaranteed to converge, it is not always guaranteed to converge to the global optimum (in practice, this algorithm is run multiple times to overcome this problem). However, for the application under consideration, namely, incumbent protection, since the transmitters do not change their location and/or their

transmit power very often (it is important to keep in mind that clustering is proposed to be used solely for TV protection), the algorithm does not need to run too many times.

Figure 73 presents the proposed clustering algorithm in detail. Initially, no clustering is present in the network and each CPE conducts measurements in the incumbent TV channels (as discussed in 21.1.3) as per requested by the BS. Based on these measurements, each CPE constructs its own incumbent profile (see Figure 74) and transmits it to the BS through the BLM-REP management message. Then, the BS performs clustering as described in Figure 73. This process can be repeated intermittently as CPEs update the BS with their incumbent profiles.

21.4.3 Implementation

Clustering is implemented through the use of MCA-REQ and MCA-RSP messages. The BS issues MCA-REQ messages to the CPEs who respond with MCA-RSP. Since CPEs are fixed terminals, cluster membership is not likely to suffer major changes over time, which further improves the efficiency of this mechanism.

The BS shall use multicast management CIDs for the purpose of implementing clustering (either physical or logical). The multicast management CID would allow CPEs to filter whether or not a measurement request is addressed to itself. Once clusters are constructed, the BS specifies the intended multicast management CID in the MAC header (see 6.1.1) whenever it requests CPE to perform a coexistence task. For example, in case of incumbents, whenever the BS sends a BLM-REQ to the CPEs it shall specify the multicast management CID in the MAC header. In case of self-coexistence (i.e., beacon transmission performed by CBP), the multicast management CID would be specified whenever the BS schedules a Coexistence IUC. For self-coexistence, however, a multicast management CID would be associated with a logical cluster since this would reduce the likelihood that those CPEs contend for the same shared medium, and hence decrease the collision probability at the receiver.

GOAL: Cluster CPEs with similar incumbent profiles.
TERMINOLOGY:

n: Number of devices in a cell, scalar

f: Total number of channels, scalar

k: Number of clusters, scalar

i: Device index, *1 … n*

j: Cluster index, *1 … k*

$x_i^{(j)}$: Measurement vector for device *i,* belonging to cluster *j*, size $1 * f$

$\overline{m_j}$: Cluster mean, vector, size $1 * f$

J: Scalar objective function to be minimized

J^*: Maximum allowed value for scalar objective function (input).

*k** : Optimal number of clusters, (output)

$\| \ \|$: Distance (in feature space) between devices

INPUT: J^*

OUTPUT: (a) Number of clusters *k**, and (b) each device assigned to one of the clusters (see Figure 3).

ALGORITHM: k-means clustering algorithm consists of the following steps.

1. Initialize *k=2*.
2. Randomly assign *k* values from $x_i^{(j)}$ to $\overline{m_j}$. These serve as initial guess for the k means.
3. Until there is no change in means, perform the following steps
   a. For each device, determine which of the means $\overline{m_j}$ is closest to its measurements vector $x_i^{(j)}$, and assign the device to that cluster, *j*.
   b. For each cluster, update $\overline{m_j}$ with the mean of all the samples for cluster *j*.
   c. Compute *J*.
      1. If *J < J** then GOTO step 4.
      2. Else, increment *k* by one and GOTO step 2.
4. STOP. Obtain outputs *k** and $x_i^{(j)}$.

OBJECTIVE FUNCTION: Mathematically, step 3 can be achieved by minimizing the following scalar objective function.

$$J = \sum_{j=1}^{k} \sum_{i=1}^{n} \left\| x_i^{(j)} - \overline{m}_j \right\|^2$$

where, $\left\| x_i^{(j)} - \overline{m}_j \right\|^2$ indicates the distance between the measurement vector $x_i^{(j)}$ of device *i* tentatively belonging to cluster *j*, and its cluster mean $\overline{m}_j$ in feature space.

Figure 73 – The k-means algorithm for clustering CPEs

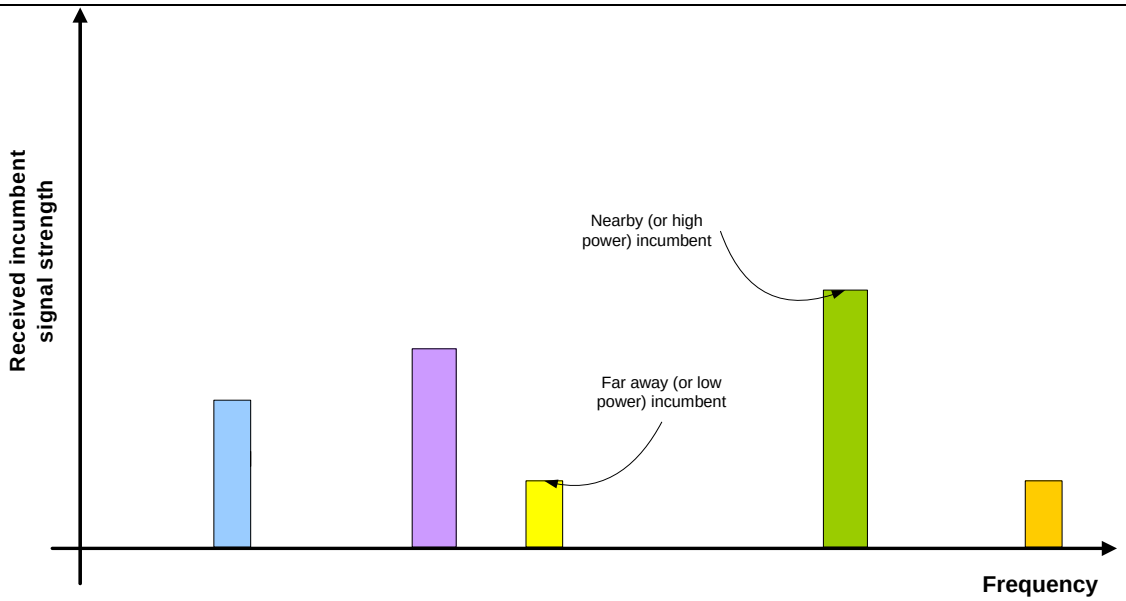

Figure 74 – Incumbent profile at a CPE

21.4.4 Discussion

There are a number of advantages of clustering CPEs. These advantages relate either to sharing measurement function within the wireless network or to efficient dissemination of measurement information.

Once CPEs are clustered together based on similar incumbent profiles, they do not have to make repeated measurement of the entire available spectrum. It is the responsibility of the BS to make the optimal distribution of measurements within a network, which involves the following trade-off. If too few CPEs in the entire network make measurements, an incumbent might be missed. On the other hand, if each CPE searches every channel, the total amount of time it takes to determine which channels are available could be very large. This approach of clustering provides an intelligent tool to make such a trade-off as shown in Figure 75.

If all the devices measure all the channels and disseminate this information over the network, the load on the network could be significant. By decimating the number of measurements made, the dissemination overhead can be considerably reduced. Note that how often a given channel must be measured for occupation by a primary depends not on the duty cycle of the primary (which may be of the order of a day), but rather on the vacation time[23], which is of the order of a few seconds or less. When vacation time is small (as it is often the case), unless dissemination overhead is efficiently managed, it could consume a significant part of the total available radio resources. This is especially true when contention-based access mechanisms are used to disseminate measurement information.

Finally, for self-coexistence clustering provides an efficient way to choose which CPEs must transmit coexistence beacons, thus making efficient use of radio resources.

[23] The vacation time is defined as the duration in which the cognitive radio network must vacate a channel after an incumbent comes ON on that channel.

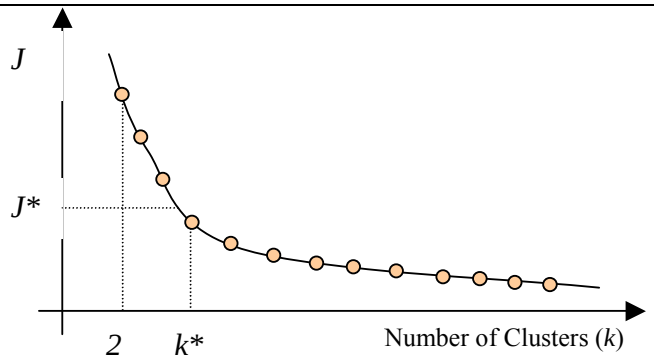

Figure 75 – The number of clusters is increased to *k as to meet the acceptable objective function value *J****

21.5 Quiet Periods

In CMAC, both in-band and out-of-band incumbent measurements are scheduled and coordinated by the BS. In case of in-band incumbent measurements, these shall always be performed when the BS schedules quiet periods in the cell. This is not to say, however, that CPEs shall only sense the spectrum during scheduled quiet period. Whenever not engaged in communication with its BS during normal cell operation, CPEs shall be free to perform either in-band or out-of-band incumbent detection in any channel. In fact, the BS may also specifically request CPEs to perform incumbent measurements even during normal 802.22 cell operation. This may be useful, for example, for the BS to keep track of potential backup channels (see 21.1.5) in case of an UCS. For self-coexistence purposes, quiet periods are not needed although they do increase the probability of detection.

Quiet periods can be scheduled in either the explicit mode, which is done through the use of CHQ-REQ (8.21.7), or in the embedded mode as specified in 8.1.1 (for more information on explicit and embedded channel management, see 21.6). For out-of-band measurements, quiet periods are not necessary and hence the BS can allow, if desired, a certain level of autonomy to the CPE to decide when to perform these measurements. Alternatively, the CPE can use the TMO-REQ message (see 8.24.1) to request the BS a timeout from the cell, and this could be used for the purpose of out-of-band measurements.

21.5.1 Synchronization of Overlapping Quiet Periods

Due to the possibility of multiple overlapping 802.22 cells, it is desirable that the quiet periods of a cell be synchronized as much as possible with its overlapping cells. This will considerably improve the reliability of detection of incumbent signals, and will also enable better self-coexistence amongst overlapping 802.22 cells. Therefore, BSs shall attempt to synchronize their quiet periods with other overlapping cells, which can be done from the TTQP and DQP fields available from both the SCH (see Table 1) as well as CBP packets (see 6.1.2). The BS shall be responsible for setting these fields whenever transmitting a SCH.

The BS who receives information about other collocated 802.22 cells (either directly or reported by CPEs), shall use a random mechanism to attempt synchronization of quiet periods, which will considerably mitigate the *ping-pong effect*. For example, consider that BS_1 received information about SCH_2 transmitted by a collocated BS_2. A basic rule is used by the BS_1 to decide whether to synchronize its quiet period with that of BS_2:

- BS_1 shall only continue with the synchronization if $TTQP_{BS1} > TTQP_{BS2}$. In other words, a BS shall only continue with the synchronization procedure if the remaining time to its next quiet period is larger than its overlapping BS.

If this rule is validated, BS_1 can proceed with the synchronization of its quiet period with that of BS_2. To this end, BS_1 shall schedule the change in its quiet period to take place N frames away, where N = rand(0, Q_{Thresh}). Here, rand(a, b) is a function which returns an integer number t, where $a \leq t < b$, and Q_{Thresh} is defined in units of superframe. If up until N superframes later BS_1 does not receive any more information regarding BS_2, it shall proceed with its quiet period change to accomplish better synchronization. This is done by modifying the values of TTQP and DQP in the SCH when initiating the new superframe, and by transmitting an updated CHQ-REQ command.

If, in the meantime, BS_1 receives information about BS_2 which indicates that BS_2 has already changed its quiet period to align with that of BS_1, BS_1 shall then cancel its scheduled quiet period change (this has a negligible likelihood, however, given the rule above). Another possibility is that the new information about the quiet period that BS_1 receives about BS_2 changed since the last notification. In this case, BS_1 shall cancel the current scheduled change of its quiet period and reschedule it if appropriate (using the same procedure as described above), taking into consideration the new parameters received from BS_2. BS_1 shall proceed with changing its quiet period in all other cases.

The synchronization mechanism discussed in 21.3 makes the quiet period synchronization procedure considerably simpler.

21.5.2 Two-Stage Mechanism for Quiet Period Management

Although the appearance of incumbent users on a channel is sporadic (e.g., TV stations do not come on the air every hour), an 802.22 network shall meet the Channel Detection Time [3] for the diction of the presence of incumbents. Typically, this leads to the need of periodic network-wide quiet periods in order to offer proper protection to incumbent users. During these quiet periods, all network traffic is suspended and stations (BS and CPEs) sense the channel on the look out for primary users. Depending on the algorithm used for these quiet periods, they may be quiet long (e.g., in the order of tens of milliseconds).

Paradoxically, users of the 802.22 network expect the same type of QoS available in existing wireless networks, including support to voice and video traffic. However, since these services demand stringent QoS requirement (e.g., could be as low as 20ms for voice [3]), an obvious question arises: how can a 802.22 network protect incumbents through quieting channels while, at the same time, supporting the expected QoS required by 802.22 users?

With this in mind, here we introduce a mechanism that efficiently addresses this issue. It is based on a two-stage sensing approach: fast sensing and fine sensing. The fast sensing is done before the fine sensing, and typically uses a quick and simple detection algorithm such as energy detection. It is done primarily over in in-band channels, and the outcome of these measurements determine the need and the duration of the upcoming fine sensing.

The proposed mechanism for quiet period management is depicted in Figure 76. It is comprised of two stages both realized through the use of network-wide quiet periods, but which have very different time scales. They are:

- **Fast Sensing**: The fast sensing stage is comprised of one or more fast sensing periods as depicted in Figure 76, where each fast sensing period can be associated with a different frequency channel. During this stage, a fast sensing algorithm is employed (e.g., simple energy detection). Typically, this is done in a few microseconds (e.g., 20μsec in IEEE 802.11b and 9μsec in IEEE 802.11a) and so can be made to be highly efficient. The outcome of the measurements done by all CPEs and the BS during this stage are consolidated in the BS, who then decides on the need for the following fine sensing stage (see next bullet). For example, if during the fast sensing stage it is concluded that energy in the affected channel is always below the threshold, the BS may decide to cancel the next scheduled fine sensing period. Not only that, since each of the many fast sensing periods can be associated with different frequency channels (e.g.,

with total of 3 fast sensing periods, one could be done on channel N, another on N-1, and another on N+1), it may be possible that energy is detected in some channels but not in others. In this case, the BS may decide to change the length of time of the next fine sensing so as to match what really needs more careful measurement. No more, no less.

- **Fine Sensing**: During this stage, more detailed sensing is performed on the target channels. Typically, algorithms executed during this stage can take in the order of milliseconds (e.g., 25ms in the case of field-sync detection for ATSC) for each single frequency channel, since they look for particular signatures of the primary user transmitted signal. In other words, the fine sensing could be over 3 orders of magnitude larger than the fast sensing, which may be unacceptable to QoS sensitive traffic such as voice and video. Therefore, in order to meet the stringent QoS requirements of secondary users while protecting primary users, the presence as well as length of the fine sensing stage has to be dynamic. This is done through the feedback obtained during the previous fast sensing stage.

It is important to note (see Figure 76) that the two-stage mechanism repeats itself after every Channel Detection Time, which is a parameter defined by primary users (e.g., TV broadcasters) or maybe regulatory bodies (e.g., FCC in the Unites States). In addition, contrary to the fast sensing stage, fine sensing stage takes place at most once during the Channel Detection Time given that it takes considerably longer time. Figure 77 shows in detail the how this two-stage mechanism is implemented at the BS and the CPE.

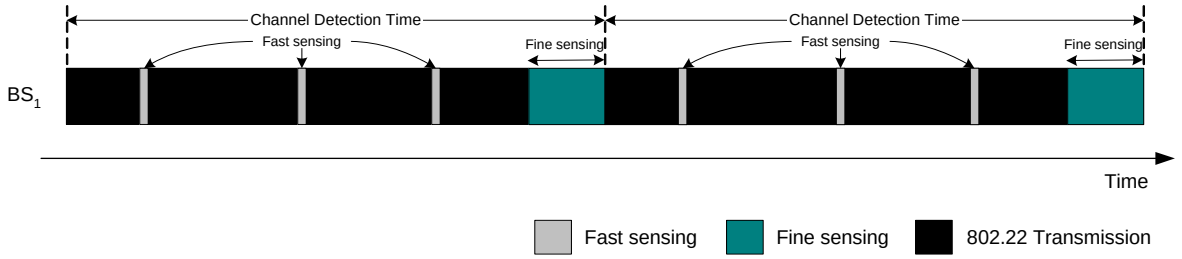

Figure 76 – Structure of the two-stage mechanism for quiet period management

As we can see, this mechanism can offer numerous advantages since it distributes the sensing overhead. It provides proper protection to primary users since the fast sensing (e.g., energy detection) stage will be able to detect the presence of any signal (obviously, above a given threshold such as the noise floor of the measuring device) in the measured channel within the required Channel Detection Time. At the same time, it also enables the support of better QoS to secondary users, since the time consuming task of fine sensing is only performed when it is really needed.

There is one more issue that needs to be addressed by this scheme, namely, when multiple overlapping secondary networks operate in the same frequency channel. Clearly, in this case if the fast sensing stage is implemented through simple energy detection, a secondary network may likely detect the energy of another secondary network and not of the intended primary user. This will disrupt this two-stage approach, as the fine sensing stage would always take place.

To overcome this problem, we use the synchronization mechanism described in 21.3, and which results in the scenario depicted in Figure 78. By doing this, overlapping cells will synchronize not only their frames but also their quiet periods. This will ensure that the result of the fast sensing is highly efficient, since all secondary networks will quiet the channel at the same time and only the signal from the primary user remains in that channel.

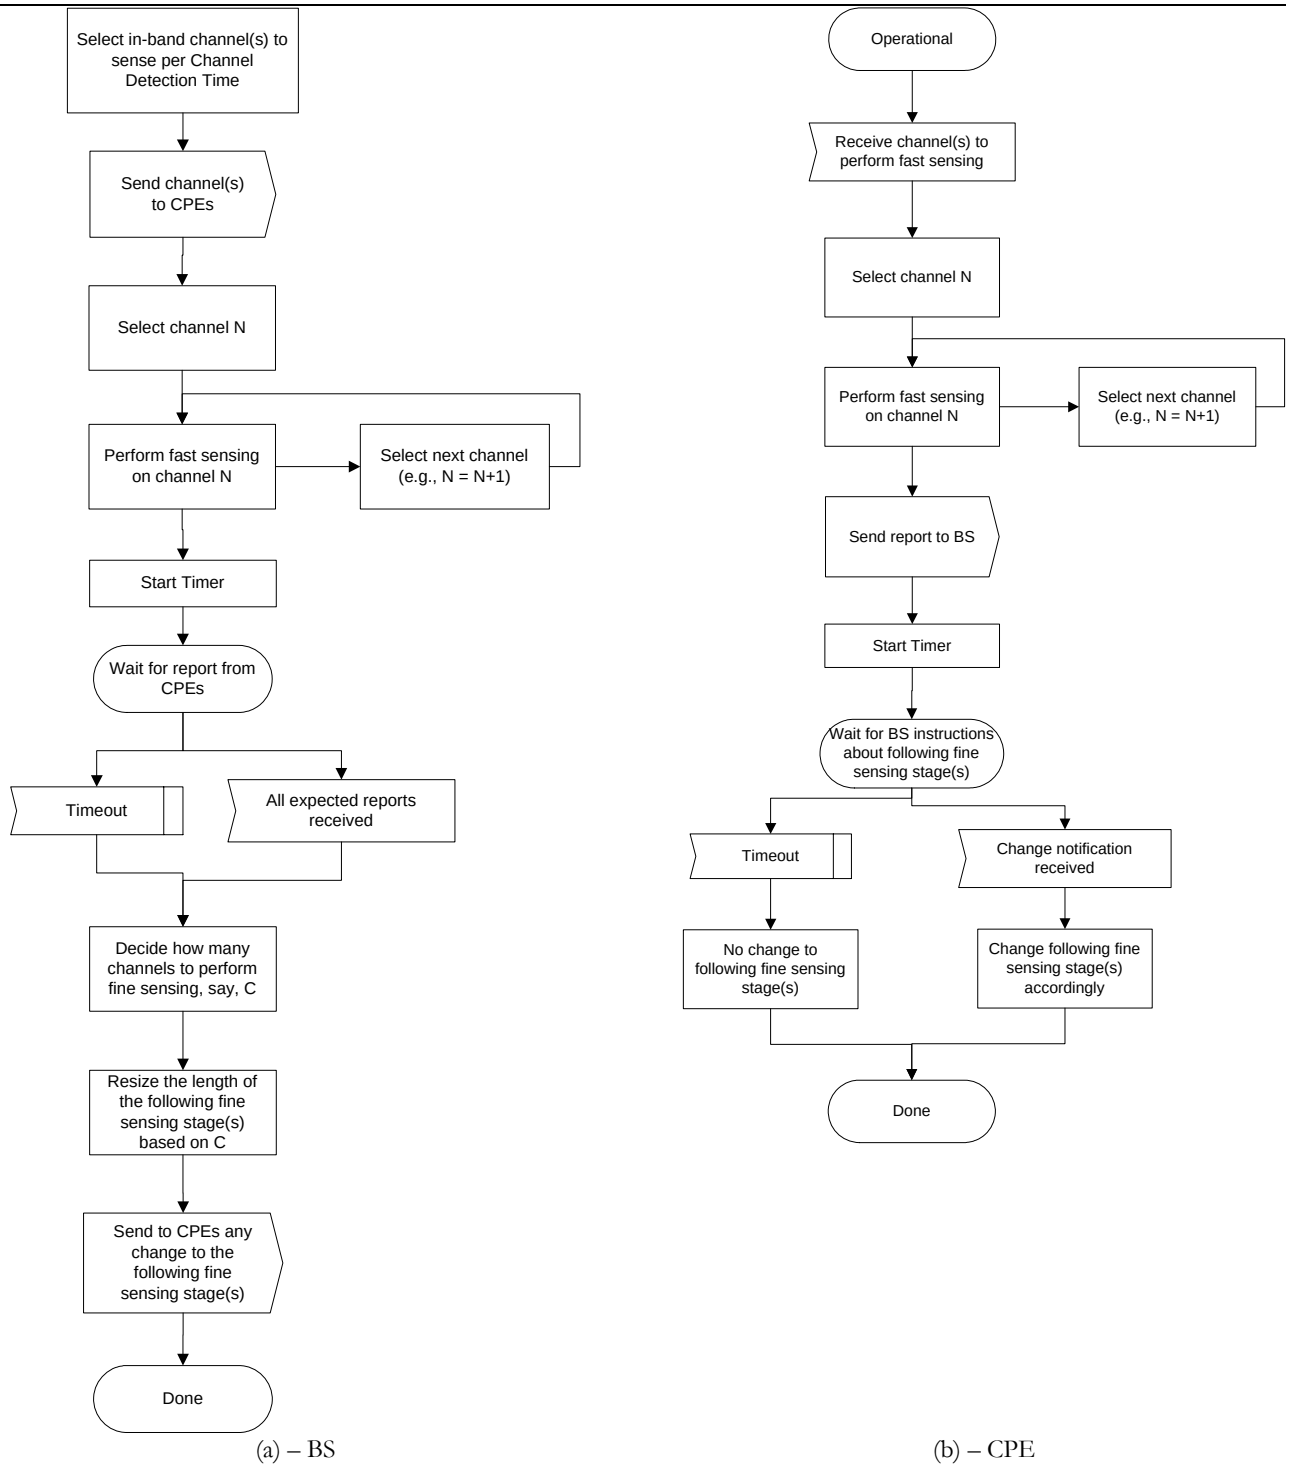

Figure 77 – The two-stage sensing procedure at the BS (a) and CPE (b)

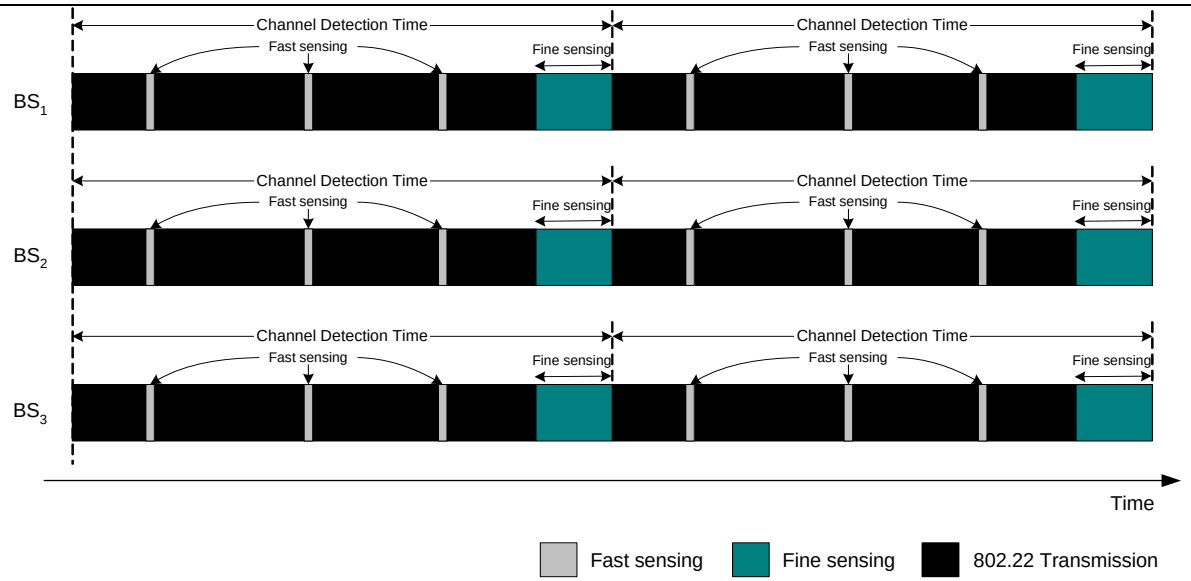

Figure 78 – Quiet period mechanism with multiple overlapping cells in a single channel

21.5.2.1 Fast Sensing Allocation

One of the critical aspects of this approach is that the fast sensing be carried out during a period of time that is common across overlapping networks. This will ensure that all networks perform fast sensing at the same time, and so increase its effectiveness in detecting the primary user.

Here, we propose the fast sensing to be done at the end of the MAC frame as depicted in Figure 79. Since frames of overlapping networks are synchronized, so will the end of frames from different networks be synchronized. Therefore, this guarantees that during this time fast sensing can be performed. This is not the case, for example, if such sensing is performed during the TTG window as shown in Figure 79, as overlapping networks have different ratios between upstream and downstream traffic.

More specifically, we propose to use the RTG window to perform the fast sensing, as this it will be more than sufficient to, for example, perform a simple energy detection. This is shown in Figure 79, where at the end of a frame there is the Sensing RTG slot.

For this scheme to be implemented, the BS has to inform CPEs in which frame fast sensing is to be performed. Not only that, the BS should also specify in which channel to perform fast sensing and how large the Sensing RTG window has to be. This is done using the DCD message as specified in Table 26.

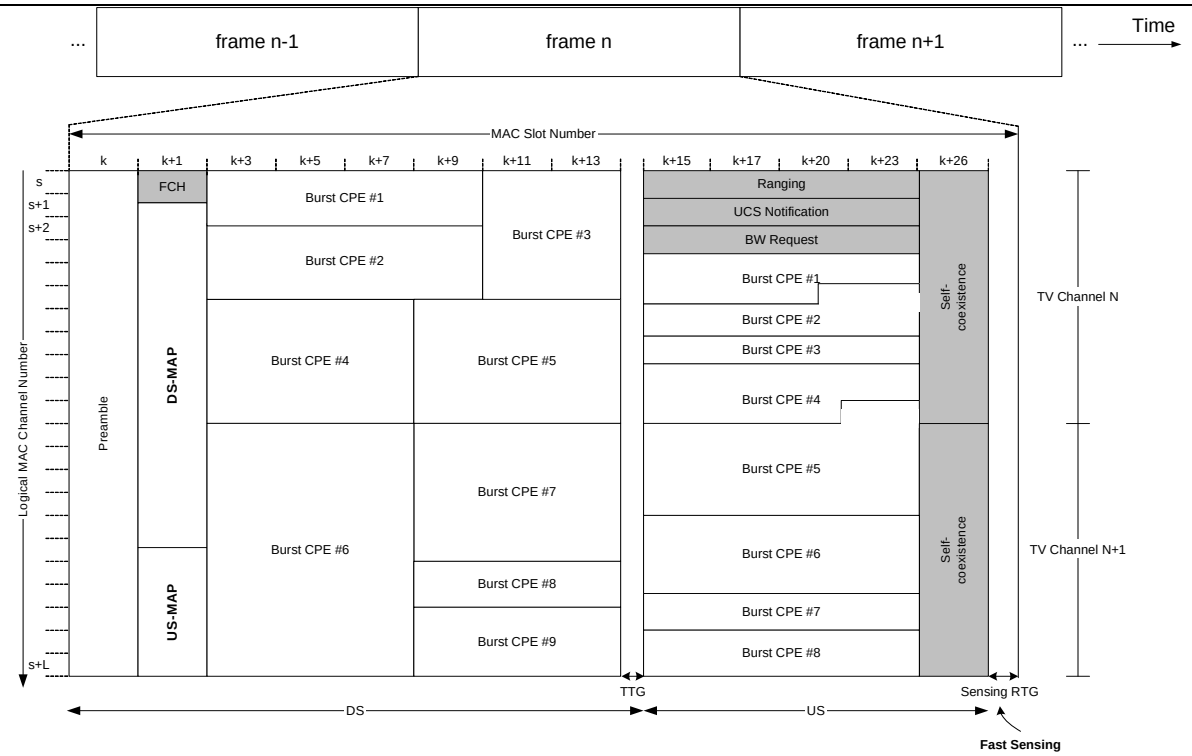

Figure 79 – Frame structure with fast sensing period

21.5.2.2 CPE Report

Once fast sensing stage is completed, the CPE needs to report the measurement results back to the BS. This can be done either after each fast sensing period or just once after all fast sensing periods within the Channel Detection Time. Regardless of that, the method the CPE uses to report such measurement results is critical.

Motivated by the fact that the outcome of energy detection measurements is a simple YES/NO answer, it does not require a detailed measurement report scheme. Therefore, we propose to use the Consolidated Spectrum Occupancy Measurement Report described in 8.22.3 for this purpose. With this report, the CPE can send to the BS in a single message the results of all of its fast sensing measurements done per Channel Detection Time period, and so results in an extremely bandwidth efficient scheme.

21.5.2.3 Fine Sensing Allocation

Once the BS receives the reports from enough CPEs about their fast sensing measurements, it can make a decision with respect to the following fine sensing stage(s). This is done through channel management messages and measurement requests that adaptively control the quiet periods. To this end, the mechanisms described in 8.21.7 and 8.22.1 are used.

21.6 Channel Management

A robust and efficient channel management component is a critical feature for any 802.22 MAC proposal. In fact, the channel management component incorporated in CMAC allows 802.22 systems to efficiently and dynamically use the available channels as the radio environment utilization changes. In CMAC, two modes of channel management are possible: embedded (see 8.1.1) and explicit (see 8.21).

The embedded mode of channel management has the advantage that individual channel management commands need not be sent (as it is the case in the explicit mode), and hence better spectrum utilization can be achieved. Another advantage of the embedded mode is that it addresses all CPEs in a cell, and hence is an effective way to take corrective actions in case an incumbent user starts operating in a channel occupied by all CPEs in an 802.22 cell.

The explicit approach to channel management, on the other hand, provides greater flexibility and is relatively independent of the MAC protocol used. Furthermore, this allows channel management to be implemented in different granularities, that is, these standalone messages could be sent by the BS to CPEs, for example, through unicast (i.e., destined to a single CPE), multicast (i.e., destined to a group of CPEs), or broadcast (i.e., all CPEs in a cell). These messages would also allow the BS to request confirmation of receipt, in case guaranteed delivery is required. In addition, contrary to the embedded mode where the BS has to wait for the next MAC frame in order to send a channel management command, this mode of operation of supporting individual channel management frames allows the MAC to rapidly react to changes in the radio spectrum occupation. This quick reaction is critical especially when we consider incumbent protection.

Irrespective of the channel management mode, CPEs and BSs shall treat channel management commands with high priority, especially when they refer to the protection of incumbent services. Once the BS detects or receives the report about the presence of an incumbent system in-band (see 21.1.1), it should send channel management commands to the effected receivers (e.g., this can be done through clustering). Depending of the urgency of the channel management command (the requirements of certain incumbent users may be more strict than others), the BS shall calculate the expected time the channel management message should arrive at the CPEs and appropriately set the scheduling fields (e.g., count, offset, duration, period) in the message header. This will dictate the urgency of the message, and how it shall be treated at the receivers.

Also, depending upon the situation, correct reception of channel management frames by all CPEs involved may be required by the BS. In this case, the BS shall set the Confirmation Needed field existing in the header of the channel management frame, which allows for the BS to specifically request the CPEs to send back a confirmation message. If we take the CHS-REQ message, Figure 80 depicts the message flow between BS and CPE when the Confirmation Needed field is set.

Once the destination receives the channel management command from the BS, it shall give higher priority for these types messages. As such, the receiver shall inspect the message control fields and proceed as instructed by the BS. If a confirmation is needed (in particular, this may be useful in case of transmissions of individual messages as these are unreliable), the receiver shall immediately send back a response message to the BS with the appropriate Confirmation Code. If a required acknowledgement message is not received within a pre-determined timeout, the BS may send another channel management message to the receiver in question. The receivers shall also check the scheduling fields in the management message in order to ascertain how urgent the message is, and how it should be treated internally. The receiver shall then change its operating parameters, at the scheduled time, as instructed by the BS.

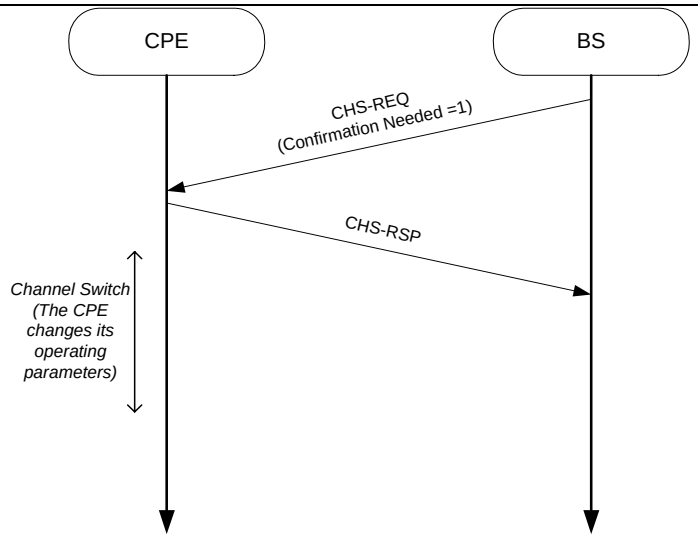

Figure 80 – Message flow between BS and CPE when confirmation is required

21.6.1 Channel Classification

For efficient channel management, four different channel sets are defined: active set 1, active set 2, candidate set, occupied set, and null set. These sets are respectively maintained in BS and CPE. If needed, each set is updated in every quiet period. Note that a set of channels that do not belong to either active set 1, active set 2, candidate set or occupied set is defined as a null set and it is maintained only the BS. Each channel set is defined as follows:

- Active set 1: a set of used channels for a certain CPE
- Active set 2: a set of used channels for a certain BS
- Candidate set: a set of clean channels available for a certain CPE or BS
- Occupied set: a set of occupied channels by incumbent user which a certain CPE finds
- Disallowed Set: a set of channels whose access are not allowed by regulation
- Null set: a set of channels that are not classified as one of above five sets

Note that the allowed set is defined by the union of the candidate set and the null set depending on a channel's SIR level. Also, the disallowed set is not considered in channel management scenarios because it does not change frequently. Additionally, in TDD case, Active set1 and Active set2 are not distinguishable. Therefore, Active set includes both Active set 1 and 2 in the case.

In order to maintain the channel sets, each BS maintains five channel sets (i.e., Active 1, Active 2, Occupied, Candidate, and Null). Also, each CPE maintains four channel sets (i.e., Active 1, Active 2, Candidate, and Occupied). These individual sets are updated after every quiet period either at a periodic interval or aperiodic interval.

21.6.2 Transition Diagram for Channel Sets

Any channel belongs to one of possible channel sets in the BS. At the end of quiet period, it may transit to other set as shown by the state transition diagram in Figure 81. Depending on the activities of incumbent users and channel quality, 8 possible transitions can be observed. The condition for each transition is described as follows:

- The channel in null set becomes useless as incumbent service appears.
- Incumbent service releases the channel and its quality is better than an existing member of the candidate set, then it is classified as a member of candidate set.

- Incumbent service releases the channel and its quality is worse than all member of the candidate set, then it is classified as a member of null set.
- If the channel quality is better than an existing member of the candidate set, then it replaces the member of candidate set.
- The channel becomes active by new connection of a WRAN service.
- The channel is switched as a member of null set as its quality is worse than newly selected candidate set from null (4), active(7) and occupied (2) sets.
- The channel is released due to the finish of its usage and its quality is better than an existing member of the candidate set.
- The channel is classified as a member of null set as the WRAN service releases the channel and its quality is worse than all member of the candidate set.

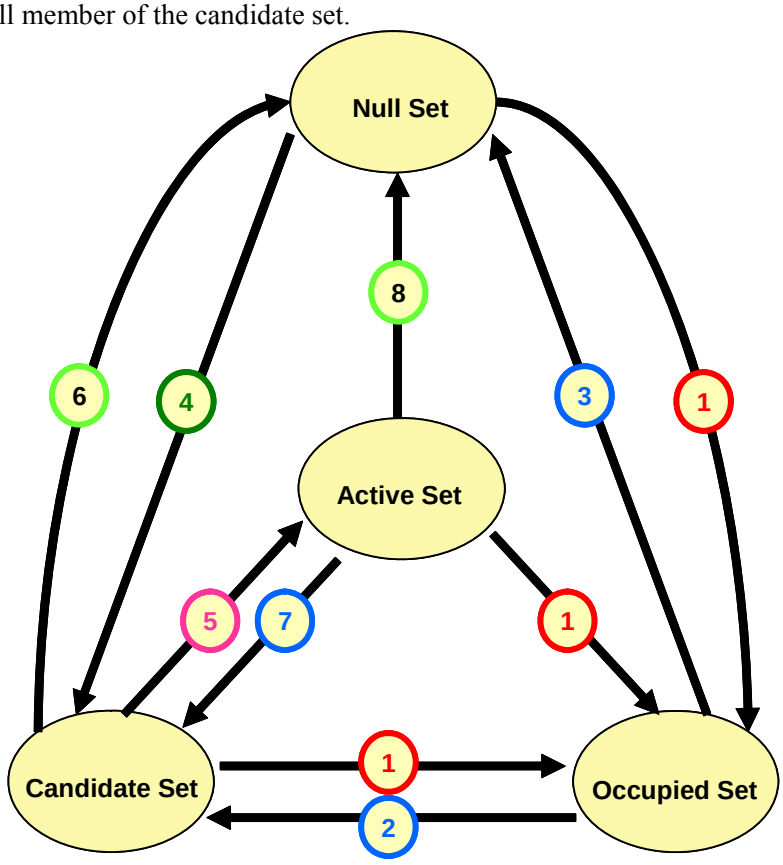

Figure 81 – Channel set transition diagram

21.6.3 Channel Switch Procedure

When incumbent users and other WRAN systems are detected in the current operating channel, the base station should select a channel CHselect from its candidate channel list (either randomly or based on some algorithm), randomly selects a wait time twait from a time window [tmin, tmax] and start a wait timer with twait as the expire time. Then, it advertises the channel selection (e.g., using a backhaul channel or WRAN air interface) before jumping to CHselect. Meanwhile, the WRAN system sense CHselect for incumbent signals and other WRAN systems. If the channel CHselect is still idle/available, it jumps to CHselect when the wait timer expires. However, if it is detected that the incument signals or other WRAN systems exist in CHselect, it goes back to the beginning to select another channel from the candidate channel list or its previously operated channel if it is not occupied by incumbents. If the collision occurs, the BS increases tmax and goes back to the beginning to select another channel from the candidate channel list or its previously operated channel if it is not occupied by the incumbents. For example, it may double tmax . The details of this procedure are shown in Figure 82.

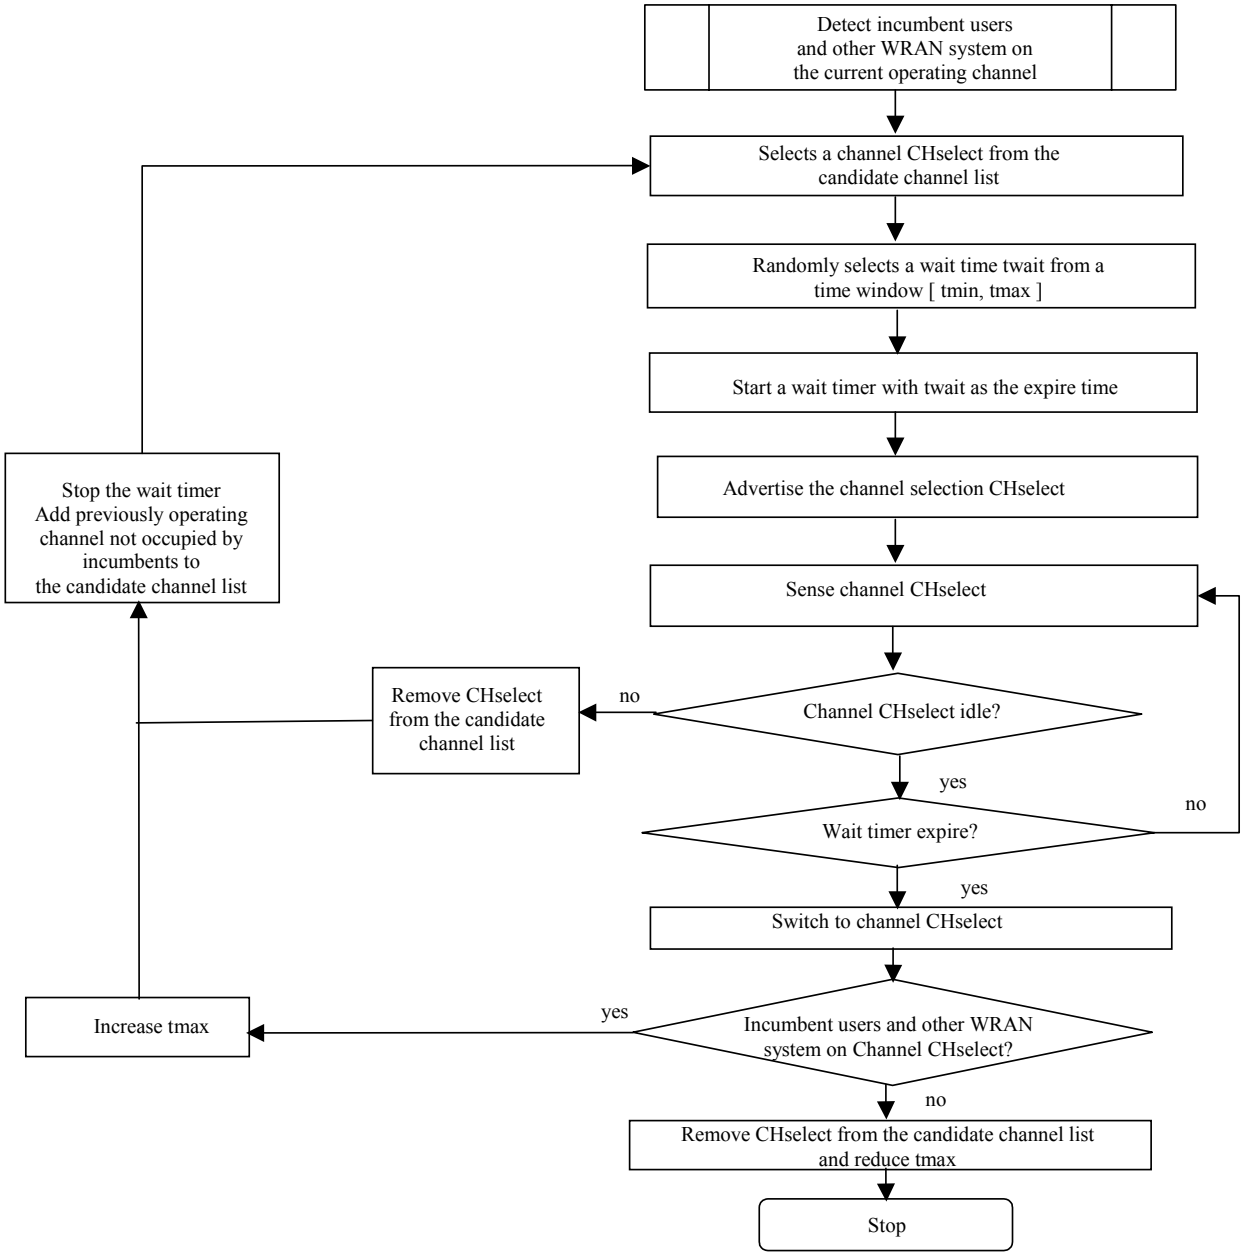

Figure 82 – The channel switch procedure

22. Security

The proposed CMAC security sublayer is in many respects inspired by the IEEE 802.16e/D12 draft [7], but it provides some simplifications in order to meet the 802.22 functional requirements.

22.1 Overview

The security sublayer provides subscribers with privacy, authentication or confidentiality across the broadband wireless network. It does this by applying cryptographic transforms to MPDUs carried across connections

between CPE and BS. In addition, the security sublayer provides operators with strong protection from theft of service. The BS protects against unauthorized access to these data transport services by securing the associated service flows across the network. The security sublayer employs an authenticated client/server key management protocol in which the BS, the server, controls distribution of keying material to client CPE.

Additionally, the basic security mechanisms are strengthened by adding digital-certificate-based CPE device-authentication to the key management protocol.

Security has two component protocols as follows:

1. An encapsulation protocol for securing packet data across the BWA network. This protocol defines:
    a. A set of supported cryptographic suites, i.e., pairings of data encryption and authentication algorithms,
    b. And the rules for applying those algorithms to a MAC PDU payload.
2. A key management protocol (called PKM, for Privacy Key Management protocol) providing the secure distribution of keying data from the BS to the CPE. Through this key management protocol, CPE and BS synchronize keying data; in addition, the BS uses the protocol to enforce conditional access to network services.

22.2 Authentication

The PKM protocol allows for mutual authentication: it is mostly based on the PKMv2 defined in the IEEE 802.16e/D12 draft. It also supports periodic re-authentication/reauthorization and key refresh. The key management protocol uses either EAP [8], or X.509 digital certificates [9] together with RSA public-key encryption algorithm [10] or a sequence starting with RSA or EAP authentication and followed by another EAP authentication: this third option can be used to facilitate device authentication (via RSA) followed by user authentication (based on EAP). It uses strong encryption algorithms to perform key exchanges between a CPE and BS.

The PKM authentication protocol establishes a shared secret (called an Authorization Key (AK)) between the CPE and the BS. The shared secret is then used to secure subsequent PKM exchanges of TEKs (Traffic Encryption Key). This two-tiered mechanism for key distribution permits refreshing of TEKs without incurring the overhead of computation-intensive operations.

A BS authenticates a client CPE during the initial authorization exchange. Each CPE presents its credentials, which will be a unique X.509 digital certificate issued by the CPE s manufacturer (in the case of RSA authentication) or an operator-specified credential (in the case of EAP-based authentication).
Since the BS authenticates the CPE, it may protect against an attacker employing a cloned CPE, masquerading as a legitimate subscriber's CPE.

The traffic-key management portion of the PKM protocol adheres to a client/server model, where the CPE (a PKM client) requests keying material, and the BS (a PKM server) responds to those requests, ensuring that individual CPE clients receive only keying material for which they are authorized.

The PKM protocol uses MAC management messaging, i.e., PKM-REQ and PKM-RSP messages defined in section 6.3.2.3.9 of IEEE 802.16e/D12 draft.

22.2.1 PKM RSA Authentication

The PKM RSA authentication protocol uses X.509 digital certificates and the RSA public-key encryption algorithm that binds public RSA encryption keys to MAC addresses of CPEs.

A BS authenticates a client CPE during the initial authorization exchange. Each CPE carries a unique X.509 digital certificate issued by the CPEs manufacturer. The digital certificate contains the CPEs Public Key and CPE MAC address. When requesting an AK, a CPE presents its digital certificate to the BS. The BS verifies the digital certificate, and then uses the verified Public Key to encrypt an AK, which the BS then sends back to the requesting CPE.

All CPEs using RSA authentication shall have factory-installed RSA private/public key pairs or provide an internal algorithm to generate such key pairs dynamically. If a CPE relies on an internal algorithm to generate its RSA key pair, the CPE shall generate the key pair prior to its first AK exchange, described in section 7.2.1 of the IEEE 802.16e/D12 draft. All CPEs with factory-installed RSA key pairs shall also have factory-installed X.509 certificates. All CPEs that rely on internal algorithms to generate an RSA key pair shall support a mechanism for installing a manufacturer-issued X.509 certificate following key generation.

22.2.2 PKM EAP Authentication

PKM EAP Authentication uses Extensible Authentication Protocol in conjunction with an operator-selected EAP Method (e.g. EAP-TLS [11]). The EAP method will us a particular kind of credential — such as an X.509 certificate in the case of EAP-TLS, or a Subscriber Identity Module in the case of EAP-SIM [12].

The particular credentials and EAP methods that are to be used are outside of the scope of this specification. However, the EAP method selected should fulfil the mandatory criteria listed in section 2.2 of RFC 4017 [13]. Use of an EAP method not meeting these criteria may lead to security vulnerabilities.

During re-authentication, the EAP transfer messages are protected with an HMAC[14]/CMAC[15] tuple. The BS and CPE must discard unprotected EAP transfer messages or EAP transfer messages with invalid HMAC/CMAC digests during re-authentication.

22.3 Authorization

The authorization process can be considered as a part of the previously defined authentication process: it is indeed the process of the BS associating a CPE's authenticated identity to a paying subscriber, and hence to the data services that subscriber is authorized to access. Thus during network entry, with the AK exchange, the BS determines the authenticated identity of a client CPE and the services (i.e., specific TEKs) the CPE is authorized to access.

These services are also known under the concept of service flows (already defined in section 20 of this document). Service flows can either be permanently provisioned on the BS, or dynamically created during connection.

The authorization process is controlled by the Authorization state machine, which consists of six states and eight distinct events (including receipt of messages) that can trigger state transitions. The Authorization finite state machine (FSM) can be found in section 7.2.4 of the IEEE 802.16-2004 standard.

22.4 Privacy

22.4.1 Data (Payload) Encryption

Encryption services are defined as a set of capabilities within the MAC Security Sublayer. MAC Header information specific to encryption is allocated in the generic MAC header format. Encryption is applied to the MAC PDU payload when required by the selected cipher suite; the generic MAC header is not encrypted.

Different cryptographic algorithms and key sizes will be used by the PKM protocol as methods of packet data encryption:

- The CBC mode of the US Data Encryption Standard (DES) algorithm [16]. Note that this method is supported by the IEEE 802.16e/D12 specification, but might not be selected for IEEE 802.22.
- The CCM mode of the US Advanced Encryption Standard (AES) algorithm [17].
- The CBC mode of the US Advanced Encryption Standard (AES) algorithm [18].

22.4.2 Protection of Network Control Information

MAC (message authentication code) keys are used to sign management messages in order to validate the authenticity of these messages and protect their integrity. All MAC management messages shall be sent in the clear to facilitate registration, ranging, and normal operation of the MAC.

The MAC to be used is negotiated at CPE Basic Capabilities negotiation. Two different message authentication codes can be used:

- A BS or CPE may support management message integrity protection based on Cipher-based MAC – together with the AES block cipher. The CMAC construction as specified in [15] shall be used.
- Or they may support a management message integrity protection based on HMAC with the secure hash algorithm SHA-1 [19].

There is a different key for US and DS messages. Also, a different message authentication key is generated for a multicast message (this is DS direction only) and for a unicast message. In general, the message authentication keys used to generate the CMAC value and the HMAC-Digest are derived from the AK.

22.5 Protection Against Deny of Service and Other Attacks

One of the great advantage to base the security sublayer of IEEE 802.22 on the one defined in IEEE 802.16e/D12 is that this last one has been deeply studied and corrected by various security experts, including members of the IETF and IEEE security groups. This sublayer can thus be considered as well secured.

A typical and frequent kind of deny of service attack consists in forging and sending legitimate management frames. Here, all critical management packets are digitally signed, and their integrity is checked by the receiver before further use: there is thus no mean for an attacker to craft such a packet.

Yet it may be possible to resend one of them, as long as the key used to sign it is still active: this case has also been taken into account, since all the critical negotiation phases (authentication or key exchange) are protected against replay. This has been done by including random numbers chosen by each peer, and included as challenge/response in the different packets used for the negotiation. Moreover, these negotiations are done via 3-way handshakes, meaning that the peer requesting the operation must send a confirmation upon reception of the response. This prevents an attacker from creating a deny of service by resending a legitimate request: since he won't be able to send the confirmation (which must include the random number added by the peer making the response), none of the legitimate peers will take into account the attack attempt.

Another common attack consists in creating "noise" by replaying formerly captured data packets, or crafting false ones. If AES in CCM mode is chosen to encrypt MAC PDUs, there is no mean for an attacker to replay formerly captured data packets: a window has indeed to be used for PN (packet number) values in order to validate the freshness and uniqueness of the packet.

One fear that comes with EAP use is the fact that some critical EAP packets are generally non-signed: this makes it possible for an attacker to forge them, and create a denial of service by sending them judiciously. Here this can

be avoided since EAP packets can be protected if a first authentication has already been performed (this is typically the case when devices authenticate themselves first with RSA or EAP, and then the subscriber can authenticate himself with the network via EAP): a CMAC or a HMAC Digest can be included to allow the binding of previous authorization and following EAP authentication by authenticating the EAP payload. The key used to calculate these message authentication digests is derived during the first authentication process.

Another fear that comes with successive EAP authentication procedures is the necessity to cryptographically bind previous EAP authentication and following EAP authentication session, while protecting second EAP messages. In order to prevent man-in-the-middle attacks, the first and second EAP methods should fulfil the mandatory criteria listed in section 2.2 of RFC 4017, such as EAP-AKA [20].

23. Parameter Configuration

23.1 General

Table 226 – Global parameter setting

Entity	Name	Time reference	Minimum value	Default value	Maximum value
BS	DCD Interval	Time between transmission of DCD messages			10 s
BS	UCD Interval	Time between transmission of UCD messages			10 s
BS	UCD Transition	The time the BS shall wait after repeating a UCD message with an incremented Configuration Change Count before issuing a US-MAP message referring to Upstream_Burst_Profiles defined in that UCD message	2 MAC frames		
BS	DCD Transition	The time the BS shall wait after repeating a DCD message with an incremented Configuration Change Count before issuing a DS-MAP message referring to Downstream_Burst_Profiles defined in that DCD message	2 MAC frames		
BS	Per Channel Transmission	If multiple channels are in use, this specifies the maximum time before a BS must transmit control messages per channel. This would allow the support of single channel CPEs (i.e., those that cannot support multiple channel operation).		300ms	
BS	Max MAP Pending	Maximum validity of map		End of next frame	
BS	Initial Ranging Interval	Time between Initial Ranging regions assigned by the BS			2 s
BS	CLK-CMP Interval	Time between the clock compare measurements used for the generation of CLK-CMP messages.	50 ms	50 ms	50 ms
CPE	Lost DS-MAP Interval	Time since last received DS-MAP message before downstream			600 ms

		synchronization is considered lost			
CPE	Lost US-MAP Interval	Time since last received US-MAP message before upstream synchronization is considered lost			600 ms
CPE	Lost SCH	Number of SCH that can be lost until synchronization is considered lost			15
CPE	Contention Ranging Retries	Number of retries on contention Ranging Requests	16		
CPE, BS	Invited Ranging Retries	Number of retries on inviting Ranging Requests	16		
CPE	Request Retries	Number of retries on bandwidth allocation requests	16		
CPE	Registration Request Retries	Number of retries on registration requests	3		
BS	T_{proc}	Time provided between arrival of the last bit of a US-MAP at an CPE and effectiveness of that map	5 symbols		
BS	CPE Ranging Response Processing Time	Time allowed for an CPE following receipt of a ranging response before it is expected to reply to an invited ranging request	10 ms		
CPE, BS	DSx Request Retries	Number of Timeout Retries on DSA/DSC/DSD Requests		3	
CPE, BS	DSx Response Retries	Number of Timeout Retries on DSA/DSC/DSD Responses		3	
CPE	T1	Wait for DCD timeout			5 * DCD interval maximum value
CPE	T2	Wait for broadcast ranging timeout			5 * ranging interval
CPE	T3	Ranging Response reception timeout following the transmission of a Ranging Request		200 ms	200 ms
CPE	T4	Wait for unicast ranging opportunity. If the pending-until-complete field was used earlier by this CPE, then the value of that field shall be added to this interval.	30 s		35 s
BS	T5	Wait for Upstream Channel Change response			2 s
CPE	T6	Wait for registration response			3 s
CPE, BS	T7	Wait for DSA/DSC/DSD Response timeout			1 s
CPE, BS	T8	Wait for DSA/DSC Acknowledge timeout			300 ms
BS	T9	Registration Timeout, the time allowed between the BS sending a RNG-RSP (success) to an CPE, and receiving a CBC-REQ from that same CPE	300 ms	300 ms	
CPE, BS	T10	Wait for Transaction End timeout			3 s
CPE	T12	Wait for UCD descriptor			5 * UCD Interval maximum value
BS	T13	The time allowed for an CPE,	15 min	15 min	

		following receipt of a REG-RSP message to send a TFTP-CPLT message to the BS			
CPE	T14	Wait for DSX-RVD Timeout			200 ms
BS	T15	Wait for MCA-RSP	20 ms	20 ms	
CPE	T16	Wait for bandwidth request grant	10 ms		Service QoS dependent
BS	T17	Time allowed for CPE to complete CPE Authorization and Key Exchange	5 min	5 min	
CPE	T18	Wait for CBC-RSP timeout		5 ms	<< T9
CPE	T19	Time DS-channel remains unusable			
CPE	T20	Time the CPE searches for preambles on a given channel	2 MAC frames		
CPE	T21	Time the CPE searches for DS-MAP on a given channel			10 s
CPE, BS	T22	Wait for ARQ-Reset			0.5 s
BS	T27 as idle timer	Maximum time between unicast grants to CPE when BS believes CPE upstream transmission quality is *good enough*	CPE Ranging Response Processing Time		
BS	T27 as active timer	Maximum time between unicast grants to CPE when BS believes CPE upstream transmission quality is *not good enough*	CPE Ranging Response Processing Time		
CPE	CBC Request Retries	Number of retries on CBC Request	3	3	16
CPE	CPE downstream management message processing time	Maximum time between reception of Fast Power Control management message and compliance to its instructions by CPE			200 s
CPE	T28	Wait for BLM-ACK timeout	10 ms		300 ms
CPE	BLM-REP Retries	Number of Timeout Retries on BLM-REP messages		3	

23.2 Well-Known CIDs

In CMAC, CIDs are reserved for specific purposes and shall not be used outside their pre-determined scope. Table 227 provides the allocation of CIDs and their corresponding use.

Table 227 – CID Allocation (m = maximum number of supported CPEs)

CID	Value	Description
Initial Ranging CID	0x0000	Used by CPE and BS during initial ranging process.
Basic CID	0x0001 – m	The same value is assigned to both the DS and US connection.
Primary Management CID	m+1 – 2m	The same value is assigned to both the DS and US connection.
Multicast Management CID	2m+1 – 3m	Shall only be used in the DS direction by the BS. Allows management commands to be addressed to a set of CPEs.
Transport CIDs and Secondary Mgmt CIDs	3m+1 – 0xFEFE	For the secondary management connection, the same value is assigned to both the DS and US connection.
Multicast Polling CIDs	0xFF00 – 0xFFFD	A CPE may be included in one or more multicast polling groups for the purposes of obtaining bandwidth via polling. These connections have no associated service flow.
Padding CID	0xFFFE	Used for transmission of padding information by CPE and BS.
Broadcast CID	0xFFFF	Used for broadcast information that is transmitted on a downstream to all CPE.

24. Definitions

For the purposes of this proposal, the following terms and definitions apply. *The Authoritative Dictionary of IEEE Standards Terms*, Seventh Edition [1], should be referenced for terms not defined in this clause.

Base station (BS): A generalized equipment set providing connectivity, management, and control of the CPE.

Cell: A 802.22 cell (or simply, a cell) is defined as formed by a single 802.22 BS and zero or more 802.22 CPEs associated with and under control by this 802.22 BS, whose coverage area extends up to the point where the transmitted signal from the 802.22 BS can be received by associated 802.22 CPEs with a given minimum SINR quality.

Channel: This is the logical channel number used by the MAC, and does not necessarily reflect a physical channel. This can be mapped into any representation as defined by the PHY. *See also:* **TV channel**.

Consumer premise equipment (CPE): A generalized equipment set providing connectivity between subscriber equipment and a BS.

Downstream: The direction from a BS to the CPE.

Downstream channel descriptor (DCD): A MAC message that describes the PHY characteristics of a downstream channel.

Downstream interval usage code (DIUC): An interval usage code specific to a downstream. *See also*: **interval usage code**.

Downstream map (DS-MAP): A MAC message that defines burst start times for both time division multiplex and time division multiple access (TDMA) by a CPE on the downstream.

Interval usage code: A code identifying a particular burst profile that can be used by a downstream or upstream transmission interval.

Security association (SA): The set of security information a base station (BS) and one or more of its client subscriber stations (SSs) share in order to support secure communications. This shared information includes traffic encryption keys (TEKs) and cipher block chaining (CBC) initialization vectors.

Security association identifier (SAID): An identifier shared between the base station (BS) and subscriber station that uniquely identifies a security association (SA).

Self-Coexistence: Coexistence between wireless systems of the same type. In the case of 802.22, this means coexistence of multiple overlapping 802.22 cells.

Time division duplex (TDD): A duplex scheme where upstream and downstream transmissions occur at different times but may share the same frequency.

Time division multiple access (TDMA) burst: A contiguous portion of the upstream or downstream using PHY parameters, determined by the Downstream Interval Usage Code (DIUC) or Upstream Interval Usage Code (UIUC), that remain constant for the duration of the burst. TDMA bursts are separated by preambles and are separated by gaps in transmission if subsequent bursts are from different transmitters.

Time division multiplexing (TDM) burst: A contiguous portion of a TDM data stream using PHY parameters, determined by the Downstream Interval Usage Code (DIUC), that remain constant for the duration of the burst. TDM bursts are not separated by gaps or preambles.

Transport connection: A connection used to transport user data.

TV channel: Refers to a specific physical TV Channel as defined by TV broadcast communication standards.

Upstream: The direction from a CPE to the BS.

Upstream channel descriptor (UCD): A medium access control message that describes the PHY characteristics of an upstream.

Upstream interval usage code (UIUC): An interval usage code specific to an upstream.

Upstream map (US-MAP): A set of information that defines the entire access for a scheduling interval.

25. Abbreviations and Acronyms

AAS	Adaptive Antenna System
AES	Advanced Encryption Protocol
AK	Authorization Key
ARQ	Automatic Repeat Request
BCH	Burst Control Header
BE	Best Effort
BS	Base Station
BTC	Block Turbo Code
CBC-MAC	Cipher Block Chaining Message Authentication Code
CBP	Coexistence Beacon Protocol
CCM	CTR mode with CBC-MAC
CDMA	Code Division Multiple Access
CID	Connection Identifier
CINR	Carrier to Interference and Noise Ratio
CMAC	Cognitive MAC
CMAC	Cipher-based Message Authentication Code
CoS	Class of Service
CPE	Consumer Premise Equipment
CRC	Cyclic Redundancy Check
CS	Convergence Sublayer
CTR	Counter Mode Encryption
DAMA	Demand Assigned Multiple Access
DFS	Dynamic Frequency Selection
DHCP	Dynamic Host Configuration Protocol
DIUC	Downstream Interval Usage Code
DS	Downstream
EAP	Extensible Authentication Protocol
FBWA	Fixed Broadband Wireless Access
FCH	Frame Control Header
FDD	Frequency Division Duplex
ID	Identification
IDRP	Incumbent Detection Recovery Protocol
IE	Information Element
IETF	Internet Engineering Task Force
IP	Internet Protocol
IUC	Interval Usage Code
LAN	Local Area Network
LC	Logical Cluster
LLC	Link Layer Control
MAC	Medium Access Control Layer
MAC	Message Authentication Code
MIC	Message Integrity Check
MPEG	Moving Pictures Experts Group
nrtPS	Non-real-time Polling Service
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
PAPR	Peak-to-Average Power Ratio
PC	Physical Cluster
PDU	Protocol Data Unit
PER	Packet Error Rate

PHY	Physical Layer
PKM	Privacy Key Management protocol
PMP	Point to Multipoint
PN	Packet Number
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
QPSK	Quadrature Phase-Shift Keying
RS	Reed-Solomon
RSS	Received Signal Strength
RTG	Receive/Transmit Transition Gap
rtPS	Real-time Polling Service
SA	Security Association
SAID	Security Association Identifier
SAP	Service Access Point
SCH	Superframe Control Header
SDU	Service Data Unit
SFID	Service Flow ID
SINR	Signal to Interference and Noise Ratio
SM	Spectrum Manager
SNMP	Simple Network Management Protocol
SNR	Signal to Noise Ratio
SSS	Sliding Self-coexistence Slots
SUB	Spectrum Usage Bitmap
TAG	Technical Advisory Group
TCP	Transmission Control Protocol
TDD	Time Division Duplex
TDM	Time Division Multiplex
TDMA	Time Division Multiple Access
TEK	Traffic Encryption Key
TFTP	Trivial File Transfer Protocol
TPC	Transmit Power Control
TTG	Transmit/Receive Transition Gap
UCS	Urgent Coexistence Situation
UGS	Unsolicited Grant Service
UIUC	Upstream Interval Usage Code
US	Upstream
WG	Working Group
WMB	Wireless Microphone Beacon
WRAN	Wireless Regional Area Network

26. References

[1] IEEE 100™, *The Authoritative Dictionary of IEEE Standards Terms*, Seventh Edition.

[2] IEEE 802.22 Working Group on Wireless Regional Area Networks, http://www.ieee802.org/22/.

[3] IEEE 802.22 Working Group on Wireless Regional Area Networks, "IEEE 802.22 Functional Requirements," September 2005.

[4] IEEE 802.22 Working Group on Wireless Regional Area Networks, "WRAN Reference Model," doc.: IEEE 802.22-04/0002r12.

[5] IEEE Std 802.16™ Specification, IEEE Standard for Local and Metropolitan Area Networks, Part 16: Air Interface for Fixed Broadband Wireless Access Systems, 1 October 2004.

[6] M. McHenry, "Report on Spectrum Occupancy Measurements," *Shared Spectrum Company*, http://www.sharedspectrum.com/?section=nsf_summary.

[7] IEEE P802.16e/D12, Draft IEEE Standard for Local and metropolitan area networks, Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems, Amendment for Physical and Medium Access Control Layers for Combined Fixed and Mobile Operation in Licensed Bands, 14/10/2005.

[8] IETF RFC 3748, Extensible Authentication Protocol, June 2004.

[9] IETF RFC 3280, Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile, April 2002.

[10] PKCS #1: RSA Cryptography Standard, June 2002.

[11] IETF RFC 2716, PPP EAP TLS Authentication Protocol, October 1999.

[12] Draft in progress, Extensible Authentication Protocol Method for GSM Subscriber Identity Modules (EAP-SIM), December 2004 (draft-haverinen-pppext-eap-sim-16.txt).

[13] IETF RFC 4017, Extensible Authentication Protocol (EAP) Method Requirements for Wireless LANs, March 2005.

[14] IETF RFC 2104, "HMAC: Keyed-Hashing for Message Authentication," H. Krawczyk, M. Bellare, R. Canetti, February 1997.

[15] NIST Special Publication 800-38B - Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication, May 2005.

[16] FIPS 46-3, FIPS 74, FIPS 81, DES Modes of Operation, December 1980.

[17] NIST Special Publication 800-38C, FIPS-197, Advanced Encryption Standard (AES)

[18] NIST Special Publication 800-38A, FIPS-197, Advanced Encryption Standard (AES)

[19] FIPS 180-1, Secure Hash Standard (SHS), April 1995.

[20] Draft in progress, Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA), draft-arkko-pppext-eap-aka-15.txt

[21] Niels Hoven and Anant Sahai, "Power Scaling for Cognitive Radio," in proceedings of IEEE WirelessCom 05 Symposium on Emerging Networks, Technologies and Standards, Hawaii, USA, June 13-16, 2005.

27. Appendix A – Power control for the protection of incumbents

27.1 Multiple CPEs Joint TPC

A joint power constraint among CPEs located in a neighbouring area to avoid creating interference above the interference threshold into TV receivers during simultaneous access to the channel in the same frequency band might be required, assuming each CPE transmits power over the whole TV band (so for example this applies only in OFDMA when CPEs share the same subcarriers or sub-channels on the uplink). In particular, this scenario applies to sharing of the spectrum by WRAN cells or operators.

The maximum transmit power for each CPE must be controlled by the base station in order to protect incumbents in the designated areas where TV receivers must be protected. This maximum transmit power constraint is independent of the adaptation performed by the WRAN base station to meet the service requirements of its customers. This power constraint comes prior to any other control over the transmit power of the CPEs. The envisioned module is meant to protect incumbent services in the protected areas in the case of simultaneous transmissions of multiple CPEs located in an area close to the boundary of the protected contour of the incumbent. We will designate this area as a constraint area. For the time being, we will only concern ourselves with TV operations. However the same principles could be applied for other incumbents. The simultaneous CPE transmissions occur in the following scenarios:

- Uplink sharing of space-time-frequency resource
- WRAN cells sharing of a space-time-frequency resource (same WRAN or coexistence with another WRAN)
- Contention-based access to a WRAN base station.
- Coexistence with other license-exempt users whose transmissions are modelled as another CPE transmitting simultaneously with one CPE of the WRAN.

Note that the joint transmit power constraint rule could also apply as a special and simpler case to the transmit power of two or more base stations overlapping over some coverage area and operating on the same TV band, in which case they would need to transmit simultaneously.

An example of a scenario of operation close to a NSTC TV operation is shown in Figure 83. Notice that a CPE located closer than 4.7 km away from the Grade B contour could operate at less than 4 W EIRP and still meet the interference D/U constraints at the Grade B contour. However we will not allow that CPE to transmit if it is closer than d_{min} to the Grade B contour, mainly because it would require such a low power that it would not be possible to maintain acceptable communication with the WRAN base station.

The application of the CPE power constraint for the protection of incumbent comes before radio resource management optimization. It basically determines the capacity of the channel of each CPE. Radio resource management might decide to allocate even less power to a CPE or leave it at its maximum transmit power, and adapt modulation and coding according to channel quality indicator of short-term channel conditions.

The concept of constraint area is illustrated in Figure 83. These boundaries are explained as follows by taking the example co-channel operation to an NTSC TV station outside the Grade B contour. Table 223 and Table 224 list the maximum transmit power constraint in the case of the operation of a single CPE at full transmit power of 4W over a 6 MHz band. That CPE could still meet the interference requirements into NTSC TV receivers located at the edge of the Grade B contour by transmitting with less power or less bandwidth if it operates closer to d_o km away from the contour. However, in order to meet the WRAN services requirements, there will be a minimum distance of d_{min} km that guarantees that the CPE transmit power is large enough. In the case of simultaneous transmission of CPEs in the same time-frequency-space resource, their radiated powers will add at the nearest TV receiver. Therefore, even outside the d_o km radius, if more than one CPE transmits at full power over 6 MHz, the interference requirements into TV receivers will not be met at the Grade B contour. In this case, a joint transmit power rule is needed for CPEs that belong to the same constraint area. This constraint area would be limited to a few km^2.

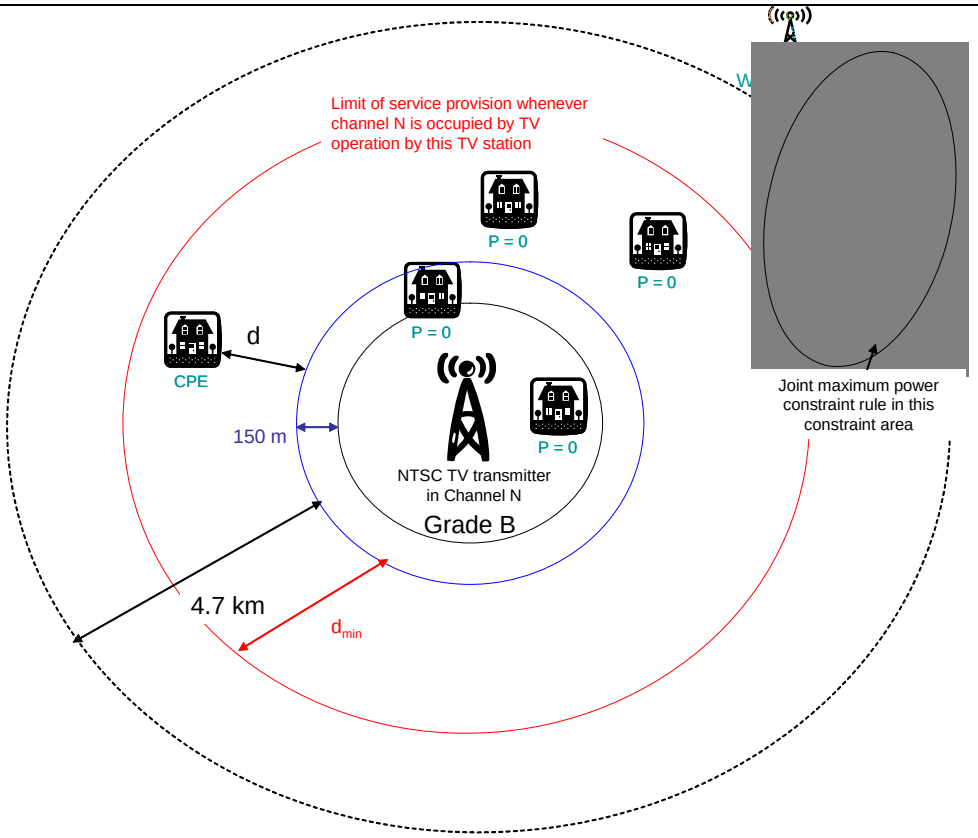

Figure 83 – Illustration of transmit power boundary constraints and constraint area in the case of NTSC TV co-channel operation

In the case of a single CPE transmitting at full power P_t over a 6 MHz bandwidth at d_k km away from the Grade B contour, co-channel to a NTSC TV operation, the signal radiated by that CPE number k is received by a TV receiver at the Grade B contour with power:

$$P_{r,\max} = P_{t,k,\max}(d_k) d_k^{-a}.$$

d_k is the distance of the CPE to the Grade B contour, and a is the path loss exponent. Fading is not included in this model, and could only decrease the power received at the TV receiver. Thus the actual path loss exponent can be chosen after analysis of the fading statistics to meet some performance target with some outage probability.

Let n be the density of CPEs in the constraint area A. These are the CPEs that share the same time-frequency-space resource. In the case of coexistence with other license-exempt users, these users should be included in that density. Then the power received from multiple transmitting CPEs is given by:

$$P_r = \sum_{CPE\ \#\ k \in constraint\ area} P_{t,k} d_k^{-a}$$

It is known that a sea of CPEs transmitting at maximum power results in a decrease of effective path loss by 2. As a consequence, the interference requirements at the Grade B contour cannot be met anymore.

We proposed the power rule for CPE number k as:

$$P_{t,k} = P_t\left(n, d_k, P_{t,k,\max}(d_k)\right).$$

The expression of the individual power constraint for an individual CPE is:

$$P_{t,k,\max}(d_k) = P_o \left(\frac{d_o}{d_k} \right)^{a}$$

d_o is 4.7 km (NTSC) or 10 km (DTV) and $P_o = 4$ W .

The total power received from all CPEs at the nearest TV receiver is:

$$P_r = \sum_k P_t\left(n, d_k, P_{t,k,\max}(d_k)\right) d_k^{-a}$$

The power rule should ensure that:

- $P_r \le P_{r,\max}$
- $P_{t,k} \le P_{t,k,\max}(d_k)$
- $\lim_{n \to 1/A} P_{t,k} = P_{t,k,\max}(d_k)$
- $\lim_{n \to n_{\max}} P_{t,k} = 0$
- With the goal that $P_r \approx P_{r,\max}$ for all $n \in [1/A, n_{\max}]$ in order to provide CPEs with enough transmit power to meet the WRAN services requirements.

Rule #1: $P_{t,k} = \frac{P_o}{K}$, K is the number of CPEs in the constraint area.

The naïve power constraint rule #1 that consists to divide the individual power constraint of each CPE by the total number of CPEs in a constraint area is too stringent with respect to meeting the interference requirements at the nearest TV receiver, unless the CPEs are all closely located at 4.7 km from the noise-protected contour. As the number of CPEs increases, even if their average distance from each other is very large, they would have to unnecessarily decrease their power, thus rendering any communication with the base station impossible. This rule is only worthwhile for a small number of CPEs equally distant from the noise-protected contour.

Rule #2: $P_{t,k} = \frac{P_o d_o^{\,a}}{\sum_{CPEs} d_k^{\,a}}$.

This power rule exactly meets the interference requirements. This expression assumes that the distances are taken with respect to the same TV receiver, thus it is pessimistic for the WRAN since for a given TV receiver, one CPE will be at the said distance, but other CPEs will be farther away, since that TV receiver will not be the closest point in the noise-protected contour for all these CPEs simultaneously. This power control will therefore meet the interference requirements with a margin.

In the case of the coexistence of multiple WRANs in a given constraint area that has been determined jointly, each WRAN only needs to compute a weight based on its knowledge of the distance of its own CPEs to the Grade B/noise-protected contour:

$$w_i = \sum d_k^{\,a}$$

Then WRANs exchange this information so they can scale the maximum individual power constraint of their CPEs as:

$$P_{t,k} = \frac{P_o d_o^{\,a}}{\sum w_i}$$

Rule #2 treats every CPE equally, thus the CPEs that are located closer to the noise-protected contour will have to reduce their transmit power as much as the CPEs that are located closer to the WRAN base station, thus they

might suffer from loss of service. Further considerations would involve radio resource management related to Quality of Service constraints, which would be an implementation issue. Therefore, the purpose of this appendix is to provide guidelines for more advanced CPE maximum transmit power constraint features for the IEEE 802.22 standard.