ETSI TS 103 311-2 V1.1.1 (2017-02)

TECHNICAL SPECIFICATION

Integrated broadband cable telecommunication networks (CABLE); Fourth Generation Transmission Systems for Interactive Cable Television Services - IP Cable Modems; Part 2: Physical Layer; DOCSIS® 3.1

Reference

DTS/CABLE-00017-2

Keywords

access, broadband, cable, data, DOCSIS, IP, IPCable, modem

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Important notice

The present document can be downloaded from:
http://www.etsi.org/standards-search

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the only prevailing document is the print of the Portable Document Format (PDF) version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at
https://portal.etsi.org/TB/ETSIDeliverableStatus.aspx

If you find errors in the present document, please send your comment to one of the following services:
https://portal.etsi.org/People/CommiteeSupportStaff.aspx

Copyright Notification

No part may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm except as authorized by written permission of ETSI.
The content of the PDF version shall not be modified without the written authorization of ETSI.
The copyright and the foregoing restriction extend to reproduction in all media.

© European Telecommunications Standards Institute 2017.
All rights reserved.

DECT™, **PLUGTESTS**™, **UMTS**™ and the ETSI logo are Trade Marks of ETSI registered for the benefit of its Members.
3GPP™ and **LTE**™ are Trade Marks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners.
GSM® and the GSM logo are Trade Marks registered and owned by the GSM Association.

Contents

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (https://ipr.etsi.org/).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Foreword

This Technical Specification (TS) has been produced by ETSI Technical Committee Integrated broadband cable telecommunication networks (CABLE).

The present document is part 2 of a multi-part deliverable. Full details of the entire series can be found in part 1 [21].

DOCSIS® is a registered Trade Mark of Cable Television Laboratories, Inc., and is used in the present document with permission.

Modal verbs terminology

In the present document "**shall**", "**shall not**", "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the ETSI Drafting Rules (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

1 Scope

The present document is part of a series of specifications that defines the fourth generation of high-speed data-over-cable systems, commonly referred to as the DOCSIS 3.1 specifications. The present document was developed for the benefit of the cable industry, and includes contributions by operators and vendors from North and South America, Europe and Asia.

This generation of the DOCSIS specifications builds upon the previous generations of DOCSIS specifications (commonly referred to as the DOCSIS 3.0 and earlier specifications), leveraging the existing Media Access Control (MAC) and Physical (PHY) layers, but with the addition of a new PHY layer designed to improve spectral efficiency and provide better scaling for larger bandwidths (and appropriate updates to the MAC and management layers to support the new PHY layer). It includes backward compatibility for the existing PHY layers in order to enable a seamless migration to the new technology.

There are differences in the cable spectrum planning practices adopted for different networks in the world. For the new PHY layer defined in the present document, there is flexibility to deploy the technology in any spectrum plan; therefore, no special accommodation for different regions of the world is required for this new PHY layer.

However, due to the inclusion of the DOCSIS 3.0 PHY layers for backward compatibility purposes, there is still a need for different region-specific physical layer technologies. Therefore, three options for physical layer technologies are included in the present document, which have equal priority and are not required to be interoperable. One technology option is based on the downstream channel identification plan that is deployed in North America using 6 MHz spacing. The second technology option is based on the corresponding European multi-program television distribution. The third technology option is based on the corresponding Chinese multi-program television distribution. All three options have the same status, notwithstanding that the document structure does not reflect this equal priority. The first of these options is defined in clauses 5 and 6, whereas the second is defined by replacing the content of those sections with the content of Annex C. The third is defined by replacing the content of those sections with the content of Annex D. Correspondingly, [18] and [2] apply only to the first option, and [8] apply to the second and third. Compliance with the present document requires compliance with one of these implementations, but not with all three. It is not required that equipment built to one option shall interoperate with equipment built to the other.

Compliance with frequency planning and EMC requirements is not covered by the present document and remains the operators' responsibility. In this respect, [14] and [15] are relevant to the USA; [1] and [i.5] to Canada; [7], [9], [10], [11], [12] and [13] are relevant to the European Union; [16] and [i.4] are relevant to China.

The present document defines the interface for the physical layer, and corresponds to the CableLabs specification [i.8].

2 References

2.1 Normative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at https://docbox.etsi.org/Reference/.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are necessary for the application of the present document.

[1] IEC CISPR 22:2008 (2008): "Information technology equipment - Radio disturbance characteristics - Limits and methods of measurement".

[2] CEA-542-D (2013): "CEA Standard: Cable Television Channel Identification Plan".

[3] ETSI EN 302 878-3: "Access, Terminals, Transmission and Multiplexing (ATTM); Third Generation Transmission Systems for Interactive Cable Television Services - IP Cable Modems; Part 3: Downstream Radio Frequency Interface; DOCSIS 3.0".

[4] ETSI TS 103 311-3: "Integrated broadband cable telecommunication networks (CABLE); Fourth Generation Transmission Systems for Interactive Cable Television Services - IP Cable Modems Part 3: MAC and Upper Layer Protocols Interface; DOCSIS® 3.1".

[5] ETSI EN 302 878-2: "Access, Terminals, Transmission and Multiplexing (ATTM); Third Generation Transmission Systems for Interactive Cable Television Services - IP Cable Modems; Part 2: Physical Layer; DOCSIS 3.0".

[6] ETSI EN 302 769 (V1.2.1): "Digital Video Broadcasting (DVB); Frame structure channel coding and modulation for a second generation digital transmission system for cable systems (DVB-C2)".

[7] ETSI EG 201 212 (V1.2.1): "Electrical safety; Classification of interfaces for equipment to be connected to telecommunication networks".

[8] ETSI EN 300 429 (V1.2.1): "Digital Video Broadcasting (DVB); Framing structure, channel coding and modulation for cable systems".

[9] CENELEC EN 50083-1 (2002): "Cabled distribution systems for television, sound and interactive multimedia signals - Part 1: Safety requirements".

[10] CENELEC EN 50083-2 (2005): "Cable networks for television signals, sound signals and interactive services -- Part 2: Electromagnetic compatibility for equipment".

[11] CENELEC EN 50083-7 (1996): "Cable networks for television signals, sound signals and interactive services -- Part 7: System performance".

[12] CENELEC EN 61000-6-4 (2001): "Electromagnetic compatibility (EMC) -- Part 6-4: Generic standards - Emission standard for industrial environments".

[13] CENELEC EN 61000-6-3 (2003): "Electromagnetic compatibility (EMC) -- Part 6-3: Generic standards - Emission standard for residential, commercial and light-industrial environments".

[14] Code of Federal Regulations, Title 47, Part 15 (October 2005).

[15] Code of Federal Regulations, Title 47, Part 76 (October 2005).

[16] Standardization Administration of People's republic of China (SAC): "Audio, video and similar electronic apparatus-Safety requirements".

NOTE: Available at www.sac.gov.cn.

[17] ISO/IEC-61169-24 (2001): "Radio-frequency connectors - Part 24: Sectional specification - Radio frequency coaxial connectors with screw coupling, typically for use in 75 ohm cable distribution systems (type F)".

[18] Annex B to Recommendation ITU-T J.83 (2007): "Digital multi-program systems for television sound and data services for cable distribution".

[19] ANSI/SCTE 02, Specification for "F" Port, Female Indoor.

[20] SCTE Measurement Recommended Practices for Cable Systems, Fourth Edition (March 2012).

[21] ETSI TS 103 311-1: "Integrated broadband cable telecommunication networks (CABLE); Fourth Generation Transmission Systems for Interactive Cable Television Services - IP Cable Modems; Part 1: General; DOCSIS® 3.1".

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] G. Castagnoli, S. Bräuer, and M. Herrmann: "Optimization of Cyclic Redundancy-Check Codes with 24 and 32 Parity Bits", IEEE Transactions on Communications, vol. 41, No. 6, pp. 883-892, June 1993.

[i.2] Cable Television Laboratories, Inc.: "DOCSIS 3.1 CCAP Operations Support System Interface Specification", CM-SP-CCAP-OSSIv3.1-I01-140808, August 8, 2014.

[i.3] Cable Television Laboratories, Inc.: "DOCSIS 3.1 Cable Modem Operations Support System Interface Specification", CM-SP-CM-OSSIv3.1-I01-140619, June 19, 2014.

[i.4] China Zhijian Publish House SAC: "Equipments and components used in cabled distribution systems primarily intended for television and sound signals--Part 1: Generic specifications".

[i.5] Information Technology Equipment (ITE): "Limits and methods of measurement".

[i.6] Cable Television Laboratories, Inc.: "Number and size of codewords versus grant sizes", CM-PHYv3.1-CODECHECK-131029.

NOTE: Available at https://community.cablelabs.com/wiki/plugins/servlet/cablelabs/alfresco/download?id=f039b211-b5dc-4b0d-81b7-bbf1dd782a64;1.0.

[i.7] PHYv3.1QAM Mapping, bit to constellation symbol mapping for DOCSIS 3.1, April 2014.

NOTE: Available at https://community.cablelabs.com/wiki/plugins/servlet/cablelabs/alfresco/download?id=794df4ae-500c-49dc-bfd9-1ddace26c235;1.0.

[i.8] Cable Television Laboratories, Inc.: "DOCSIS 3.1 Physical Layer Specification", CM-SP-PHYv3.1-I03-140610, June 10, 2014.

3 Definitions, symbols and abbreviations

3.1 Definitions

For the purposes of the present document, the following terms and definitions apply:

active channel: any channel which has been assigned to a cable modem's transmit channel set either in a registration response message or a dynamic bonding request message, and prior to registration

NOTE: After registration, the set of active channels also is called the transmit channel set. If the CMTS needs to add, remove, or replace channels in the cable modem's transmit channel set, it uses the dynamic bonding request message with transmit channel configuration encodings to define the de sired new transmit channel set. Note that the set of channels actually bursting upstream from a cable modem is a subset of that cable modem's active channels. In many instances one or all of a cable modem's active channels will not be bursting, but such quiet channels are still considered active channels for that cable modem.

active subcarrier:

1) In a downstream OFDM channel, any subcarrier other than an excluded subcarrier.

2) In an upstream OFDMA channel, any subcarrier other than an excluded subcarrier (subcarriers in zero-valued minislots as defined in OFDMA profiles, and unused subcarriers are considered active subcarriers because they are used in probes).

adaptive equalizer: circuit in a digital receiver that compensates for channel response impairments

NOTE: In effect, the circuit creates a digital filter that has approximately the opposite complex frequency response of the channel through which the desired signal was transmitted.

adaptive equalizer tap: See Tap.

adaptive pre-equalizer: circuit in a DOCSIS 1.1 or newer cable modem that pre-equalizes or pre-distorts the transmitted upstream signal to compensate for channel response impairments

NOTE: In effect, the circuit creates a digital filter that has approximately the opposite complex frequency response of the channel through which the desired signal is to be transmitted.

additive white Gaussian noise (AWGN): See Thermal Noise.

availability: ratio of time that a service, device, or network is available for use to total time, usually expressed as a percentage of the total time

NOTE: For example, four-nines availability, written as 99,99 %, means that a service is available 8 759,12 hours out of 8 760 total hours in a year.

BCH: class of error correction codes named after the inventors Raj Bose, D. K. Ray-Chaudhuri, and Alexis Hocquenghem

Binary Phase Shift Keying (BPSK): form of digital modulation in which two phases separated by 180 degrees support the transmission of one bit per symbol

Bit Error Rate (BER): See Bit Error Ratio.

Bit Error Ratio (BER): ratio of errored bits to the total number of bits transmitted, received, or processed over a defined amount of time

NOTE: Mathematically, BER = (number of errored bits)/(total number of bits) or BER = (error count in measurement period)/(bit rate * measurement period). Usually expressed in scientific notation format. Also called Bit Error Rate.

bit loading: technique of assigning the optimum number of bits (modulation order) for transmission per subcarrier in an OFDM or OFDMA system

burst: single continuous RF signal from the cable modem upstream transmitter, from transmitter on to transmitter off

burst noise:

1) Another name for impulse noise.

2) A type of noise comprising random and sudden step-like changes between levels, often occurring in semiconductors.

NOTE: Sometimes called Popcorn Noise.

Cable Modem (CM): modulator-demodulator at the subscriber premises intended for use in conveying data communications on a cable television system

Cable Modem Termination System (CMTS): device located at the cable television system headend or distribution hub, which provides complementary functionality to the cable modems to enable data connectivity to a wide-area network

Carrier-to-Noise Ratio (CNR or C/N):

1) The ratio of signal (or carrier) power to noise power in a defined measurement bandwidth.

2) For OFDM and OFDMA signals, the ratio of average signal power (PSIGNAL) in the occupied bandwidth to the average noise power in the occupied bandwidth given by the noise power spectral density integrated over the same occupied bandwidth, expressed mathematically as $CNR = 10\ log_{10}[P_{SIGNAL} / \int N(f)df]$ dB. This is a lower bound on the actual received signal-to-noise ratio.

3) For SC-QAM, the ratio of the average signal power (PSIGNAL) to the average noise power in the QAM signal's symbol rate bandwidth (NSYM), and expressed mathematically as CNR = 10 log10(PSIGNAL/NSYM) dB or equivalently for an AWGN channel as CNR = 10 log10(ES/N0) dB.

NOTE 1: For an AWGN channel, P_SIGNAL⁄(N_SYM=((E_S⁄T_S))⁄((N_0 B_N)))=((E_S⁄T_S))⁄((N_0⁄T_S)=E_S⁄N_0), where ES and TS are the symbol energy and duration respectively, N0 is the noise power spectral density, and BN is the noise bandwidth equal to the symbol rate bandwidth 1/TS. 4) For analog television signals, the ratio of visual carrier peak envelope power during the transmission of synchronizing pulses (PPEP) to noise power (N), where the visual carrier power measurement bandwidth is nominally 300 kHz and the noise power measurement bandwidth is 4 MHz for NTSC signals. For the latter, the noise measurement bandwidth captures the total noise power present over a 4 MHz band centered within the television channel, and is expressed mathematically as CNR = 10 log10(PPEP/N) dB.

NOTE 2: For analog PAL and SECAM channels, the noise measurement bandwidth is a larger value than the 4 MHz specified for NTSC (4,75 MHz, 5,00 MHz, 5,08 MHz, or 5,75 MHz, depending on the specific system).

CEA-542: Consumer Electronics Association standard that defines a channel identification plan for 6 MHz-wide channel frequency allocations in cable systems

ceiling: mathematical function that returns the lowest-valued integer that is greater than or equal to a given value

channel: portion of the electromagnetic spectrum used to convey one or more RF signals between a transmitter and receiver

NOTE: May be specified by parameters such as centre frequency, bandwidth, or CEA channel number. May be referred to as RF Channel.

codeword: forward error correction data block, comprising a combination of information bytes and parity bytes

Codeword Error Ratio (CER): ratio of errored codewords to the total number of codewords transmitted, received, or processed over a defined amount of time

NOTE: Mathematically, CER = (number of errored codewords)/(total number of codewords). Usually expressed in scientific notation format.

coefficient: complex number that establishes the gain of each tap in an adaptive equalizer or adaptive pre-equalizer

Common Path Distortion (CPD): clusters of second and third order distortion beats generated in a diode-like nonlinearity such as a corroded connector in the signal path common to downstream and upstream

NOTE: The beats tend to be prevalent in the upstream spectrum. When the primary RF signals are digitally modulated signals instead of analog TV channels, the distortions are noise-like rather than clusters of discrete beats.

complementary pilots: subcarriers that carry data, but with a lower modulation order than other data subcarriers in a given minislot

NOTE: Complementary pilots allow phase tracking along the time axis for frequency offset and phase noise correction, and may be used by the CMTS upstream receiver to enhance signal processing, such as improving the accuracy of centre frequency offset acquisition.

Composite Second Order (CSO): clusters of second order distortion beats generated in cable network active devices that carry multiple RF signals

NOTE: When the primary RF signals are digitally modulated signals instead of analog TV channels, the distortions are noise-like rather than clusters of discrete beats.

Composite Triple Beat (CTB): clusters of third order distortion beats generated in cable network active devices that carry multiple RF signals

NOTE: When the primary RF signals are digitally modulated signals instead of analog TV channels, the distortions are noise-like rather than clusters of discrete beats.

continuous pilots: pilots that occur at the same frequency location in every OFDM symbol, and which are used for frequency and phase synchronization

convolution: process of combining two signals in which one of the signals is time-reversed and correlated with the other signal

NOTE: The output of a filter is the convolution of its impulse response with the input signal.

convolutional interleaver: interleaver in which symbols are sequentially shifted into a bank of "I" registers and each successive register has "J" symbols more storage than the preceding register

NOTE: The first interleaver path has zero delay, the second has a J symbol period of delay, the third 2 x J symbol periods of delay, etc. up to the Ith path which has (I - 1) x J symbol periods of delay. This process is reversed in the receiver's deinterleaver so that the net delay of each symbol is the same through the interleaver and deinterleaver. See also interleaver.

correlation:

1) A process of combining two signals in which the signals are multiplied sample-by-sample and summed; the process is repeated at each sample as one signal is slid in time past the other.

2) Cross-correlation is a measure of similarity between two signals.

cross modulation (XMOD): form of television signal distortion in which modulation from one or more television signals is imposed on another signal or signals

Customer Premises Equipment (CPE): device such as a cable modem or set-top at the subscriber's or other end user's location. May be provided by the end user or the service provider

cyclic prefix (CP): copy of the end of a symbol that is added to the beginning of the same symbol, in order to help mitigate the effects of micro-reflections and similar impairments

decibel (dB): ratio of two power levels expressed mathematically as dB = 10log10(P1/P2)

decibel carrier (dBc): ratio of the power of a signal to the power of a reference carrier, expressed mathematically as dBc = 10log10(Psignal/Pcarrier)

decibel millivolt (dBmV): unit of RF power expressed in terms of voltage, defined as decibels relative to 1 millivolt, where 1 millivolt equals 13,33 nanowatts in a 75 ohm impedance

NOTE: Mathematically, dBmV = 20log10(value in mV/1 mV).

decibel reference (dBr): ratio of a signal level to a reference signal level

NOTE: When the signals are noise or noise-like, the measurement bandwidth for the two signals is the same. When both signal levels are in the same units of power, the ratio is expressed mathematically as dBr = 10log10(Psignal/Preference). When both signal levels are in the same units of voltage, and assuming the same impedance, the ratio is expressed mathematically as dBr = 20log10(Vsignal/Vreference).

Discrete Fourier Transform (DFT): part of the family of mathematical methods known as Fourier analysis, which defines the "decomposition" of signals into sinusoids

NOTE: Discrete Fourier transform defines the transformation from the time to the frequency domain. See also inverse discrete Fourier transform.

distortion: See linear distortion and nonlinear distortion.

distribution hub: facility in a cable network which performs the functions of a headend for customers in its immediate area, and which receives some or all of its content for transmission from a master headend in the same metropolitan or regional area

DOCSIS: Data-Over-Cable Service Interface Specifications

NOTE: A group of specifications that defines interoperability between cable modem termination systems and cable modems.

DOCSIS 1.x: Abbreviation for DOCSIS versions 1.0 or 1.1.

downstream: direction of RF signal transmission from headend or hub site to subscriber

NOTE: In North American cable networks, the downstream or forward spectrum may occupy frequencies from just below 54 MHz to as high as 1 002 MHz or more. The DOCSIS 3.1 downstream is 258 MHz (optional 108 MHz) to 1 218 MHz (optional 1 794 MHz).

downstream channel: portion of the electromagnetic spectrum used to convey one or more RF signals from the headend or hub site to the subscriber premises

NOTE: For example, a single CEA channel's bandwidth is 6 MHz, and a downstream DOCSIS 3.1 OFDM channel's bandwidth can be up to 192 MHz.

downstream modulated spectrum: encompassed spectrum minus the excluded subcarriers within the encompassed spectrum, where excluded subcarriers include all the individually excluded subcarriers and all the subcarriers comprising excluded subbands

NOTE: This also is the spectrum comprising all active subcarriers. For this definition, the width of an active or excluded subcarrier is equal to the subcarrier spacing.

downstream occupied bandwidth: sum of the bandwidth in all standard channel frequency allocations (e.g. 6 MHz spaced CEA channels) that are occupied by the OFDM channel

NOTE: Even if one active subcarrier of an OFDM channel is placed in a given standard channel frequency allocation, that standard channel frequency allocation in its entirety is said to be occupied by the OFDM channel.

drop: coaxial cable and related hardware that connects a residence or other service location to a tap in the nearest coaxial feeder cable

NOTE: Also called drop cable or subscriber drop.

Dynamic Host Configuration Protocol (DHCP): protocol that defines the dynamic or temporary assignment of Internet protocol addresses, so that the addresses may be reused when they are no longer needed by the devices to which they were originally assigned

Dynamic Range Window (DRW): in DOCSIS 3.1, the range, in decibels, of the maximum difference in power per 1,6 MHz between multiple transmitters in a cable modem that is operating in multiple transmit channel mode

NOTE: In DOCSIS 3.0, the range, in decibels, of the maximum power difference between multiple transmitters in a cable modem that is operating in multiple transmit channel mode.

encompassed spectrum:

1) For an OFDM or OFDMA channel, the range of frequencies from the centre frequency of the channel's lowest active subcarrier minus half the subcarrier spacing, to the centre frequency of the channel's highest active subcarrier plus half the subcarrier spacing.

2) For an SC-QAM channel, the encompassed spectrum is the signal bandwidth (i.e. 6 MHz or 8 MHz in the downstream; 1,6 MHz, 3,2 MHz and 6,4 MHz in the upstream).

3) For the RF output of a downstream or upstream port including multiple OFDM, OFDMA, and/or SC-QAM channels, the range of frequencies from the lowest frequency of the encompassed spectrum of the lowest frequency channel to the highest frequency of the encompassed spectrum of the highest frequency channel.

equivalent legacy DOCSIS channels: within a downstream OFDM channel, an integer number equal to ceil(modulated spectrum in MHz / 6)

excluded subcarrier: subcarrier that cannot be used because another type of service is using the subcarrier's frequency or a permanent ingressor is present on the frequency

NOTE: The CMTS or cable modem is administratively configured to not transmit on excluded subcarriers.

exclusion band: set of contiguous subcarriers within the OFDM or OFDMA channel bandwidth that are set to zero-value by the transmitter to reduce interference to other co-existing transmissions such as legacy SC-QAM signals

F connector: threaded, nominally 75-ohm impedance RF connector, whose electrical and physical specifications are defined in various SCTE standards

NOTE: Commonly used on smaller sizes of coaxial cable such as 59-series and 6-series, and on mating interfaces of subscriber drop components, customer premises equipment, and some headend and test equipment.

Fast Fourier Transform (FFT): algorithm to compute the discrete Fourier transform from the time domain to the frequency domain, typically far more efficiently than methods such as correlation or solving simultaneous linear equations

NOTE: See also discrete Fourier transform, inverse discrete Fourier transform, and inverse fast Fourier transform.

FFT duration: reciprocal of subcarrier spacing. For example, with 50 kHz subcarrier spacing, FFT duration is 20 µs, and with 25 kHz subcarrier spacing, FFT duration is 40 µs

NOTE: Sometimes called "useful symbol duration." See also Symbol Duration.

fibre node: See Node.

filler subcarrier: zero bit loaded subcarrier that is inserted in an OFDM symbol when no data is transmitted, or when the number of codewords has exceeded the upper limit, or when it is not possible to begin a new codeword because of insufficient space to include a next codeword pointer

floor: mathematical function that returns the highest-valued integer that is less than or equal to a given value

forward: See Downstream.

Forward Error Correction (FEC): method of error detection and correction in which redundant information is sent with a data payload in order to allow the receiver to reconstruct the original data if an error occurs during transmission

Frequency Division Multiple Access (FDMA): multiple access technology that accommodates multiple users by allocating each user's traffic to one or more discrete frequency bands, channels, or subcarriers

Frequency Division Multiplexing (FDM): transmission of multiple signals through the same medium at the same time. Each signal is on a separate frequency or assigned to its own channel

NOTE: For example, an analog TV signal might be carried on CEA channel 7 (174 MHz-180 MHz), a 256-QAM digital video signal on channel 8 (180-186 MHz) and so on.

frequency response: complex quantity describing the flatness of a channel or specified frequency range, and which has two components: amplitude (magnitude)-versus-frequency and phase-versus-frequency

gigahertz (GHz): one billion (10^9) hertz

NOTE: See also hertz.

Group Delay (GD): negative derivative of phase with respect to frequency, expressed mathematically as GD = -(dφ/dω) and stated in units of time such as nanoseconds or microseconds

group delay ripple: group delay variation which has a sinusoidal or scalloped sinusoidal shape across a specified frequency range

Group Delay Variation (GDV) or group delay distortion: difference in group delay between one frequency and another in a circuit, device or system

guard interval: in the time domain, the period from the end of one symbol to the beginning of the next symbol, which includes the cyclic prefix and applied transmit windowing

NOTE: Also called guard time.

guard band: narrow range of frequencies in which user data is not transmitted, located at the lower and upper edges of a channel, at the lower and upper edges of a gap within a channel, or in between channels

NOTE: A guard band minimizes interference from adjacent signals, but is not needed in the case of adjoining OFDM channels that are synchronous with identical FFT size and cyclic prefix that would not mutually interfere.

Harmonic Related Carriers (HRC): method of spacing channels on a cable television system defined in [2], in which visual carriers are multiples of 6,0003 MHz

NOTE: A variation of HRC channelization used in some European cable networks is based upon multiples of 8 MHz.

headend: central facility that is used for receiving, processing, and combining broadcast, narrowcast and other signals to be carried on a cable network

NOTE 1: Somewhat analogous to a telephone company's central office. See also Distribution Hub.

NOTE 2: Location from which the DOCSIS cable plant fans out to subscribers.**header:** protocol control information located at the beginning of a protocol data unit

hertz (Hz): unit of frequency equivalent to one cycle per second

hum modulation: amplitude distortion of a signal caused by the modulation of that signal by components of the power source (e.g. 60 Hz) and/or its harmonics

Hybrid Fibre/Coax (HFC): broadband bidirectional shared-media transmission system or network architecture using optical fibres between the headend and fibre nodes, and coaxial cable distribution from the fibre nodes to the subscriber locations

impedance: combined opposition to current in a component, circuit, device, or transmission line that contains both resistance and reactance

NOTE: Represented by the symbol Z and expressed in ohms.

impulse noise: noise that is bursty in nature, characterized by non-overlapping transient disturbances

NOTE: May be periodic (e.g. automobile ignition noise or high-voltage power line corona noise), or random (e.g. switching noise or atmospheric noise from thunderstorms). Generally of short duration - from about 1 microsecond to a few tens of microseconds - with a fast risetime and moderately fast falltime.

Incremental Related Carriers (IRC): method of spacing channels on a cable television system defined in [2], in which all visual carriers except channels 5 and 6 are offset +12,5 kHz with respect to the standard channel plan

NOTE: Channels 5 and 6 are offset +2,0125 MHz with respect to the standard channel plan. See also standard frequencies.

in-phase (I): real part of a vector that represents a signal, with 0 degrees phase angle relative to a reference carrier

NOTE: See also Quadrature (Q).

interleaver: subset or layer of the forward error correction process, in which the data to be transmitted is rearranged or mixed such that the original bits, bytes, or symbols are no longer adjacent

NOTE: The latter provides improved resistance to various forms of interference, especially burst or impulse noise. Interleaving may be performed in the time domain, frequency domain, or both. A de-interleaver in the receiver rearranges the bits, bytes, or symbols into their original order prior to additional error correction.

International Electrotechnical Commission (IEC): organization that prepares and publishes international standards for electrical, electronic and related technologies

International Organization for Standardization (ISO): organization that develops voluntary international standards for technology, business, manufacturing, and other industries

Internet Engineering Task Force (IETF): body responsible, among other things, for developing standards used in the Internet

Internet Protocol (IP): network layer protocol that supports connectionless internetwork service, and which contains addressing and control information that allows packets to be routed

internet protocol detail record: record formatter and exporter functions of the CMTS that provides information about Internet protocol-based service usage, and other activities that can be used by operational support systems and business support systems

Inverse Discrete Fourier Transform (IDFT): part of the family of mathematical methods known as Fourier analysis, which defines the "decomposition" of signals into sinusoids

NOTE: Inverse discrete Fourier transform defines the transformation from the frequency to the time domain. See also discrete Fourier transform.

Inverse Fast Fourier Transform (IFFT): algorithm to compute the inverse discrete Fourier transform from the frequency domain to the time domain, typically far more efficiently than methods such as correlation or solving simultaneous linear equations

NOTE: See also discrete Fourier transform, fast Fourier transform, and inverse discrete Fourier transform.

jitter: fluctuation in the arrival time of a regularly scheduled event such as a clock edge or a packet in a stream of packets. Jitter is defined as fluctuations above 10 Hz

kilohertz (kHz): one thousand (10^3) hertz

NOTE: See also hertz.

latency:

1) The time taken for a signal element to propagate through a transmission medium or device.

2) The delay between a device's request for network access and when permission is granted for transmission.

3) The delay from when a frame is received by a device to when the frame is forwarded via the device's destination port.

layer: one of seven subdivisions of the Open System Interconnection reference model

linear distortion: class of distortions that occurs when the overall response of the system (including transmitter, cable plant, and receiver) differs from the ideal or desired response

NOTE: This class of distortions maintains a linear, or 1:1, signal-to-distortion relationship (increasing signal by 1 dB causes distortion to increase by 1 dB), and often occurs when amplitude-versus-frequency and/or phase-versus-frequency depart from ideal. Linear distortions include impairments such as micro-reflections, amplitude ripple, and group delay variation, and can be corrected by an adaptive equalizer.

Low Density Parity Check (LDPC): error correction code used in DOCSIS 3.1. LDPC is more robust than Reed-Solomon error correction codes

MAC address: "built-in" hardware address of a device connected to a shared medium

MAC frame: MAC header plus optional protocol data unit

MAC Management Message (MMM): unclassified traffic between the CMTS and cable modem

NOTE: Examples include MAC domain descriptor, ranging-request, ranging-response, and upstream channel descriptor messages.

Media Access Control (MAC): sublayer of the Open Systems Interconnection model's data link layer (Layer 2), which manages access to shared media such as the Open Systems Interconnection model's physical layer (Layer 1)

Megahertz (MHz): one million (10^6) hertz

NOTE: See also hertz.

micro-reflection: short time delay echo or reflection caused by an impedance mismatch

NOTE: A micro-reflection's time delay is typically in the range from less than a symbol period to several symbol periods.

microsecond (µs): one millionth (10^{-6}) of a second

millisecond (ms): one thousandth (10^{-3}) of a second

millivolt (mV): one thousandth (10^{-3}) of a volt

minislot: group of dedicated subcarriers, all with the same modulation order, for upstream transmission by a given cable modem in DOCSIS 3.1 OFDMA applications

NOTE: In DOCSIS 3.0 and earlier TDMA applications, a unit of time for upstream transmission that is an integer multiple of 6,25 µs units of time called "ticks." For both TDMA and OFDMA, a cable modem may be assigned one or more minislots in a transmission burst.

modulated spectrum: See either Downstream Modulated Spectrum or Upstream Modulated Spectrum.

NOTE: For this definition, the width of a transmitted subcarrier is equal to the subcarrier spacing.

Modulation Error Ratio (MER): ratio of average signal constellation power to average constellation error power - that is, digital complex baseband signal-to-noise ratio - expressed in decibels

NOTE: In effect, MER is a measure of how spread out the symbol points in a constellation are. More specifically, MER is a measure of the cluster variance that exists in a transmitted or received waveform at the output of an ideal receive matched filter. MER includes the effects of all discrete spurious, noise, carrier leakage, clock lines, synthesizer products, linear and nonlinear distortions, other undesired transmitter and receiver products, ingress, and similar in-channel impairments.

modulation rate: signalling rate of the upstream modulator (for example, 1 280 to 5 120 kHz)

NOTE: In S-CDMA it is the chip rate. In TDMA, it is the channel symbol rate.

Multiple Transmit Channel (MTC) [Mode]: operational mode in a cable modem that enables the simultaneous transmission of more than one upstream channel

nanosecond (ns): one billionth (10^{-9}) of a second

National Television System Committee (NTSC): committee that defined the analog television broadcast standards (black and white in 1941, colour in 1953) used in North America and some other parts of the world

NOTE: The NTSC standards are named after the committee.

Next Codeword Pointer (NCP): message block used to identify where a codeword begins

node: optical-to-electrical (RF) interface between a fibre optic cable and the coaxial cable distribution network

NOTE: Also called Fibber Node.

noise: typically any undesired signal or signals - other than discrete carriers or discrete distortion products - in a device, communications circuit, channel or other specified frequency range

NOTE: See also Impulse Noise, Phase Noise, and Thermal Noise.

nonlinear distortion: class of distortions caused by a combination of small signal nonlinearities in active devices and by signal compression that occurs as RF output levels reach the active device's saturation point

NOTE: Nonlinear distortions generally have a nonlinear signal-to-distortion amplitude relationship - for instance, 1:2, 1:3 or worse (increasing signal level by 1 dB causes distortion to increase by 2 dB, 3 dB or more). The most common nonlinear distortions are even order distortions such as composite second order, and odd order distortions such as composite triple beat. Passive components can generate nonlinear distortions under certain circumstances.

occupied bandwidth: See either Downstream Occupied Bandwidth or Upstream Occupied Bandwidth.

orthogonal: distinguishable from or independent such that there is no interaction or interference

NOTE: In OFDM, subcarrier orthogonality is achieved by spacing the subcarriers at the reciprocal of the symbol period (T), also called symbol duration time. This spacing results in the sinc (sin x/x) frequency response curves of the subcarriers lining up so that the peak of one subcarrier's response curve falls on the first nulls of the lower and upper adjacent subcarriers' response curves. Orthogonal subcarriers each have exactly an integer number of cycles in the interval T.

Orthogonal Frequency Division Multiple Access (OFDMA): OFDM-based multiple-access scheme in which different subcarriers or groups of subcarriers are assigned to different users

OFDMA channel bandwidth: occupied bandwidth of an upstream OFDMA channel

NOTE: See also Occupied Bandwidth.

Orthogonal Frequency Division Multiplexing (OFDM): data transmission method in which a large number of closely-spaced or overlapping very-narrow-bandwidth orthogonal QAM signals are transmitted within a given channel

NOTE: Each of the QAM signals, called a subcarrier, carries a small percentage of the total payload at a very low data rate.

OFDM channel bandwidth: occupied bandwidth of a downstream OFDM channel

NOTE: See also Occupied Bandwidth.

OFDMA frame: group of a configurable number, K, of consecutive OFDMA symbols

NOTE: A frame comprises either a group of probing symbols or a column of minislots across the spectrum of the OFDMA channel. Multiple modems can share the same OFDMA frame simultaneously by transmitting data and pilots on allocated subcarriers within the frame.

phase noise: rapid, short-term, random fluctuations in the phase of a wave, caused by time domain instabilities

PHY Link Channel (PLC): set of contiguous OFDM subcarriers (eight for 4K FFT and 16 for 8K FFT), constituting a "sub-channel" of the OFDM channel, which conveys physical layer parameters from the CMTS to the cable modem

PHY link channel frame: in downstream OFDM transmission, a group of 128 consecutive OFDM symbols, beginning with the first OFDM symbol following the last OFDM symbol containing the PLC preamble

physical layer (PHY): layer 1 in the Open System Interconnection architecture; the layer that provides services to transmit bits or groups of bits over a transmission link between open systems and which entails electrical, mechanical and handshaking procedures

picosecond (ps): one trillionth (10^{-12}) of a second

pilot: dedicated OFDM subcarrier that may be used for such purposes as channel estimation (measurement of channel condition), synchronization, and other purposes

NOTE: See also Complementary Pilots, Continuous Pilots, and Scattered Pilots.

pload_min_set: reduction of channel n maximum transmit power with respect to Phi_n, when multiple transmit channel mode is enabled

NOTE: The value of Pload_min_set is sent by the CMTS to the cable modem.

pload_n: reduction of channel n reported transmit power with respect to Phi_n, where Phi_n = min(Pmax - Gconst) over all burst profiles used by the cable modem in channel n

preamble: data sequence transmitted at or near the beginning of a frame, allowing the receiver time to achieve lock and synchronization of transmit and receive clocks

pre-equalizer: See Adaptive Pre-equalizer.

profile: set of parameters that defines how information is transmitted from a CMTS to a cable modem, or from a cable modem to a CMTS

NOTE: Examples of some of the parameters defined in a profile include modulation order, forward error correction, preamble, and guard interval.

protocol: description of a set of rules and formats that specify how devices on a network exchange data

Pseudo Random Binary Sequence (PRBS): deterministic sequence of bits that appears to be random, that is, with no apparent pattern

NOTE: Also called Pseudo Random Bit Stream.

QAM signal: analog RF signal that uses quadrature amplitude modulation to convey information such as digital data

quadrature (Q): imaginary part of a vector that represents a signal, with 90 degrees phase angle relative to a reference carrier

NOTE: See also In-Phase (I).

Quadrature Amplitude Modulation (QAM): modulation technique in which an analog signal's amplitude and phase vary to convey information, such as digital data

NOTE: The name "quadrature" indicates that amplitude and phase can be represented in rectangular coordinates as in-phase (I) and quadrature (Q) components of a signal.

Quadrature Phase Shift Keying (QPSK): form of digital modulation in which four phase states separated by 90 degrees support the transmission of two bits per symbol

NOTE: Also called 4-QAM.

Radio Frequency (RF): that portion of the electromagnetic spectrum from a few kilohertz to just below the frequency of infrared light

randomizer: subset or layer of the forward error correction process, in which the data to be transmitted is randomized using a PRBS scrambler

NOTE: Randomization spreads out the energy across the spectrum, ensures uniform population of all of the data constellation points, and minimizes the likelihood of long strings of all zeros or ones. Also referred to as a scrambler.

Receive Channel Set (RCS): combination of legacy SC-QAM and OFDM channels that the cable modem has been configured to receive by the CMTS

Reed-Solomon (R-S): class of error correction codes named after the inventors Irving Reed and Gustave Solomon. The forward error correction in DOCSIS 3.0 and earlier uses Reed-Solomon error correction codes

return: See upstream.

return loss (R): ratio of incident power PI to reflected power PR, expressed mathematically as R = 10log10(PI/PR), where R is return loss in decibels

reverse: See upstream.

RF channel: See channel.

Roll-off Period (RP): duration in microseconds, or the equivalent number of IFFT output sample periods, used for the ramping up (or ramping down) transition region of the Tukey raised-cosine window, which is applied at the beginning (and end) of an OFDM symbol

NOTE: A sampling rate of 102,4 MHz is assumed for the upstream and 204,8 MHz is assumed for the downstream. The roll-off period contains an even number of samples with weighting coefficients between, but not including, 0 and 1. The roll-off, which ramps down at the end of a symbol, overlaps the mirror-image roll-off which ramps up at the beginning of the following symbol, and the two segments add to unity. In the case of no transmit windowing, the roll-off duration is zero and there are no samples in the roll-off period.

root mean square: statistical measure of the magnitude of a varying quantity such as current or voltage, where the RMS value of a set of instantaneous values over, say, one cycle of alternating current is equal to the square root of the mean value of the squares of the original values

scattered pilots: pilots that do not occur at the same frequency in every symbol, and which may be used for channel estimation

NOTE: The locations of scattered pilots change from one OFDM symbol to another.

signal-to-composite noise: ratio of signal power to composite noise power in a defined measurement bandwidth, where composite noise is the combination of thermal noise and composite intermodulation distortion (noise-like distortion)

single carrier quadrature amplitude modulation: data transmission method used in DOCSIS 1.x, 2.0 and 3.0, in which each downstream or upstream RF channel slot carries only one QAM signal

OFDM/OFDMA channel spectral edge: centre frequency of the channel's lowest active subcarrier minus half the subcarrier spacing; and the centre frequency of the channel's highest active subcarrier plus half the subcarrier spacing

OFDM/OFDMA exclusion band channel spectral edge: centre frequency of the channel's highest active subcarrier plus half the subcarrier spacing adjacent to the beginning of an exclusion band, and the centre frequency of the channel's lowest active subcarrier minus half the subcarrier spacing adjacent to the end of the same exclusion band

standard frequencies: method of spacing channels on a cable television system defined in [2]

NOTE: Channels 2-6 and 7-13 use the same frequencies as over-the-air channels 2-6 and 7-13. Other cable channels below Ch. 7 down to 91,25 MHz and above Ch. 13 are spaced in 6 MHz increments.

subcarrier: one of a large number of closely spaced or overlapping orthogonal narrow-bandwidth data signals within an OFDM channel

NOTE: Also called a Tone. See also Excluded Subcarrier, Unused Subcarrier, and Used Subcarrier.

sublayer: subdivision of a layer in the Open System Interconnection reference model

subscriber: end user or customer connected to a cable network

subslot: subdivision or subunit of a minislot that fits within a minislot's boundaries, used to provide multiple transmission opportunities for bandwidth requests

NOTE: Data subcarriers within a subslot use QPSK, and are not FEC encoded.

symbol duration: sum of the FFT duration and cyclic prefix duration

NOTE: Symbol duration is greater than FFT duration, because symbol duration includes a prepended cyclic prefix.

synchronous code division multiple access: multiple access physical layer technology in which different transmitters can share a channel simultaneously

NOTE: The individual transmissions are kept distinct by assigning each transmission an orthogonal "code." Orthogonality is maintained by all transmitters being precisely synchronized with one another.

tap:

1) In the feeder portion of a coaxial cable distribution network, a passive device that comprises a combination of a directional coupler and splitter to "tap" off some of the feeder cable RF signal for connection to the subscriber drop. So-called self-terminating taps used at feeder ends-of-line are splitters only and do not usually contain a directional coupler. Also called a multitap.

2) The part of an adaptive equalizer where some of the main signal is "tapped" off, and which includes a delay element and multiplier.

3) One term of the difference equation in a finite impulse response or an infinite impulse response filter.

NOTE 1: The gain of the multipliers is set by the equalizer's coefficients. Also referred to as an Equalizer Tap.

NOTE 2: The difference equation of a FIR follows: $y(n) = b0x(n) + b1x(n-1) + b2x(n-2) + \dots + bNx(n-N)$.

thermal noise: fluctuating voltage across a resistance due to the random motion of free charge caused by thermal agitation

NOTE: Also called Johnson-Nyquist noise. When the probability distribution of the voltage is Gaussian, the noise is called Additive White Gaussian Noise (AWGN).

time division multiple access: multiple access technology that enables multiple users to access, in sequence, a single RF channel by allocating unique time slots to each user of the channel

transit delay: time required for a signal to propagate or travel from one point in a network to another point in the network

NOTE: For example, from the CMTS to the most distant cable modem. Also called Propagation Delay.

transmit channel set: combination of legacy SC-QAM channels and OFDMA channels that may be transmitted by a cable modem

transmit pre-equalizer: See Adaptive Pre-Equalizer.

under-grant hold bandwidth: minimum grant bandwidth that can be allocated beyond which the spurious emissions limits (in dBc) are no longer relaxed as the based on grant size continues to decrease

NOTE: Defined mathematically as UGHB = (100 % grant spectrum)/(under-grant hold number of users).

under-grant hold number of users: maximum number of equal-size grants that can be allocated beyond which the spurious emissions limits (in dBc) are no longer relaxed as the number of based on grants size continues to increase

NOTE: Defined mathematically as UGHU = floor[0,2 + 10((-44 - SpurFloor)/10)].

unused subcarrier: subcarriers in an upstream OFDMA channel which are not excluded, but are not assigned to minislots

NOTE: For example, unused subcarriers may occur when the number of subcarriers between excluded subcarriers is not divisible by the fixed number of consecutive subcarriers which comprise every OFDMA minislot. Thus, after constructing minislots from a group of consecutive non-excluded subcarriers, the remainder will become unused subcarriers. Unused subcarriers are not used for data transmission, but still carry power during probe transmission.

upstream:

1) The direction of RF signal transmission from subscriber to headend or hub site. In most North American cable networks, the legacy upstream spectrum occupies frequencies from 5 MHz to as high as 42 MHz.

2) The DOCSIS 3.1 upstream is 5-204 MHz, with support for 5-42 MHz, 5-65 MHz, 5-85 MHz and 5-117 MHz.

NOTE: Also called Return or Reverse.

upstream channel: portion of the electromagnetic spectrum used to convey one or more RF signals from the subscriber premises to the headend or hub site

NOTE: For example, a commonly used DOCSIS 3.0 upstream channel bandwidth is 6,4 MHz. A DOCSIS 3.1 upstream OFDMA channel bandwidth may be as much as 96 MHz.

upstream channel descriptor: MAC management message used to communicate the characteristics of the upstream physical layer to the cable modems

upstream modulated spectrum: spectrum comprising all non-zero-valued subcarriers of a cable modem's OFDMA transmission, resulting from the exercised transmit opportunities

NOTE: This also is the spectrum comprising all active subcarriers. For this definition, the width of a transmitted subcarrier is equal to the subcarrier spacing.

upstream occupied bandwidth:

a) For a single OFDMA channel, the sum of the bandwidth in all the subcarriers of that OFDMA channel which are not excluded. The upstream occupied bandwidth is calculated as the number of subcarriers which are not excluded, multiplied by the subcarrier spacing.

b) For the transmit channel set, the sum of the occupied bandwidth of all OFDMA channels plus the bandwidth of the legacy channels (counted as 1,25 times the modulation rate for each legacy channel) in a cable modem's transmit channel set.

NOTE: The combined bandwidth of all the minislots in the channel is normally smaller than the upstream occupied bandwidth due to the existence of unused subcarriers. The bandwidth occupied by an OFDMA probe with a skip value of zero is equal to the upstream occupied bandwidth.

used subcarrier: upstream subcarrier that is part of a minislot

NOTE: The cable modem transmits data, ranging, and probes on these subcarriers when instructed to do so by MAP messages. It is a MULPI term.

useful symbol duration: See FFT Duration.

vector: quantity that expresses magnitude and direction (or phase), and is represented graphically using an arrow

windowing: technique to shape data in the time domain, in which a segment of the start of the IFFT output is appended to the end of the IFFT output to taper or roll-off the edges of the data using a raised cosine function

NOTE: Windowing maximizes the capacity of the channel by sharpening the edges of the OFDM/A signal in the frequency domain.

word: information part of a codeword, without parity

NOTE: See also Codeword.

zero bit-loaded-subcarrier: subcarrier with power but not carrying user data, although it could be modulated by a pseudo-random binary sequence

zero-valued minislot: minislot composed of zero-valued subcarriers and no pilots

zero-valued subcarrier: subcarrier with no power

NOTE: See also Excluded Subcarrier.

3.2 Symbols

For the purposes of the present document, the following symbols apply:

η	Eta (lower case); a measure of modulation
μ	Mu/micron (lower case); represents 10^{-6}
τ	Tau (lower case); a measure of time
μs	Microsecond

3.3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

ADC	Analog-to-Digital Converter
ANSI	American National Standards Institute
AWGN	Additive White Gaussian Noise
BCH	Bose, Ray-Chaudhuri, Hocquenghem [codes]
BER	Bit Error Ratio
BER	Bit Error Rate
BPSK	Binary Phase Shift Keying
BW	BandWidth
CableLabs	Cable Television Laboratories, Inc.
CDMA	Code Division Multiple Access
CEA	Consumer Electronics Association
ceil	ceiling
CENELEC	European Committee for Electrotechnical Standardization
CER	codeword error ratio
CL	Convergence Layer
CL	CableLabs
CM	Cable Modem
CMCI	Cable Modem-to-Customer premises equipment Interface
CMTS	Cable Modem Termination System
CNR	Carrier to Noise Ratio
CP	Cyclic Prefix
CP	Complementary Pilot
CPD	Common Path Distortion
CPE	Customer Premises Equipment
CPU	Central Processing Unit
CRC	Cyclic Redundancy Check
CS	Cyclic Suffix
CSO	Composite Second Order
CTB	Composite Triple Beat
CW	Continuous Wave
CW	CodeWord
DAC	Digital to Analog Converter
dB	decibel
dBc	decibel carrier
dBmV	decibel millivolt
dBr	decibel reference
DCID	downstream channel identifier
DEPI	downstream external PHY interface
DFT	discrete Fourier transform
DHCP	dynamic host configuration protocol
DLS	DOCSIS light sleep [mode]
DOCSIS 1.x	Data-Over-Cable Service Interface Specifications version 1.0 or 1.1
DOCSIS 2.0	Data-Over-Cable Service Interface Specifications version 2.0
DOCSIS 3.0	Data-Over-Cable Service Interface Specifications version 3.0
DOCSIS 3.1	Data-Over-Cable Service Interface Specifications version 3.1
DOCSIS	Data-Over-Cable Service Interface Specifications
DRFI	[DOCSIS] Downstream Radio Frequency Interface [Specification]
DRW	dynamic range window
DS	downstream
DSG	DOCSIS Set-top Gateway [Interface Specification]
DTI	DOCSIS Timing Interface [Specification]
DUT	Device Under Test
DVB	Digital Video Broadcasting [Project]
DVB-C2	"Digital Video Broadcasting (DVB); Frame structure channel coding and modulation for a second generation digital transmission system for cable systems (DVB-C2)"
eDOCSIS	embedded Data-Over-Cable Service Interface Specifications
EM	energy management [message]

EMC	electromagnetic compatibility
EN	European Standard
EQAM	edge quadrature amplitude modulation [modulator]
ERMI	Edge Resource Manager Interface [Specification]
FCC	(U.S.) Federal Communications Commission
FDM	frequency division multiplexing
FDMA	frequency division multiple access
FEC	forward error correction
FFT	fast Fourier transform
FIR	finite impulse response
FR	fine ranging
ft	1) foot; 2) feet
FTTH	fibre to the home
GB	[Chinese] national standard (*guobiao*)
GB/T	[Chinese] recommended national standard (*guobiao tuijian*)
GF	Galois field
GHz	gigahertz
GT	guard time
HFC	hybrid fiber/coax
HRC	harmonic related carriers
Hz	hertz
I	in-phase
ICI	inter-carrier interference
I-CMTS	integrated cable modem termination system
ID	identifier
IDFT	inverse discrete Fourier transform
IEC	International Electrotechnical Commission
IETF	Internet Engineering Task Force
IFFT	inverse fast Fourier transform
IP	Internet protocol
IPDR	Internet protocol detail record
IPv4	Internet protocol version 4
IPv6	Internet protocol version 6
IR	initial ranging
IRC	incremental related carriers
ISI	inter-symbol interference
ISO	International Organization for Standardization
ITU	International Telecommunication Union
ITU-T	ITU Telecommunication Standardization Sector
kb	kilobit
kHz	kilohertz
L2VPN	layer 2 virtual private network
LAN	local area network
LDPC	low-density parity check
LFSR	linear feedback shift register
LLR	log-likelihood ratio
log	logarithm
LSB	least significant bit
LTE	long term evolution
MAC	media access control
MAP	Bandwidth Allocation Map
MB	message block
MC	message channel
M-CMTS	modular cable modem termination system
MER	modulation error ratio
MHz	megahertz
ms	millisecond
MSB	most significant bit
MSM	maximum scheduled minislots
Msym/s	megasymbols per second
MTC	multiple transmit channel [mode]
MULPI	MAC and upper layer protocols interface

NCP	next codeword pointer
ND	Neighbor Discovery
NMS	Network Management System
NPR	Noise Power Ratio
ns	nanosecond
NSI	Network Side Interface
NTSC	National Television System Committee
NULL	Null
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division mMultiple access
OOB	Out-Of-Band
OPT	Optional
OSSI	Operations Support System Interface
P	pilot
PAPR	Peak-to-Average Power Ratio
PDU	Protocol Data Unit
PER	Packet Error Ratio
PHY	PHYsical layer
PIE	Probe Information Element
pk-pk	peak-to-peak
Pkt	Packet
PLC	PHY Link Channel
PN	Pseudorandom Number
PNM	Proactive Network Maintenance
PRBS	Pseudo-Random Binary Sequence
Pre-eq	Pre-equalization
ps	picosecond
PSD	Power Spectral Density
Ptr	Pointer
Q	Quadrature
QAM	Quadrature Amplitude Modulation
QC-LDPC	Quasi-Cyclic Low-Density Parity Check
QoS	Quality of Service
QPSK	Quadrature Phase Shift Keying
RC	Raised Cosine
RCS	Receive Channel Set
REQ	REQuest
RF	Radio Frequency
RFC	Request For Comments
RFI	Radio Frequency Interface
RFoG	Radio Frequency over Glass
RMS	Root Mean Square
RP	Roll-off Period
R-S	Reed-Solomon
RX	receive
RX	receiver
s	second
SAC	Standardization Administration of the People's Republic of China
S-CDMA	Synchronous Code Division Multiple Access
SCN	Signal-to-Composite Noise [ratio]
SC-QAM	Single Carrier Quadrature Amplitude Modulation
SCTE	Society of Cable Telecommunications Engineers
SEC	SECurity
SID	Service IDentifier
SNMP	Simple Network Management Protocol
SNR	Signal-to-Noise Ratio
TCC	Transmit Channel Configuration
TCM	Trellis Coded Modulation
TCP	Transmission Control Protocol
TCS	Transmit Channel Set
TDM	Time Division Multiplexing
TDMA	Time Division Multiple Access

TEI	TDM emulation interface
TLV	Type/Length/Value
TS	Time Stamp
TV	TeleVision
TX	Transmit
TX	Transmitter
UCD	Upstream Channel Descriptor
UGHB	Under-Grant Hold Bandwidth
UGHU	Under-Grant Hold Number of Users
UID	Unique identifier
URL	Uniform Resource Locator
US	UpStream
VSA	Vector Signal Analyzer
XMOD	cross MODulation
XOR	exclusive OR

4 Requirements and Conventions

4.1 Requirements

Throughout the present document, the words that are used to define the significance of particular requirements are capitalized. These words are:

"SHALL" This word means that the item is an absolute requirement of the present document.

"SHALL NOT" This phrase means that the item is an absolute prohibition of the present document.

"SHOULD" This word means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before choosing a different course.

"SHOULD NOT" This phrase means that there may exist valid reasons in particular circumstances when the listed behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.

"MAY" This word means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

The present document defines many features and parameters, and a valid range for each parameter is usually specified. Equipment (CM and CMTS) requirements are always explicitly stated. Equipment will comply with all mandatory (shall and shall not) requirements to be considered compliant with the present document. Support of non-mandatory features and parameter values is optional.

4.2 Conventions

In the present document, the following convention applies any time a bit field is displayed in a figure. The bit field should be interpreted by reading the figure from left to right, then top to bottom, with the most-significant bit (MSB) being the first bit read, and the least-significant bit (LSB) being the last bit read.

5 Overview and Functional Assumptions

5.0 Cable Plant Characteristics

This clause describes the characteristics of a cable television plant, assumed to be for the purpose of operating a data-over-cable system.

The cable plants have very diverse physical topologies. These topologies range from fiber to the home node architectures as well as fiber nodes with many actives in cascade. The plant characteristics described in this clause covers the great majority of plant scenarios.

This clause is not a description of CMTS or CM parameters. The data-over-cable system shall be interoperable within the environment described in this clause.

Whenever a reference in this clause to frequency plans, or compatibility with other services, conflicts with any legal requirement for the area of operation, the latter shall take precedence. Any reference to National Television System Committee (NTSC) analog signals in 6 MHz channels does not imply that such signals are physically present.

5.1 Overview

The present document defines the PHY layer protocol of DOCSIS 3.1. It also describes the channel assumptions over which DOCSIS 3.1 systems are expected to operate.

DOCSIS 3.1 ultimate service goal of multi-gigabit per second in the downstream direction and gigabit per second in the upstream direction resulted in significant changes in the PHY layer approach compared to earlier DOCSIS versions in addition to changes on the cable network assumptions. DOCSIS 3.1 focuses on the eventual use of the entire spectrum resources available in cable environment by the CMTS and CM and on scalable cost effective techniques to achieve full spectrum use.

DOCSIS 3.1 assumes Orthogonal Frequency Division Multiplexing (OFDM) downstream signals and Orthogonal Frequency Division Multiple Access (OFDMA) upstream signals to achieve robust operation and provide more efficient use of the spectrum than previous DOCSIS versions. In order to reach the target service goal in the upstream direction, plant changes on the upstream/downstream spectrum split are expected. The DOCSIS 3.1 system will have options of several split configurations that can be exercised based on traffic demand, services offered and the capability of the cable plant. In the downstream direction, HFC plant spectrum extended beyond the current 1 002 MHz is expected.

The DOCSIS 3.0 systems and earlier versions are sometimes referred to in the present document as single carrier QAM (SC-QAM) systems in contrast to the multicarrier DOCSIS 3.1 OFDM/OFDMA system.

The OFDM downstream multicarrier system is composed of a large number of subcarriers that have either 25 kHz or 50 kHz spacing. These subcarriers are grouped into independently configurable OFDM channels each occupying a spectrum of up to 192 MHz in the downstream, totalling 7 680 (25 kHz) subcarriers or 3 840 (50 kHz) subcarriers; of which up to 7 600 (25 kHz) or 3 800 (50 kHz) active subcarriers span 190 MHz. The OFDMA upstream multicarrier system is also composed of either 25 kHz or 50 kHz subcarriers. In the upstream, the subcarriers are grouped into independently configurable OFDMA channels each occupying a spectrum of up to 96 MHz, totalling 3 840 (25 kHz) spaced subcarriers or 1 920 (50 kHz) spaced subcarriers. Many parameters of these channels can be independently configured thereby optimizing configuration based on channel conditions.

The cable topologies have been evolving to smaller node architectures with fewer amplifiers in cascade. This natural cable network evolution and gradual reduction in node sizes bring a corresponding improvement in channel conditions. A DOCSIS 3.1 goal is to leverage the expected network improvements and achieve higher efficiency levels corresponding to improvement in channel conditions. In the downstream and in the upstream directions, profiles can be defined to match the transmission configuration to the channel conditions with greater granularity. DOCSIS 3.1 technology is able to operate in classic cable network topologies, but those networks may not be able to achieve full capabilities of DOCSIS 3.1 bandwidth efficiencies.

An assumption in the topology configuration of DOCSIS 3.1 is that the CM is predominantly placed in a gateway architecture configuration. Specifically the CM is located either at the drop-home boundary or after the first two-way splitter within the customer premises. This configuration limits potential attenuation within the home environment. This reduced attenuation is leveraged to enable higher efficiencies intended in DOCSIS 3.1.

Another intent of DOCSIS 3.1 is to leverage the granularity from multiple narrowband subcarriers to exclude unwanted spectrum regions from transmission so that interferers can be avoided. Leveraging the same mechanism, coexistence with existing systems can be implemented by carving out a portion of the spectrum to allow for another transmission to co-exist.

It is expected that CMs and CMTS are able to operate in multiple modes. A mode could be pure Single Carrier Quadrature Amplitude Modulation (SC-QAM), pure OFDM transmission or a combination of the two. This flexibility enables a smooth transition between SC-QAM and OFDM systems.

DOCSIS 3.1 uses Low-Density-parity-Check (LDPC) Forward Error Correction (FEC) coding instead of the Reed-Solomon used in DOCSIS 3.0 and earlier versions. Long FEC codeword types are defined in the upstream and downstream to optimize efficiency. LDPC FEC along with frequency and time interleaving is used to provide robustness against narrowband interferers and burst events.

In several instances equivalent characterization and metrics to those used in DOCSIS 3.0 and earlier versions have been devised to facilitate comparison of requirements among different versions of the specification.

The DOCSIS 3.1 suite of specifications includes considerations to improve and optimize operation under special modes. One mode is a light sleep mode that is introduced to minimize CM power consumption. Also, some features are defined primarily operation in fiber to the home (FTTH) network architectures configured for RFoG.

5.2 Functional Assumptions

5.2.1 Equipment Assumptions

5.2.1.1 Frequency Plan

In the downstream direction, the cable system is assumed to have a pass band with a lower edge of either 54 MHz, 87,5 MHz, 108 MHz or 258 MHz, and an upper edge that is implementation-dependent but is typically in the range of 550 to 1 002 MHz. Upper frequency edges extending to 1 218 MHz, 1 794 MHz and others are expected in future migrations of the plants. Within that pass band, NTSC analog television signals in 6 MHz channels are assumed present on the standard, HRC or IRC frequency plans of [2], as well as other narrowband and wideband digital signals.

In the upstream direction, the cable system may have a 5 - 42 MHz, 5 - 65 MHz, 5 - 85 MHz, 5 - 117 MHz, 5 - 204 MHz or pass bands with an upper band edge beyond 204 MHz. NTSC analog television signals in 6 MHz channels may be present, as well as other signals.

5.2.1.2 Compatibility with Other Services

The CM shall coexist with any services on the cable network.

The CMTS shall coexist with any services on the cable network.

In particular:

- The CMTS shall be interoperable in the cable spectrum assigned for CMTS and CM interoperation while the balance of the cable spectrum is occupied by any combination of television and other signals.
- The CM shall be interoperable in the cable spectrum assigned for CMTS and CM interoperation while the balance of the cable spectrum is occupied by any combination of television and other signals.
- The CMTS shall not cause harmful interference to any other services that are assigned to the cable network in spectrum outside of that allocated to the CMTS.
- The CM shall not cause harmful interference to any other services that are assigned to the cable network in spectrum outside of that allocated to the CM.

Harmful interference is understood as:

- no measurable degradation (highest level of compatibility);

- no degradation below the perceptible level of impairments for all services (standard or medium level of compatibility); or
- no degradation below the minimal standards accepted by the industry (for example, FCC for analog video services) or other service provider (minimal level of compatibility).

5.2.1.3 Fault Isolation Impact on Other Users

As CMTS transmissions are on a shared-media, point-to-multipoint system, fault-isolation procedures should take into account the potential harmful impact of faults and fault-isolation procedures on numerous users of the data-over-cable, video and other services.

For the interpretation of harmful impact, see clause 5.2.1.2.

5.2.1.4 Cable System Terminal Devices

The CM is expected to meet and preferably exceed all applicable regulations for Cable System Termination Devices and Cable Ready Consumer Equipment as defined in FCC Part 15 [14] and Part 76 [15]. None of these specific requirements may be used to relax any of the specifications contained elsewhere within the present document.

5.2.2 RF Channel Assumptions

5.2.2.0 Cable Networks Interoperability

The data-over-cable system, configured with at least one set of defined physical-layer parameters (e.g. modulation, interleaver depth, etc.) from the range of configurable settings described in the present document, is expected to be interoperable on cable networks having characteristics defined in this clause. This is accomplished in such a manner that the forward error correction provides for equivalent operation in a cable system both with and without the impaired channel characteristics described below.

5.2.2.1 Transmission Downstream

The RF channel transmission characteristics of the cable network in the downstream direction are described in table 5.1. These numbers assume total average power of a digital signal in a 192 MHz channel bandwidth for subcarrier levels unless indicated otherwise. For impairment levels, the numbers in table 5.1 assume average power in a bandwidth in which the impairment levels are measured in a standard manner for cable TV systems. For analog signal levels, the numbers in table 5.1 assume peak envelope power in a 6 MHz channel bandwidth. All conditions are present concurrently. It is expected that the HFC plant will have better conditions for DOCSIS 3.1 to provide the higher throughput and capacities anticipated.

Table 5.1: Typical Downstream RF Channel Transmission Characteristics

Parameter	Value
Frequency range	Cable system normal downstream operating range is from 54 MHz to 1 002 MHz. Extended operating ranges include lower downstream edges of 108 MHz and 258 MHz and upper downstream edges of 1 218 MHz and 1 794 MHz.
RF channel spacing (design bandwidth)	24 to 192 MHz
One way transit delay from headend to most distant customer	≤ 0,400 ms (typically much less)
Signal to Composite Noise Ratio	≥ 35 dB
Carrier-to-Composite triple beat distortion ratio	Not less than 41 dB
Carrier-to-Composite second order distortion ratio	Not less than 41 dB
Carrier-to-Cross-modulation ratio	Not less than 41 dB
Carrier-to-any other discrete interference (ingress)	Not less than 41 dB
Maximum amplitude variation across the 6 MHz channel (digital channels)	≤ 1,74 dB pk-pk/6 MHz
Group Delay Variation	≤ 113 ns over 24 MHz
Micro-reflections bound for dominant single echo	-20 dBc for echos ≤ 0,5 µs -25 dBc for echos ≤ 1,0 µs -30 dBc for echos ≤ 1,5 µs -35 dBc for echos > 2,0 µs -40 dBc for echos > 3,0 µs -45 dBc for echos > 4,5 µs -50 dBc for echos > 5,0 µs
Carrier hum modulation	Not greater than -30 dBc (3 %)
Maximum analog video carrier level at the CM input	17 dBmV
Maximum number of analog carriers	121
NOTE: Cascaded group delay could possibly exceed the ≤ 113 ns value within approximately 30 MHz above the downstream spectrum's lower band edge, depending on cascade depth, diplex filter design, and actual band split.	

5.2.2.2 Transmission Upstream

5.2.2.2.0 RF Channel Transmission Characteristics

The RF channel transmission characteristics of the cable network in the upstream direction are described in table 5.2. No combination of the following parameters will exceed any stated interface limit defined elsewhere in the present document. Transmission is from the CM output at the customer location to the headend.

Table 5.2: Typical Upstream RF Channel Transmission Characteristics

Parameter	Value
Frequency range	Cable standard upstream frequency range is from a lower band-edge of 5 MHz to upper band-edges to 42 MHz and 65 MHz. Extended upstream frequency ranges include upper upstream band-edges of 85 MHz, 117 MHz and 204 MHz.
One way transit delay from most distant customer to headend.	≤ 0,400 ms (typically much less)
Carrier-to-interference plus ingress (the sum of noise, distortion, common-path distortion and cross modulation and the sum of discrete and broadband ingress signals, impulse noise excluded) ratio	Not less than 25 dB
Carrier hum modulation	Not greater than -26 dBc (5,0 %)
Maximum amplitude variation across the 6 MHz channel (digital channels)	≤ 2,78 dB pk-pk/6 MHz
Group Delay Variation	≤ 163 ns over 24 MHz
Micro-reflections bound for dominant single echo	-16 dBc for echos ≤ 0,5 µs -22 dBc for echos ≤ 1,0 µs -29 dBc for echos ≤ 1,5 µs -35 dBc for echos > 2,0 µs -42 dBc for echos > 3,0 µs -51 dBc for echos > 4,5 µs
Seasonal and diurnal reverse gain (loss) variation	Not greater than 14 dB min to max
NOTE: Cascaded group delay could possibly exceed the ≤163 ns value within approximately 10 MHz of the upstream spectrum's lower and upper band edges, depending on cascade depth, diplex filter design, and actual band split.	

5.2.2.2.1 Availability

Cable network availability is typically greater than 99,9 %.

5.2.3 Transmission Levels

The nominal power level of the upstream CM signal(s) will be as low as possible to achieve the required margin above noise and interference. Uniform power loading per unit bandwidth is commonly followed in setting upstream signal levels, with specific levels established by the cable network operator to achieve the required carrier-to-noise and carrier-to-interference ratios.

5.2.4 Frequency Inversion

There will be no frequency inversion in the transmission path in either the downstream or the upstream directions, i.e. a positive change in frequency at the input to the cable network will result in a positive change in frequency at the output.

6 PHY Sublayer for SC-QAM

6.1 Scope

This clause applies to cases where a DOCSIS 3.1 CM or CMTS is operating with SC-QAM operation only, with no OFDM/OFDMA operation and for the SC-QAM channels with simultaneous operation of SC-QAM and OFDM/OFDMA channels unless otherwise noted. As such, it represents backward compatibility requirements when operating with DOCSIS 3.0 systems or with the new DOCSIS 3.1 PHY disabled. It also applies only to the first technology option referred to in clause 1; for the second option refer to Annex C; and for the third option refer to Annex D.

The present document defines the electrical characteristics and signal processing operations for a CM and CMTS. It is the intent of the present document to define an interoperable CM and CMTS such that any implementation of a CM can work with any CMTS. It is not the intent of the present document to imply any specific implementation.

As the requirements for a DOCSIS 3.1 CM and CMTS are largely unchanged relative to DOCSIS 3.0 devices for SC-QAM operation, this clause is comprised primarily of references to the appropriate DOCSIS 3.0 specification sections for the specific requirements for a DOCSIS 3.1 CM and CMTS, as well as any deltas relative to those requirements (the primary difference being that a DOCSIS 3.1 CM and CMTS are required to support a minimum of 24 downstream and 8 upstream channels instead of 4 downstream and 4 upstream as in DOCSIS 3.0 devices).

A DOCSIS 3.1 CM shall comply with the referenced requirements in the PHYv3.0 and DRFI specifications noted in this clause, with the exception of any deltas called out in this clause (which will be identified with separate requirement statements). A DOCSIS 3.1 CMTS shall comply with the referenced requirements in the PHYv3.0 and DRFI specifications noted in this clause, with the exception of any deltas called out in this clause (which will be identified with separate requirement statements).

6.2 Upstream Transmit and Receive

6.2.1 Overview

See clause 6.2.1 of [5], with the exceptions noted below.

A CM shall support at least eight (8) active upstream channels (which are referred to as the Transmit Channel Set for that CM).

A CMTS shall support at least eight (8) active upstream channels.

A CMTS may support S-CDMA mode. If a CMTS implements S-CDMA mode, the CMTS shall comply with S-CDMA requirements defined in [5].

A CM may support S-CDMA mode. If a CM implements S-CDMA mode, the CM shall comply with S-CDMA requirements defined in [5].

6.2.2 Signal Processing Requirements

See clause 6.2.2 of [5].

6.2.3 Modulation Formats

See clause 6.2.3 of [5].

6.2.4 R-S Encode

See clause 6.2.4 of [5].

6.2.5 Upstream R-S Frame Structure (Multiple Transmit Enabled)

See clause 6.2.5 of [5].

6.2.6 Upstream R-S Frame Structure (Multiple Transmit Disabled)

See clause 6.2.6 of [5].

6.2.7 TDMA Byte Interleaver

See clause 6.2.7 of [5].

6.2.8 Scrambler (Randomizer)

See clause 6.2.8 of [5].

6.2.9 TCM Encoder

See clause 6.2.9 of [5].

6.2.10 Preamble Prepend

See clause 6.2.10 of [5].

6.2.11 Modulation Rates

See clause 6.2.11 of [5].

6.2.12 S-CDMA Framer and Interleaver

See clause 6.2.12 of [5].

6.2.13 S-CDMA Framer

See clause 6.2.13 of [5].

6.2.14 Symbol Mapping

See clause 6.2.14 of [5].

6.2.15 S-CDMA Spreader

See clause 6.2.15 of [5].

6.2.16 Transmit Pre-Equalizer

See clause 6.2.16 of [5].

6.2.17 Spectral Shaping

See clause 6.2.17 of [5].

6.2.18 Relative Processing Delays

See clause 6.2.18 of [5].

6.2.19 Transmit Power Requirements

Applies only to cases where a DOCSIS 3.1 CM is operating with a DOCSIS 3.0 CMTS or where a DOCSIS 3.1 CMTS is operating with a DOCSIS 3.0 CM.

See clause 6.2.19 of [5].

6.2.20 Burst Profiles

See clause 6.2.20 of [5].

6.2.21 Burst Timing Convention

See clause 6.2.21 of [5].

6.2.22 Fidelity Requirements

Applies only to cases where a DOCSIS 3.1 CM is operating with a DOCSIS 3.0 CMTS or where a DOCSIS 3.1 CMTS is operating with a DOCSIS 3.0 CM.

See clause 6.2.22 of [5].

6.2.23 Upstream Demodulator Input Power Characteristics

See clause 6.2.23 of [5].

6.2.24 Upstream Electrical Output from the CM

Applies only to cases where a DOCSIS 3.1 CM is operating with a DOCSIS 3.0 CMTS or where a DOCSIS 3.1 CMTS is operating with a DOCSIS 3.0 CM.

See clause 6.2.24 of [5].

6.2.25 Upstream CM Transmitter Capabilities

Applies only to cases where a DOCSIS 3.1 CM is operating with a DOCSIS 3.0 CMTS or where a DOCSIS 3.1 CMTS is operating with a DOCSIS 3.0 CM.

See clause 6.2.25 of [5].

6.3 Downstream Transmit

6.3.1 Downstream Protocol

See clause 6.3.1 of [3].

6.3.2 Spectrum Format

See clause 6.3.2 of [3].

6.3.3 Scaleable Interleaving to Support Video and High-Speed Data Services

See clause 6.3.3 of [3].

6.3.4 Downstream Frequency Plan

See clause 6.3.4 of [3].

6.3.5 DRFI Output Electrical

Applies only the case where a DOCSIS 3.1 device is operating in DOCSIS 3.0 mode only.

For legacy, SC-QAMs, the EQAM and CMTS shall support the electrical output requirements specified in the following clauses and tables of [3]:

- clause 6.3.5, DRFI Output Electrical;
- clause 6.3.5.1, CMTS or EQAM Output Electrical;
- table 6-3, RF Output Electrical Requirements;

- clause 6.3.5.1.1, Power per Channel CMTS or EQAM;
- table 6-4, DRFI Device Output Power;
- clause 6.3.5.1.2, Independence of Individual Channels within the Multiple Channels on a Single RF Port;
- clause 6.3.5.1.3, Out-of-Band Noise and Spurious Requirements for CMTS or EQAM;
- table 6-5, EQAM or CMTS Output Out-of-Band Noise and Spurious Emissions Requirements for $N \leq 8$;
- table 6-6, EQAM or CMTS Output Out-of-Band Noise and Spurious Emissions Requirements $N \geq 9$ and $N' \geq N/4$;
- table 6-7, EQAM or CMTS Output Out-of-Band Noise and Spurious Emissions Requirements $N \geq 9$ and $N' < N/4$.

6.3.6 CMTS or EQAM Clock Generation

Applies only the case where a DOCSIS 3.1 CMTS is operating in a DOCSIS 3.0 mode only.

See clause 6.3.6 of [3].

6.3.7 Downstream Symbol Clock Jitter for Synchronous Operation

Applies only the case where a DOCSIS 3.1 CMTS is operating in a DOCSIS 3.0 mode only.

See clause 6.3.7 of [3].

6.3.8 Downstream Symbol Clock Drift for Synchronous Operation

Applies only the case where a DOCSIS 3.1 CMTS is operating in a DOCSIS 3.0 mode only.

See clause 6.3.8 of [3].

6.3.9 Timestamp Jitter

See clause 6.3.9 of [3].

6.4 Downstream Receive

6.4.1 Downstream Protocol and Interleaving Support

See clause 6.3.1 of [5].

6.4.2 Downstream Electrical Input to the CM

See clause 6.3.2 of [5], with the exceptions noted below.

A CM shall support at least 24 active downstream channels.

A CMTS shall support at least 24 active downstream channels.

6.4.3 CM BER Performance

See clause 6.3.3 of [5].

6.4.4 Downstream Multiple Receiver Capabilities

See clause 6.3.4 of [5].

6.4.5 Non-Synchronous DS Channel Support

Applies only to the case where a DOCSIS 3.1 CM operating with a DOCSIS 3.0 CMTS.

See clause 6.3.5 of [5].

7 PHY Sublayer for OFDM

7.1 Scope

The present document defines the electrical characteristics and signal processing operations for a cable modem (CM) and Cable Modem Termination System (CMTS). It is the intent of the present document to define an interoperable CM and CMTS such that any implementation of a CM can work with any CMTS. It is not the intent of the present document to imply any specific implementation.

7.2 Upstream and Downstream Frequency Plan

7.2.0 System Spectrum Requirement

The following spectrum definitions are based on the system requirement that the downstream transmission frequencies always reside above the upstream transmission frequencies in the cable plant.

7.2.1 Downstream CM Spectrum

The CM shall support a minimum of two independently configurable OFDM channels each occupying a spectrum of up to 192 MHz in the downstream.

The CM shall support a downstream upper band edge of 1,218 GHz.

The CM may support a downstream upper band edge of 1,794 GHz.

The CM shall support a downstream lower band edge of 258 MHz.

The CM should support a downstream lower band edge of 108 MHz when the CM is configured to use an upstream upper band edge of 85 MHz or less.

7.2.2 Downstream CMTS Spectrum

The CMTS shall support a minimum of two independently configurable OFDM channels each occupying a spectrum of up to 192 MHz in the downstream.

The CMTS shall support a downstream upper band edge of 1,218 GHz.

The CMTS may support a downstream upper band edge of 1,794 GHz.

The CMTS shall support a downstream lower band edge of 258 MHz.

The CMTS should support a downstream lower band edge of 108 MHz.

7.2.3 Upstream CM Spectrum

The CM shall support a minimum of two independently configurable OFDMA channels each occupying a spectrum of up to 96 MHz in the upstream.

The CM may support more than two independently configurable OFDMA channels each occupying a spectrum of up to 96 MHz in the upstream.

The CM modulator shall support upstream transmissions from 5 to at least 204 MHz and agile placement of the OFDMA channels within that range.

Individual CM implementations may limit the spectrum over which the CM is able to transmit upstream signals. As a result, in order to be compliant with the present document a CM shall support one or more of the following upstream upper band edges, as long as one of the upstream upper band edges supported is 85 MHz or greater: 42 MHz; 65 MHz, 85 MHz, 117 MHz, and/or 204 MHz.

The CM shall be configurable to operate with any supported upstream upper band edge. The nature and operation of this configurability is vendor-specific. Possible forms of configurability include a hardware switch on the modem housing, a software-controlled diplex filter responsive to OSSI commands, or other forms.

The CM may support additional spectrum beyond 204 MHz for the upstream.

The CM shall not cause harmful interference to any downstream signals that might exist above its configured upstream upper band edge.

The CM shall be capable of transmitting 192 MHz of active channels when operating with the 204 MHz upstream upper band edge.

In DOCSIS 3.1 upstream mode, the CM shall be capable of transmitting OFDMA channels and legacy SC-QAM channels at the same time (as controlled by the CMTS). In all cases the CM is not required to transmit legacy SC-QAM channels above a frequency of 85 MHz.

7.2.4 Upstream CMTS Spectrum

The CMTS shall support a minimum of two independently configurable OFDMA channels each occupying a spectrum of up to 96 MHz in the upstream.

The CMTS may support more than two independently configurable OFDMA channels each occupying a spectrum of up to 96 MHz in the upstream.

The CMTS shall support upstream transmissions from 5 to at least 204 MHz and agile placement of the OFDMA blocks within that range.

The CMTS may support additional spectrum beyond 204 MHz for the upstream.

The CMTS shall capable of receiving 192 MHz of active channels when operating with the 204 MHz upstream upper band edge. In DOCSIS 3.1 upstream mode the CM is capable of transmitting OFDMA channels and legacy SC-QAM channels at the same time (as controlled by the CMTS). In all cases the CMTS shall not configure the CM to transmit legacy SC-QAM channels above a frequency of 85 MHz.

7.2.5 Channel Band Rules

7.2.5.0 OFDM/OFDMA Channel Planning

During OFDM/OFDMA channel planning, the following rules are to be observed to ensure proper operation of DOCSIS 3.1 CMTS and CM.

The CMTS shall ensure that the configured OFDM/OFDMA channels are aligned with the rules specified in clauses 7.2.5.2 and 7.2.5.3.

7.2.5.1 Downstream Channel Bandwidth Rules

The CMTS shall ensure that the encompassed spectrum of a 192 MHz downstream OFDM channel does not exceed 190 MHz. Therefore the CMTS shall ensure that the number of contiguous active subcarriers in a downstream OFDM channel does not exceed 3 800 for 4K FFT and 7 600 for 8K FFT. When configured for 4K FFT, the CMTS shall only use subcarriers in the range $148 \leq k \leq 3947$, where *k* is the spectral index of the subcarrier in the IDFT equation defining the OFDM signal. When configured for 8K FFT, the CMTS shall only use subcarriers in the range $296 \leq k \leq 7895$, where *k* is the spectral index of the subcarrier in the IDFT equation defining the OFDM signal.

The CMTS shall ensure that there is at least 1 MHz of exclusion band between the spectral edge of a legacy SC-QAM channel and the centre frequency of the nearest OFDM subcarrier. This SC-QAM channel may be external to the OFDM channel or may be embedded within the OFDM channel.

The CMTS shall also ensure that there is at least 2 MHz exclusion band between any two adjacent asynchronous OFDM channels. In other words the CMTS shall ensure that the frequency separation between the highest frequency active subcarrier of one OFDM channel and the lowest frequency active subcarrier of the adjacent asynchronous OFDM channel is not less than 2 MHz.

Such an exclusion band is not needed if the two OFDM channels are fully synchronous. The term synchronous here implies that both OFDM channels have the same FFT length, the same cyclic prefix, and are synchronized in time, frequency and phase. For example, the CMTS may use a single 16K FFT with a sample rate of 409,6 MHz to construct two 8K FFT OFDM channels each with sample rate 204,8 MHz. The use of a single FFT guarantees that all synchronization criteria are met. The CMTS may place 15 200 contiguous active subcarriers, with an encompassed spectrum of 380 MHz, anywhere within this 16K FFT. These 15 200 subcarriers may be partitioned equally between two adjacent downstream 8K FFT OFDM channels.

7.2.5.2 Downstream Exclusion Band Rules

The CM and CMTS are not expected to meet performance and fidelity requirements when the system configuration does not comply with the downstream exclusion band rules listed below. These rules apply to each OFDM channel and also to the composite downstream inclusive of OFDM and non-OFDM channels:

- There has to be at least one contiguous modulated OFDM bandwidth of 22 MHz or greater.
- Exclusion bands separate contiguous modulation bands.
- The minimum contiguous modulation band has to be 2 MHz.
- Exclusion bands and individually excluded subcarriers are common to all downstream profiles.
- Exclusion bands are a minimum of 1 MHz but increment above 1 MHz by granularity of individual subcarrier (25 kHz for 8k FFT and 50 kHz for 4K FFT).
- The ONLY exception to the above is for exclusion bands that are allowed to occupy the following frequency ranges in alignment with FCC regulations.
    - 121,400 MHz to 121,600 MHz
    - 156,750 MHz to 156,850 MHz
    - 242,950 MHz to 243,050 MHz
    - 405,925 MHz to 406,176 MHz

Unique spurious emissions requirements exist for these bands separate from the general exclusion bands requirements.

- Exclusion bands plus individually excluded subcarriers are limited to 20 % or less of spanned modulation spectrum, where the spanned modulation spectrum is defined as: frequency of maximum active subcarrier - frequency of minimum active subcarrier.
- The number of individually excluded subcarriers is limited by the following:
    - The total spectrum of individually excluded subcarriers cannot exceed 5 % of any contiguous modulation spectrum.
    - The total spectrum of individually excluded subcarriers cannot exceed 5 % of a 6 MHz moving window across the contiguous modulation spectrum.
    - The total spectrum of individually excluded subcarriers cannot exceed 20 % of a 1 MHz moving window across the contiguous modulation spectrum.
- The 6 MHz of contiguous spectrum reserved for the PLC cannot have any exclusion bands or excluded subcarrier.

7.2.5.3 Upstream Channel Bandwidth Rules

The CMTS shall ensure that the encompassed spectrum of a 96 MHz upstream OFDMA channel does not exceed 95 MHz. Therefore the number of contiguous active subcarriers in an upstream OFDMA channel shall not exceed 1 900 for 2K FFT and 3 800 for 4K FFT. When configured for 2K FFT, the CMTS shall only use subcarriers in the range $74 \leq k \leq 1\,973$, where k is the spectral index of the subcarrier in the IDFT equation defining the OFDMA signal. When configured for 4K FFT, the CMTS shall only use subcarriers in the range $148 \leq k \leq 3\,947$, where k is the spectral index of the subcarrier in the IDFT equation defining the OFDMA signal.

7.2.5.4 Upstream Exclusion Band Rules

These rules apply to the upstream exclusion band:

- Excluded subcarriers will only occur between minislots.
- Subcarrier exclusions between minislots can be any integer number of subcarriers. There is no minimum subcarrier exclusion restriction.

7.3 OFDM Numerology

7.3.1 Downstream OFDM Numerology

DOCSIS 3.1 uses OFDM for downstream modulation.

Two modes of operation are defined for the downstream: 4K FFT and 8K FFT modes for a sampling rate of 204,8 MHz. These are described in clause 7.5.7.

Table 7.1 summarizes the numerical values for the downstream OFDM parameters; a more detailed description of the parameters is given in the sections which follow.

Table 7.1: Downstream OFDM parameters

Parameter	4K mode	8K mode
Downstream master clock frequency	10,24 MHz	
Downstream Sampling Rate (fs)	204,8 MHz	
Downstream Elementary Period (Tsd)	1/(204,8 MHz)	
Channel bandwidths	24 MHz … 192 MHz	
IDFT size	4 096	8 192
Subcarrier spacing	50 kHz	25 kHz
FFT duration (Useful symbol duration) (Tu)	20 µs	40 µs
Maximum number of active subcarriers in signal (192 MHz channel) Values refer to 190 MHz of used subcarriers.	3 800	7 600
Maximum spacing between first and last active subcarrier	190 MHz	
Cyclic Prefix	0,9375 µs (192 * T_{sd}) 1,25 µs (256 * T_{sd}) 2,5 µs (512 * T_{sd}) 3,75 µs (768 * T_{sd}) 5 µs (1 024 * T_{sd})	
Windowing	Tukey raised cosine window, embedded into cyclic prefix 0 µs (0 * T_{sd}) 0,3125 µs (64 * T_{sd}) 0,625 µs (128 * T_{sd}) 0,9375 µs (192 * T_{sd}) 1,25 µs (256 * $T_{sd)}$	

The downstream OFDM channel bandwidth can vary from 24 MHz to 192 MHz. Smaller bandwidths than 192 MHz are achieved by zero-valuing the subcarriers prior to the IDFT, i.e. by adjusting the equivalent number of active subcarriers while maintaining the same subcarrier spacing of 25 kHz or 50 kHz.

7.3.2 Upstream OFDMA Numerology

DOCSIS 3.1 uses OFDMA (orthogonal frequency-division multiple access) for upstream modulation. OFDMA is a multi-user version of OFDM, and assigns subsets of subcarriers to individual CMs. The upstream OFDMA parameters are derived from the downstream parameters, and are summarized in table 7.2. A more detailed description of the parameters is given in the clauses that follow.

Table 7.2: Upstream OFDMA Parameters

Parameter	2k Mode	4k Mode
Upstream Sampling Rate (f_{su})	102,4 MHz	
Upstream Elementary Period Rate (T_{su})	1/102,4 MHz	
Channel bandwidths	10 MHz, …, 96 MHz	6,4 MHz, …, 96 MHz
IDFT size (depending on channel bandwidth)	2 048	4 096
Subcarrier spacing	50 kHz	25 kHz
FFT duration (Useful symbol duration) (T_u)	20 µs	40 µs
Maximum number of active subcarriers in signal (96 MHz channel) Values refer to 95 MHz of active subcarriers	1 900	3 800
Upstream Cyclic Prefix	0,9375 µs (96 * T_{su}) 1,25 µs (128 * T_{su}) 1,5625 µs (160 * T_{su}) 1,875 µs (192 * T_{su}) 2,1875 µs (224 * T_{su}) 2,5 µs (256 * T_{su}) 2,8125 µs (288 * T_{su}) 3,125 µs (320 * T_{su}) 3,75 µs (384 * T_{su}) 5,0 µs (512 *T_{su}) 6,25 µs (640 *T_{su})	
Upstream window size	Tukey raised cosine window, embedded into cyclic prefix 0 µs (0 * T_{su}) 0,3125 µs (32 * T_{su}) 0,625 µs (64 * T_{su}) 0,9375 µs (96 * T_{su}) 1,25 µs (128 * T_{su}) 1,5625 µs (160 * T_{su}) 1,875 µs (192 * T_{su}) 2,1875 µs (224 * T_{su})	

7.3.3 Subcarrier Clocking

The "locking" of subcarrier "clock and carrier" are defined and characterized by the following rules:

- Each OFDM symbol is defined with an FFT duration (equal to subcarrier clock period) of nominally 20 µs or 40 µs. For each OFDM symbol, the subcarrier clock period (µs) may vary from nominal with limits defined in clause 7.5.3.
- The number of cycles of each subcarrier generated by the CMTS during one period of the subcarrier clock (for each OFDM symbol) shall be an integer number.

The CMTS subcarrier clock shall be synchronous with the 10,24 MHz Master Clock defined by:

- Subcarrier clock frequency = (M/N)*Master Clock frequency where M = 20 or 40, and N = 8 192.
- The limitation on the variation from nominal of the subcarrier clock frequency at the output connector is defined in clause 7.5.3.
- Each OFDM symbol has a cyclic prefix which is an integer multiple of 1/64th, of the subcarrier clock period.
- Each OFDM symbol duration is the sum of one subcarrier clock period and the cyclic prefix duration.
- The number of cycles of each subcarrier generated by the CMTS during the OFDM symbol duration (of each symbol) shall be K+K*L/64, where K is an integer related to the subcarrier index and frequency upconversion of the OFDM channel, and L is an integer related to the cyclic prefix. (K is an integer related to the subcarrier index and increases by 1 for each subcarrier).
- The phase of each subcarrier within one OFDM symbol is the same, when each is assigned the same constellation point (I + jQ), relative to the Reference Time of the OFDM symbol. There is nominally no change in phase on each subcarrier for every cycle of 64 OFDM symbols, when both are assigned the same I + jQ, and referenced to the Reference Time of their respective OFDM symbol.

7.4 Upstream Transmit and Receive

7.4.1 Overview

This clause specifies the upstream transmission electrical and signal processing requirements for the transmission of OFDM modulated RF signals from the CM to the CMTS.

7.4.2 Signal Processing Requirements

7.4.2.0 OFDMA Frames

Upstream transmission uses OFDMA frames. Each OFDMA frame is comprised of a configurable number of OFDM symbols, *K*. Several transmitters may share the same OFDMA frame by transmitting data and pilots on allocated sub-carriers of the OFDMA frame. There are several pilot patterns as described in clause 7.4.17.

The structure of an OFDMA frame is depicted in figure 7.1. The upstream spectrum is divided into groups of sub-carriers called minislots. Minislots have dedicated sub-carriers, all with the same modulation order ("bit loading"). A CM is allocated to transmit one or more minislots in a Transmission Burst. The modulation order of a minislot, as well as the pilot pattern to use may change between different transmission bursts and are determined by a transmission profile.

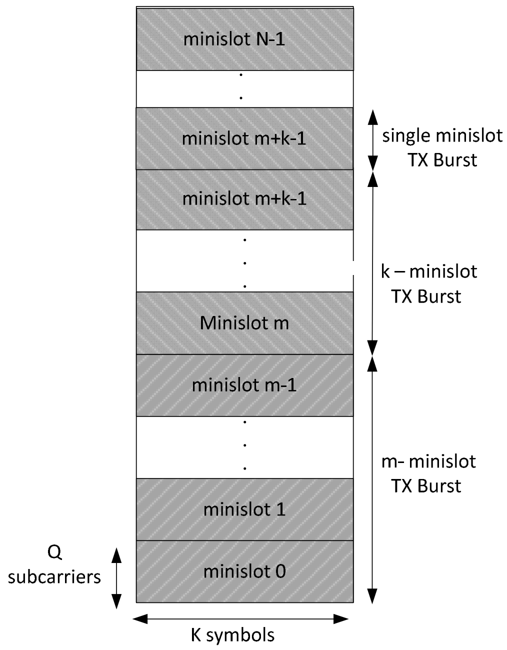

Figure 7.1: OFDMA Frame Structure

Serial data signals received from the PHY-MAC Convergence Layer are received and processed by the PHY as illustrated in figure 7.2. This process yields a transmission burst of a single or multiple OFDMA minislots, as allocated by the PHY-MAC Convergence Layer. Each minislot is comprised of pilots, complementary pilots, and data subcarriers, as described in clause 8.2.3. Subcarriers that are not used for data or pilots are set to zero.

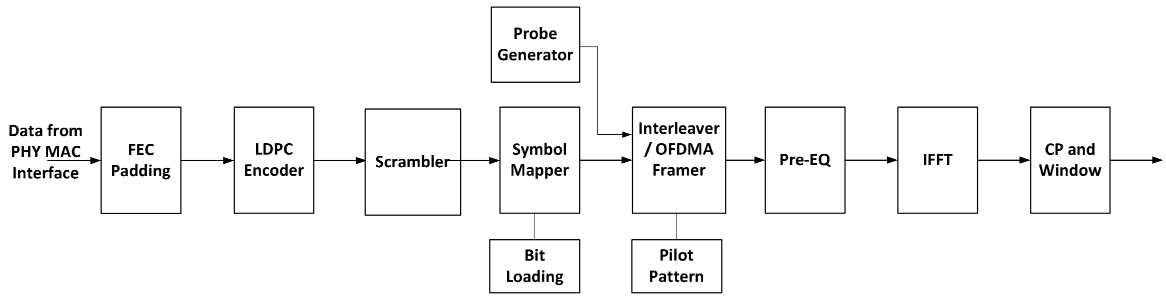

Figure 7.2: Upstream transmitter block diagram

This clause briefly describes the process and provides links to the specific requirements for each process described in the present document.

7.4.2.1 Framing

Figure 7.3 describes how the received bits from the PHY-MAC Convergence layers are framed before being converted into constellation symbols. The number of FEC codewords, corresponding codeword lengths and zero-padding bits are calculated by the PHY-MAC Convergence layer as described in clause 7.4.4.1 according to the allocation of minislot and the profile received by the grant message.

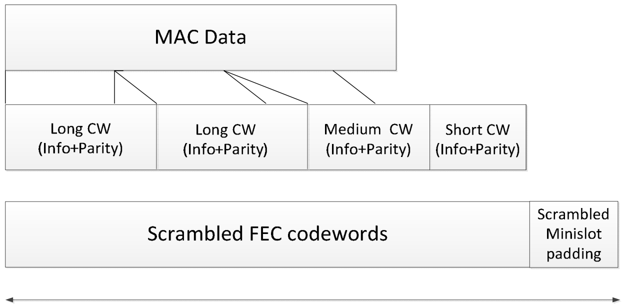

Figure 7.3: Upstream Transmitter Block Diagram

7.4.2.2 Forward Error Correction Encoding

Data received from the PHY-MAC Convergence layer interface, along with the FEC padding, is LDPC encoded. The upstream has three LDPC codes: a long, medium, and short FEC code, as described in clause 7.4.4. Prior to encoding, the transmitter is to decide on the configuration of the codewords as described in clause 7.4.4.1 and codeword shortening as described in clause 7.4.4.3. If required, zero-padding has to be applied to complete the number of minislots in the grant.

7.4.2.3 Randomizer and Symbol Mapper

The encoded bits are then randomized (scrambled) using the PRBS scrambler as specified in clause 7.4.5. The scrambler output bits are converted into constellation symbols according to the corresponding modulation order of the minislot. All subcarriers of a given type (Pilots, Complementary Pilots, Data subcarriers) in a minislot have the same modulation order. The Symbol Mapper is described in clause 7.4.7.

7.4.2.4 OFDMA framer and Interleaver

Constellation symbols then enter the OFDMA framer and Interleaver block. The OFDMA framer adds pilots according to the pilot pattern associated with the transmission burst minislot. The constellation symbols are written to subcarriers associated with the transmission burst minislots and are then interleaved in time and frequency as described in clause 7.4.6.

7.4.2.5 Pre-Equalization

The upstream transmitter applies pre-equalization as described in clause 7.4.18 in order to pre-distort the transmitted signals according to coefficients received from the CMTS to compensate for the channel response.

7.4.2.6 IFFT Transformation

In this stage each pre-equalized symbol is transformed into the time domain using the IFFT block. IFFT inputs that are not used (that is, that do not correspond to any of the minislots used by the transmission burst) are set to zero.

The transmitter converts the output of the IFFT from parallel to serial and performs cyclic prefix addition and Windowing in the time domain.

7.4.2.7 Cyclic Prefix and Windowing

A segment at the end of the IFFT output is prepended to the IFFT output; this is referred to as the Cyclic Prefix (CP) of the OFDM symbol. For windowing purposes, another segment at the start of the IFFT output is appended to the end of the IFFT output - the roll-off period (RP).

The addition of a cyclic prefix enables the receiver to overcome the effects of inter-symbol-interference caused by micro-reflections in the channel. Windowing maximizes channel capacity by sharpening the edges of the spectrum of the OFDM signal. Spectral edges occur at the two ends of the spectrum of the OFDM symbol, as well as at the ends of internal exclusion bands.

These topics are discussed in detail in clause 7.4.11.

7.4.3 Time and Frequency Synchronization

7.4.3.0 Timing and Frequency Synchronization Requirements

CM upstream frequency and timing of transmissions is based on downstream tracking, and in the case of timing, also based upon receiving and implementing timing adjustments from the CMTS. This clause describes the CM upstream transmission performance requirements on frequency and timing which are based upon tracking the downstream input to the CM, and implementing and operating upon commands from the CMTS.

7.4.3.1 Channel Frequency Accuracy

To support OFDMA upstream using legacy downstream, the CM shall provide upstream frequency accuracy, relative to the downstream reference, of ≤ ±20 Hz in at least Five Sigma of upstream grants, under all the downstream conditions detailed in clause 7.5.9.

Upstream frequency accuracy measurements are averaged over the duration of an upstream single frame grant.

7.4.3.2 Channel Timing Accuracy

For OFDMA upstream, regardless of what is used for the timing master, the ranging time offset will be given as described in [4].

Specifically, this timestamp has an integer portion of 10,24 MHz clocks. It also has an integer portion of counting 1/20ths duration of 10,24 MHz clock period (this integer portion counts up to 20), and then it has a 4 bit binary fractional portion so the CM's timing resolution shall be (1/10,24 MHz) x (1/20) x (1/16) = 305 ps.

The CMTS shall be able to send timing adjustment commands with a resolution of 305 ps or an integer submultiple of 305 ps.

The CM shall implement the OFDMA Timing Adjust to within ±10 ns. For example, the average error as measured at the CMTS over 35 s has to be within ±10 ns.

7.4.3.3 Modulation Timing Jitter

The CM shall implement the upstream timing so that the OFDMA clock timing error (with the mean error subtracted out) relative to the CMTS master clock as measured at the CMTS will be within ±10 ns in each burst measured within 35 s measurement duration.

This applies to the worst-case jitter and frequency drift specified for the CMTS Master clock and the CMTS downstream symbol clock in the requirements above.

The mean error is the result of the adjustment implemented by the CM as specified in clause 7.4.3.2.

7.4.3.4 Upstream Subcarrier Clock Frequency

The CM shall lock the Upstream Subcarrier Clock Frequency (25 kHz or 50 kHz) to the 10,24 MHz Master Clock via the Downstream Subcarrier Clock, subject to the carrier phase noise and frequency offset requirements in this clause. The subcarrier frequency is specified as locked to the 10,24 MHz reference and given with 25 kHz resolution (25 kHz = 10,24 MHz x 10/4 096).

Maximal subcarrier frequency offset depends on the downstream OFDM occupied bandwidth which determines the number and general distribution of Continuous Pilots. All upstream subcarrier frequency specifications assume a downstream input to the CM per clauses 7.5.9 and 7.5.3.

For OFDMA upstream channels, the CM's upstream subcarrier frequency offset for each subcarrier (relative to the Master Clock reference) with downstream occupied bandwidth of 192 MHz shall be less than 0,1 ppm for five sigma of upstream OFDMA transmissions; at subcarrier frequency of 204 MHz this corresponds to less than 20 Hz.

For OFDMA upstream channels, the CM's upstream subcarrier frequency offset for each subcarrier (relative to the Master Clock reference) with downstream occupied bandwidth of 24 MHz shall be less than 1 ppm for five sigma of upstream OFDMA transmissions.

The frequency offset in ppm refers to the downstream OFDM sampling clock.

7.4.4 Forward Error Correction

7.4.4.0 Upstream Codword Parameters

DOCSIS 3.1 uses three Quasi-Cyclic Low-Density Parity-Check codes (QC-LDPC) for the upstream transmission, as depicted in table 7.3.

Table 7.3: Upstream Codeword Parameters

Code	Code rate	Codeword size in bits (N_i)	Information bits (K_i)	Parity bits (P_i)
Long code	89 % (8/9)	16 200	14 400	1 800
Medium code	85 % (28/33)	5 940	5 040	900
Short code	75 % (3/4)	1 120	840	280

Before FEC encoding, the CM shall first map the input byte stream into a bit-stream such that the MSB of the first byte is the first bit of the bit-stream.

7.4.4.1 FEC Codeword Selection

7.4.4.1.0 FEC Codeword Size

The choice of codeword sizes to be used in any given burst is based on the grant in the MAP message. The grant indicates which minislots are assigned to a given burst and which upstream profile is to be used. The CM and CMTS use this information to determine the total number of bits in the grant which are available to be used for FEC information or parity.

The CM shall follow the FEC codeword selection algorithm defined by Matlab code in clause 7.4.4.1.1 to determine the exact number, type, and size of the codewords to be used, and in what order.

The CMTS shall follow the FEC codeword selection algorithm defined by the Matlab code in clause 7.4.4.1.1 to determine the exact number, type, and size of the codewords to be used, and in what order.

Codewords are filled and transmitted in the following order, with codeword shortening applied according to rules defined in clause 7.4.4.3:

- full long codewords (if present);
- shortened long codeword (if present);
- full medium codewords (if present);
- shortened medium codeword (if present);
- full short codewords (if present);
- shortened short codewords (if present);
- zero-pad (if present).

7.4.4.1.1 FEC Codeword Selection Algorithm

The FEC codeword selection algorithm is given by:

```
% The total number of bits in the grant is given by rgrant_size
% set values for codeword sizes
% total bits = size including parity
% info bits = information bits only
% thresholds - if more bits than threshold, shorten this cw instead of
% using a smaller one
```

```
% short codeword
SHORT_TOTAL_BITS = 1120;
SHORT_INFO_BITS = 840;
SHORT_PARITY_BITS = SHORT_TOTAL_BITS - SHORT_INFO_BITS;
SHORT_TOTAL_THRESH_BITS = SHORT_PARITY_BITS + 1;
SHORT_MIN_INFO_BITS = SHORT_INFO_BITS / 2;

% medium codeword
MED_TOTAL_BITS = 5940;
MED_INFO_BITS = 5040;
MED_PARITY_BITS = MED_TOTAL_BITS - MED_INFO_BITS;
MED_TOTAL_THRESH_BITS = 3421;

% long codeword
LONG_TOTAL_BITS = 16200;
LONG_INFO_BITS = 14400;
LONG_PARITY_BITS = LONG_TOTAL_BITS - LONG_INFO_BITS;
LONG_TOTAL_THRESH_BITS = 11881;

% variable rgrant_size is input
% set rgrant_size to desired input value in workspace
% initialize output variables
rlong_cws = 0;
rshortened_long_cws = 0;
rmed_cws = 0;
rshortened_med_cws = 0;
rshort_cws = 0;
rshortened_short_cws = 0;
rother_shortened_cw_bits = 0;
rshortened_cw_bits = 0;
rpad_bits = 0;

% intermediate variable to track type of last full codeword
rlast_full_cw = '';

% now begin calculation
    bits_remaining = rgrant_size;
    % if we don't have enough space to make at least a min size shortened
    % short cw, then this grant is nothing but pad bits.
    % NOTE: in the case, the CM should ignore the grant and should not
    % transmit any bits at all in the grant.

    if rgrant_size < SHORT_MIN_INFO_BITS + SHORT_PARITY_BITS
        rpad_bits = rgrant_size;
        bits_remaining = 0;
    end

    % make as many full long cws as possible
    while bits_remaining >= LONG_TOTAL_BITS
        rlong_cws = rlong_cws + 1;
        bits_remaining = bits_remaining - LONG_TOTAL_BITS;
        rlast_full_cw = 'Long';
    end

    % if remaining bits can make a shortened long codeword, do so
    if bits_remaining >= LONG_TOTAL_THRESH_BITS
        rshortened_long_cws = 1;
        rshortened_cw_bits = bits_remaining;
        bits_remaining = 0;
    end

    % if not, make as many med cws as possible with remaining bits
    while bits_remaining >= MED_TOTAL_BITS
        rmed_cws = rmed_cws + 1;
        bits_remaining = bits_remaining - MED_TOTAL_BITS;
        rlast_full_cw = 'Medium';
    end

    % if remaining bits can make a shortened med cw, do so
    if bits_remaining >= MED_TOTAL_THRESH_BITS
        rshortened_med_cws = 1;
        rshortened_cw_bits = bits_remaining;
        bits_remaining = 0;
    end

    % if not, make as many short cws as possible with remaining bits
    while bits_remaining >= SHORT_TOTAL_BITS
        rshort_cws = rshort_cws + 1;
```

```
        bits_remaining = bits_remaining - SHORT_TOTAL_BITS;
        rlast_full_cw = 'Short';
    end

    % if remaining bits can make a shortened short cw, do so
         if bits_remaining >= SHORT_TOTAL_THRESH_BITS
        rshortened_short_cws = 1;

        % at this point we are definitely making this cw; however, we need
        % at least SHORT_MIN_INFO_BITS to put in it.  If we do not have
        % that many, we will have to borrow from the last full cw, making
        % it also a shortened cw.

        if (bits_remaining - SHORT_PARITY_BITS) >= SHORT_MIN_INFO_BITS
            % no need to borrow bits
            rshortened_cw_bits = bits_remaining;
            bits_remaining = 0;
        else
            % identify type/size of last full cw, then borrow
            % SHORT_MIN_INFO_BITS from it
            switch rlast_full_cw
                case 'Long'
                    % change last full cw to a shortened cw
                    rlong_cws = rlong_cws - 1;
                    rshortened_long_cws = rshortened_long_cws + 1;
                    % number of bits in that cw is reduced by
                    % SHORT_MIN_INFO_BITS
                    rother_shortened_cw_bits = LONG_TOTAL_BITS - ...
                        SHORT_MIN_INFO_BITS;
                    % put those bits plus bits_remaining into the last
                    % shortened cw
                    rshortened_cw_bits = SHORT_MIN_INFO_BITS + ...
                        bits_remaining;
                    bits_remaining = 0;
                case 'Medium'
                    % same steps as for long, just substitute medium
                    rmed_cws = rmed_cws - 1;
                    rshortened_med_cws = rshortened_med_cws + 1;
                    rother_shortened_cw_bits = MED_TOTAL_BITS - ...
                        SHORT_MIN_INFO_BITS;
                    rshortened_cw_bits = SHORT_MIN_INFO_BITS + ...
                        bits_remaining;
                    bits_remaining = 0;
                case 'Short'
                    % again, same steps as for long - now substitute short
                    rshort_cws = rshort_cws - 1;
                    rshortened_short_cws = rshortened_short_cws + 1;
                    rother_shortened_cw_bits = SHORT_TOTAL_BITS - ...
                        SHORT_MIN_INFO_BITS;
                    rshortened_cw_bits = SHORT_MIN_INFO_BITS + ...
                        bits_remaining;
                    bits_remaining = 0;
            end
        end
    end

    % any space left over at this point has to be filled with pad bits (it
    % cannot fit any cws)
    if bits_remaining > 0
        rpad_bits = bits_remaining;
        bits_remaining = 0;
    end
```

Based on the algorithm above, the minimum grant size allowed is:

SHORT_MIN_INFO_BITS + SHORT_PARITY_BITS = SHORT_INFO_BITS / 2 + SHORT_PARITY_BITS = 420 + 280 bits = 700 bits.

This grant is sufficient for 52 bytes of information.

The CM should not transmit in any grant smaller than the minimum allowed grant size specified above.

In some cases, the total number of information bits derived from the algorithm above will not be an integer number of bytes. In such cases there are 1 - 7 leftover bits that are not enough to form the last information byte. The CM shall set the values of information bits left over after the FEC Codeword Selection Algorithm forms bytes, to 1. These bits will be discarded by the CMTS after decoding.

The FEC codeword selection algorithm follows the procedure below:

- If there are enough bits in the grant to create a full long codeword, do so. Continue creating full long codewords until there are not enough bits remaining.
- If the number of bits remaining is greater than or equal to the minimum allowed size for a shortened long codeword, create such a codeword and end the burst.
- Otherwise, if there are enough bits remaining to create a full medium codeword, do so. Continue creating full medium codewords until there are not enough bits remaining.
- If the number of bits remaining is greater than or equal to the minimum allowed size for a shortened medium codeword, create such a codeword and end the burst.
- Otherwise, if there are enough bits remaining to create a full short codeword, do so. Continue creating full short codewords until there are not enough bits remaining.
- If there are enough bits remaining to create a shortened short codeword containing at least the minimum allowed number of information bits, do so and end the burst.
- Otherwise, if there are enough bits remaining to create a shortened short codeword with fewer than the minimum allowed number of information bits, remove a number of bits equal to the minimum allowed number of short codeword information bits from the last full codeword, changing it to a shortened codeword. Add this number of bits to the bits remaining and create a shortened short codeword using these bits, and end the burst.
- Otherwise, there are not enough bits remaining to create a shortened short codeword (i.e. fewer bits than the number required for one information bit plus the applicable number of parity bits). These bits will be padded with zeros by the CM and will be ignored by the CMTS.
- If a grant does not contain enough bits to create any codewords, the CM should not transmit in the grant.

The reverse calculation to determine the grant size required to hold the desired number of bits, number of codewords and codeword sizes is given in Annex H.

7.4.4.2 FEC Encoding

All three LDPC encoders are systematic. Every encoder encodes N-M information bits $i_0,\cdots,i_{N-M-1}$ into a codeword $c=(i_0,\cdots,i_{N-M-1},p_0,\cdots,p_{M-1})$ by adding m parity bits obtained so that $Hc^T=0$, where H is an m*n parity check matrix. The parity-check matrix can be divided into blocks of L∗L submatrices, where L represents the submatrix size or lifting factor. The parity-check matrix in compact circulant form is represented by an m×n block matrix:

$$H=\begin{bmatrix} H_{1,1} & H_{1,2} & H_{1,3} & \cdots & H_{1,n} \\ H_{2,1} & H_{2,2} & H_{2,3} & \cdots & H_{2,n} \\ H_{3,1} & H_{3,2} & H_{3,3} & \cdots & H_{3,n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ H_{m,1} & H_{m,2} & H_{m,3} & \cdots & H_{m,n} \end{bmatrix}$$

Each submatrix $H_{i,j}$ is an L*L all-zero submatrix or a cyclic right-shifted identity submatrix. The last $n-m$ sub-matrix columns represent the parity portion of the matrix. In the present document, the L*L sub-matrix $H_{i,j}$ is represented by a value in {-, 0,…, L-1}, where a '-' value represents an all-zero submatrix, and the remaining values represent an L*L identity submatrix cyclically right-shifted by the specified value. The code rate is (n-m)/n and a codeword length is N=n*L bits.

The CM shall employ the following matrix table for the long code rate:

Rate= 89 % (16 200, 14 400) code, m=5 rows * n=45 columns, L=360

93	271	-	83	26	208	245	200	-	175	331	17	86	-	337	-	238	81	-	307	-	165	-	47	76	73	150	349	139	331	118	345	27	294	-	145	279	97	106	160	143	-	-	-	-
274	115	329	338	124	-	293	-	69	64	342	-	88	139	-	137	212	-	157	195	357	81	194	1	159	56	72	126	277	156	32	111	175	-	306	224	-	206	-	29	106	334	-	-	-
134	355	175	24	253	242	-	187	94	26	87	302	-	191	323	22	-	245	294	240	84	76	342	345	174	269	329	-	214	-	-	-	-	218	104	40	197	73	229	63	-	270	72	-	-
-	-	184	70	247	14	22	7	285	54	-	352	26	108	10	298	123	139	117	-	336	49	202	359	342	-	224	106	-	273	177	245	98	355	178	176	147	-	280	-	-	-	221	208	-
253	273	90	-	-	151	311	320	339	-	295	148	48	91	62	100	232	146	200	135	12	-	179	-	-	232	-	21	331	313	349	34	97	187	38	-	235	52	170	58	-	-	-	257	0

The CM shall employ the following matrix table for the medium code rate:

Rate= 85 % (5 940, 5 040) code, m=5 rows * n=33 columns, L=180

142	158	113	124	92	44	93	70	172	3	25	44	141	160	50	45	118	84	-	64	66	97	1	115	8	108	-	-	22	-	-	-	-
54	172	145	28	55	19	159	22	96	12	85	-	128	5	158	120	51	171	65	141	-	42	83	7	-	39	121	84	101	171	-	-	-
63	11	112	114	61	123	72	55	114	20	53	114	42	33	4	66	163	50	46	17	175	-	-	-	92	-	41	138	-	34	74	-	-
28	160	102	44	8	84	126	9	169	174	147	24	145	-	26	-	-	-	67	82	4	177	151	131	139	117	36	18	-	-	23	8	-
52	159	75	74	46	71	42	11	108	153	-	72	-	163	-	9	2	168	158	-	1	49	89	63	179	10	75	161	-	-	-	177	19

The CM shall employ the following matrix table for the short code rate:

Rate= 75 % (1 120, 840) code, m=5 rows * n=20 columns, L=56

5	14	12	1	2	37	45	26	24	0	3	-	34	7	46	10	-	-	-	-
0	35	1	26	0	10	16	16	34	4	2	23	0	51	-	49	20	-	-	-
12	28	22	46	3	16	51	2	25	29	19	18	52	-	37	-	34	39	-	-
0	51	16	31	13	39	27	33	8	27	53	13	-	52	33	-	-	38	7	-
36	6	3	51	4	19	4	45	48	9	-	11	22	23	43	-	-	-	14	1

7.4.4.3 Shortening of LDPC Codewords

Shortening of LDPC codewords is useful in order to optimize FEC protection for the payload. If a shortened codeword is required, the CM shall construct it as follows:

1) Binary zeros are appended to a reduced number of information bits at the input of the encoder.

2) The encoder calculates the parity bits.

3) The appended binary zeros are removed from the transmitted shortened codeword.

7.4.5 Data Randomization

The CM shall implement a randomizer in the upstream modulator shown in figure 7.4 where the 23-bit seed value is programmable.

At the beginning of each grant, the register is cleared and the seed value is loaded. The CM shall use the seed value to calculate the scrambler bit which is combined in an XOR with the first bit of data of each grant.

The CM shall configure the randomizer seed value in response to the Upstream Channel Descriptor provided by the CMTS.

The CM shall use x^23+x^18+1 for the data randomizer polynomial.

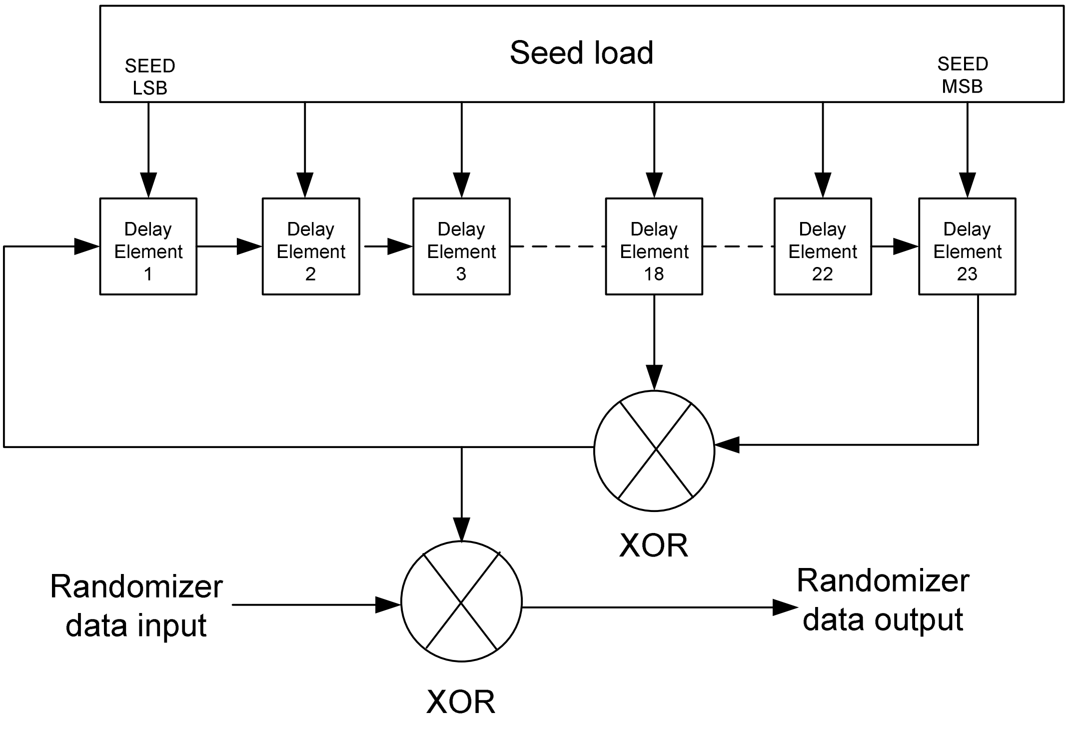

Figure 7.4: Upstream Data Randomizer

7.4.6 Time and Frequency Interleaving and De-interleaving

Upstream transmissions can be affected by burst noise that reduces the SNR of all the subcarriers of a few successive OFDMA symbols. Upstream transmissions may also be impacted by ingress, i.e. relatively narrowband interferers, that can last for several symbol periods and therefore reduce the SNR of a subset of subcarriers over an entire OFDMA frame. The purpose of the interleaver is to distribute the affected subcarriers over a number of FEC blocks, enabling the FEC at the receiver to correct the corrupt data.

Time and frequency interleaving in the upstream are applied together in the CM as a single operation and hence referred to as upstream interleaving. Similarly, time and frequency de-interleaving are performed together as a single operation in the CMTS, and hence referred to here as upstream de-interleaving.

The CM shall apply interleaving to upstream OFDMA subcarriers. The interleaving is applied after the randomizer in conjunction with the bits being allocated to QAM subcarriers, and before the OFDMA IFFT operation.

The CM shall exclude any zero-valued minislots from the interleaving process.

The CM shall apply interleaving to a sequence of minislots of an OFDMA frame of a specific grant, not exceeding 24, as described in this clause.

The CMTS shall apply de-interleaving which is the inverse of the CM interleaving function carried out by the CM.

The maximum number of minislots over which interleaving is applied is equal to 24. If the number of minislots of a specific grant in one OFDMA frame is less than or equal to 24, then interleaving is applied over all of these minislots.

If the number of minislots in a specific grant in one OFDMA frame is more than 24, say N_{MS_Total}, then the CM shall partition these N_{MS_Total} minislots into $ceil(N_{MS\_Total}/24)$ blocks of minislots, as uniformly as possible, without the number of minislots in any block exceeding 24, using the algorithm given in the flow diagram shown in figure 7.5.

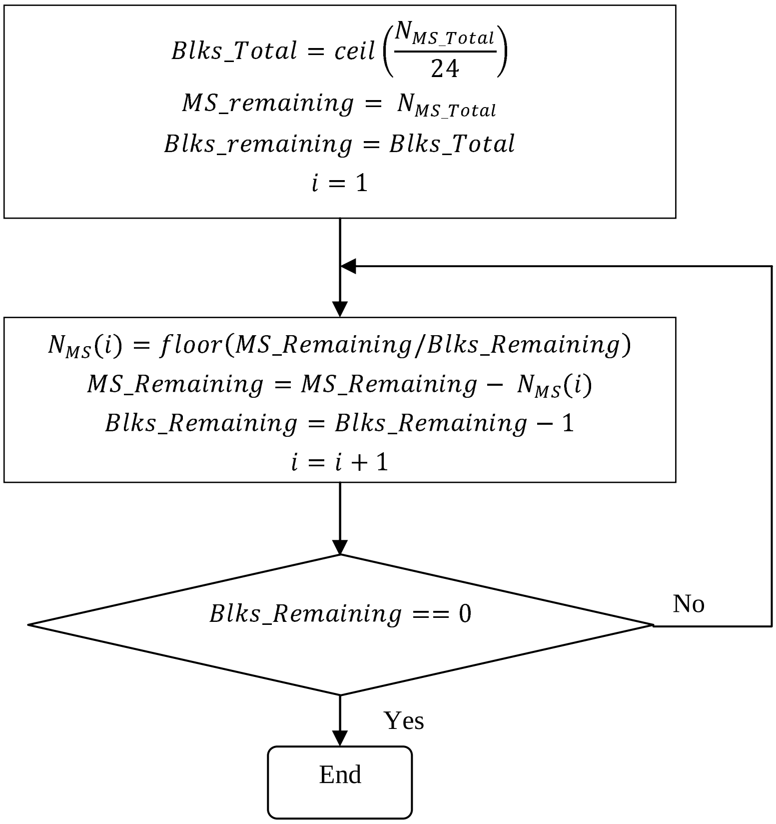

Figure 7.5: Calculating Number of Minislots in Each Block for Upstream Interleaving

The described algorithm in figure 7.5 yields the sequence:

$$\{N_{MS}(i), for\ \ i = 1, 2, \dots, Blks\_Total\}$$

There are *Blks_Total* of blocks of minislots, and in each block there are $N_{MS}(i)$ minislots. For each block of $N_{MS}(i)$ minislots the CM shall apply interleaving as described in this clause.

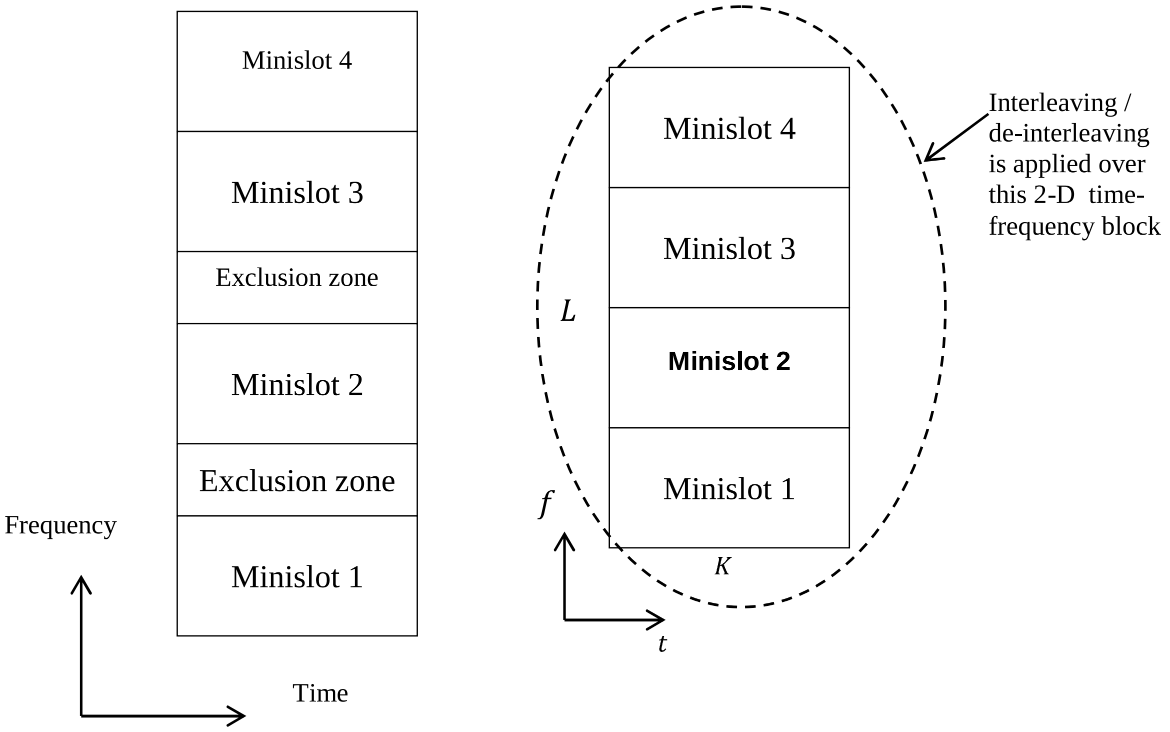

Figure 7.6: Illustrating Minislots of a Grant over which Interleaving is Performed

Figure 7.6 shows an example of a block of four minislots over which interleaving is applied. The horizontal axis is time. Every vertical column constitutes a segment of an OFDMA symbol. The vertical axis is frequency. Each horizontal line is a set of subcarriers at a specific frequency over several symbols.

In the illustration in figure 7.6, there is an exclusion zone between minislots 1 and 2. There is also an exclusion zone between minislots 2 and 3. All four minislots are merged to form a two-dimensional grant for the purpose of interleaving and de-interleaving. In the CM, the interleaving is applied first and then the exclusion zones are introduced in mapping of the interleaved subcarriers onto OFDMA symbols.

Interleaving and de-interleaving are two-dimensional operations in the time-frequency plane.

The system block diagram for interleaving is illustrated in figure 7.7.

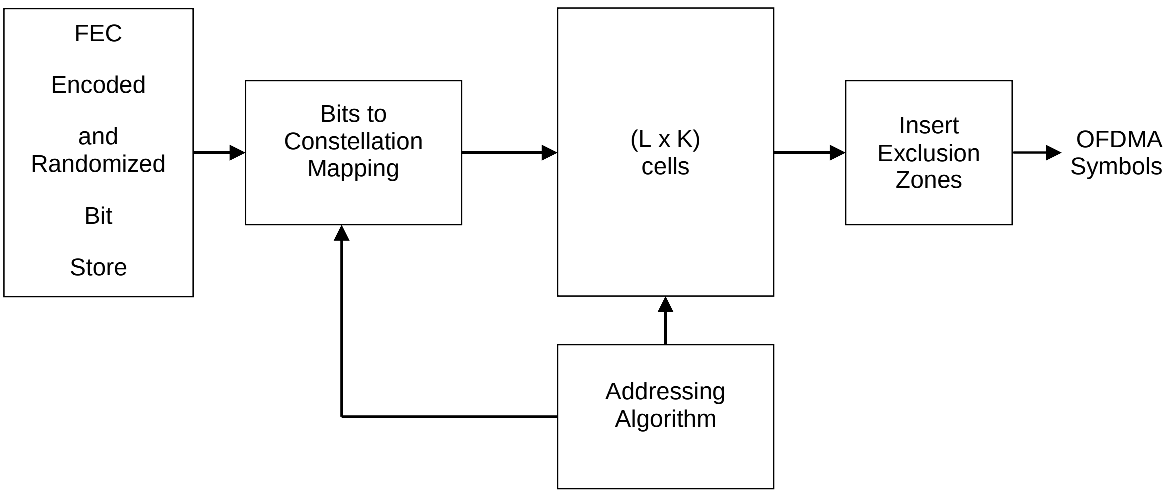

Figure 7.7: Sample Interleaver Block Diagram

The two-dimensional array is addressed by coordinate pair (t, f). The horizontal dimension is *K*, which is the number of OFDMA symbols in the frame. The vertical dimension is L, which is the total number of subcarriers in all the minislots that make up the grant in the current frame. Each element in this two-dimensional array is a member of the set:

$$\{\text{Data subcarrier, Complementary data subcarrier, Pilot}\}$$

All data subcarriers in a minislot will have the same QAM constellation. All complementary data subcarriers in a minislot will also have the same QAM constellation, but this will be lower in order than that of the data subcarriers in that minislot.

Furthermore, the QAM constellations of data and complementary pilots need not be the same for all minislots in the grant.

Interleaving involves the following two stages:

- Writing the subcarriers in the cells of the two dimensional array of size (*L* x *K*):
    - The CM shall follow the algorithm given in this clause for placing data subcarriers and complementary data subcarriers in the cells of this two-dimensional array.
    - The CM shall not place any data subcarriers or complementary pilots at locations corresponding to pilots which are also part of this two-dimensional array.
- Reading data subcarriers as well as pilots along vertical columns of the two-dimensional array, in the ascending order of the time dimension coordinate t, inserting exclusion zones, if any, and passing these segments of OFDMA symbols to the IFFT processor.

Figure 7.6 is for illustration only. An implementation may not necessarily have a separate FEC Encoded bit store. The FEC encoded and randomized output may be mapped directly into QAM subcarriers and placed in the cells of the *(L x K)* two-dimensional array. In that way the two-dimensional array may form the output buffer for the FEC encoder.

The Address Generation and the Bit Mapping algorithms need to know:

- Values of *K* and *L*
- Locations of pilots
- Locations of complementary pilots
- QAM constellations for data subcarriers of all minislots of the grant in the frame
- QAM constellations for complementary pilots of all minislots of the grant in the frame
- Minislot boundaries along the frequency dimension of the *(L x K)* array

The interleaving algorithm follows the flow diagram in figure 7.8.

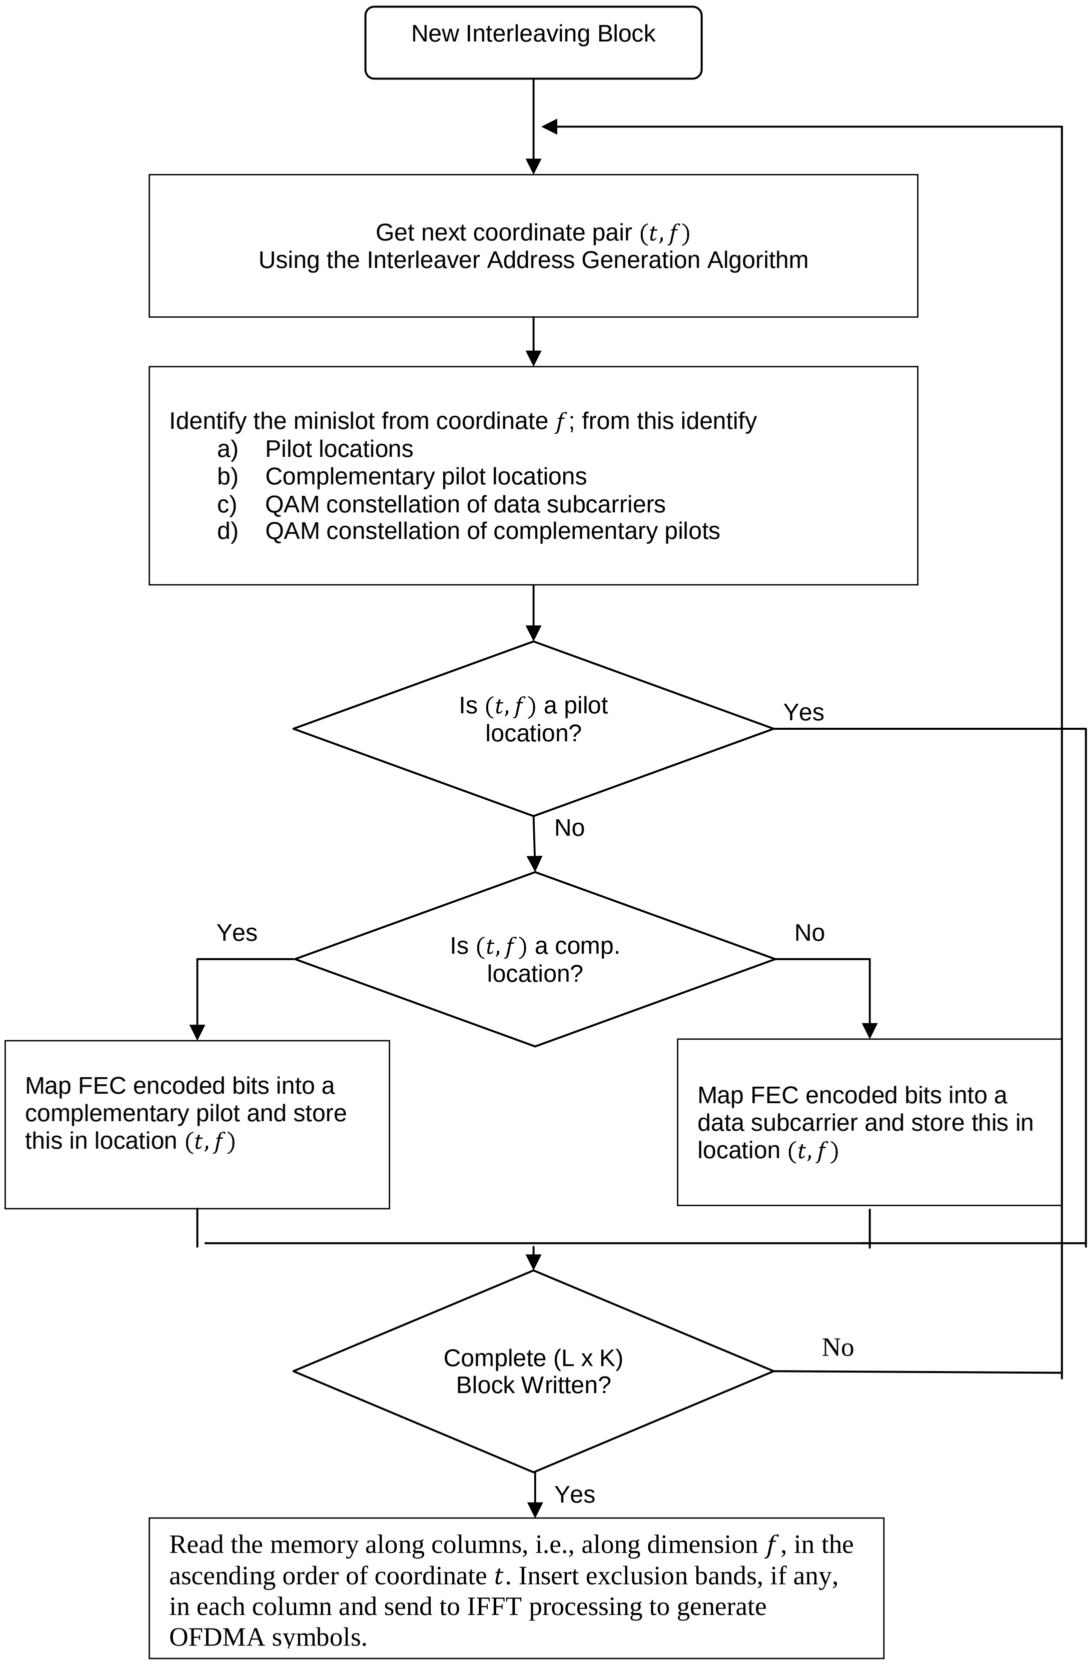

Figure 7.8: Interleaving a Grant within an OFDMA Frame

The address generation algorithm for getting the next coordinate pair *(t, f)* is described below. This makes use of three bit-reverse counters:

1) Count_t
2) Count_f
3) Count_diagonal

The third counter is used to count the diagonals. This is because subcarriers are written in the two-dimensional *t-f* array along diagonals. To write along diagonals in natural order, both the counters *Count_t* and *Count_f* have to be incremented at the same time. Once one diagonal is written, the third counter *Count_diagonal* is incremented by one.

However, in order to maximize the separation of successive subcarriers in the time-frequency plane, bit-reversed counting is used in all of the above three counters. This ensures that successive subcarriers are not written on successive locations in the diagonals.

The algorithm for generating the sequence of addresses *(t, f)* is described below with sample C code given in Annex H.

Initialize three counters, *Count_t*, *Count_f* and *Count_diagonal*, to zero.

For each value of OFDM symbol index *idx_t* going from 0 to (K-1), implement the following 4 steps:

1) For each value subcarrier index *idx_f* going from 0 to (L-1) implement the following 4 sub-steps:
   - Generate the component *t* of *(t, f)* by passing *Count_t* and parameter K to the Bit-Reverse counter defined in the flow diagram of figure 7.9. This returns *t* and a new counter value for *Count_t*.
   - Generate the component *f* of *(t, f)* by passing *Count_f* and parameter L to the Bit-Reverse counter defined in the flow diagram of figure 7.9. This returns *f* and a new counter value for *Count_f*.
   - Increment *Count_t* by one modulo K_1.
   - Increment *Count_f* by one modulo L_1.
2) Increment *Count_diagonal* by one modulo K_1.
3) Pass *Count_diagonal* and parameter K to the Bit-Reverse counter defined in the flow diagram of figure 7.9. This returns a new counter value for *Count_diagonal*.
4) Set *Count_t* to the value of *Count_diagonal*. Set *Count_f* to zero. Return to step 1.

The pseudo code given in Annex H will generate the entire sequence of addresses. This is for illustration purposes only. In the actual implementation, the code may be modified to generate one address at a time, so that data may be saved in the memory in parallel with address generation.

The pseudo code in Annex H contains a call to the function called Bit_Reverse_Count. The algorithm implemented by this function is explained below with reference to the flow diagram of figure 7.9. With no loss of generality, this explanation uses the function call for *Count_t*.

The parameter K_1 is defined as the smallest power of 2 that is equal to or greater than K. The minimum number of binary bits needed to represent K is k_1. In this case then, $K_1 = 2^{k_1}$. Similarly, parameter L_1 is defined as the smallest power of 2 greater than L.

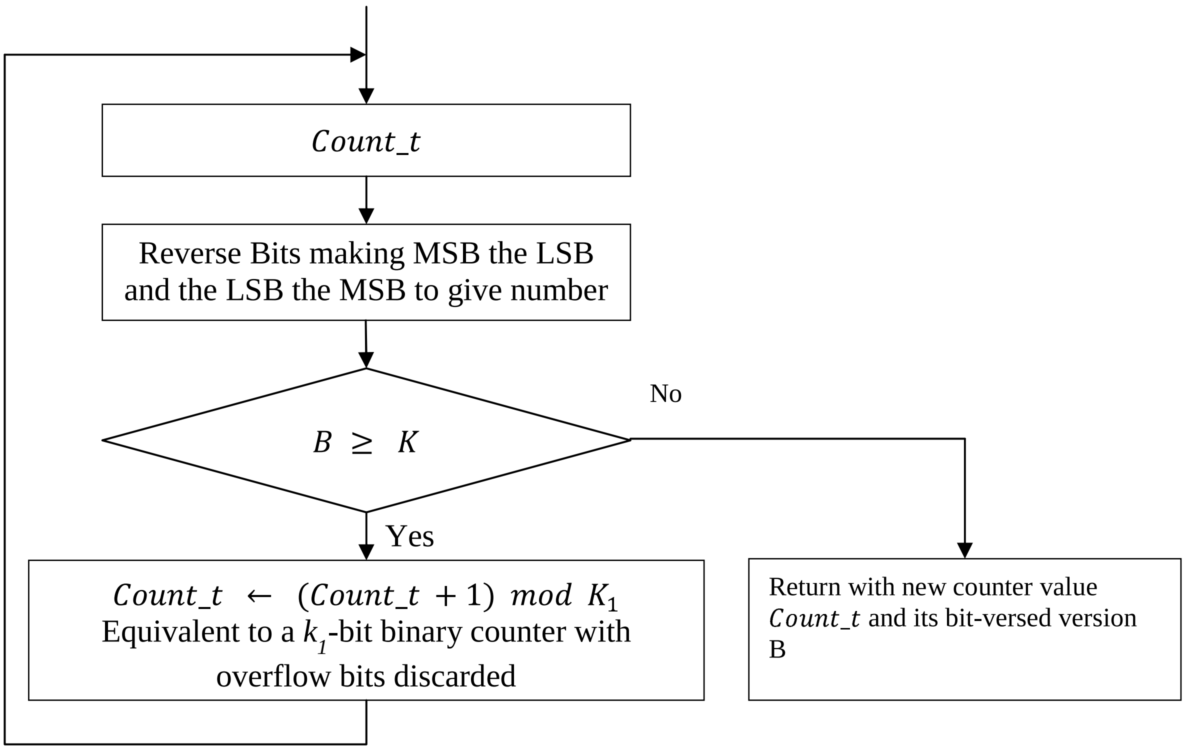

Figure 7.9: Bit-Reversed Counter Implementation

Bit-reverse counting employs a modulo 2^{k_1} counter. This is equivalent to a k_1-bit counter with overflow bits discarded. In bit-reversed counting the above counter is incremented beginning from its current value until the bit-reversed version of the counter value is in the range [0, (K-1)].

The term bit-reversion is defined below. Let A be the value of *Count_t* and let B be its bit-reversed value. Let the binary representation of A be given by:

$$A = \sum_{i=0}^{k_1-1} a_i 2^i$$

Then B is given by:

$$B = \sum_{i=0}^{k_1-1} b_i 2^i = \sum_{i=0}^{k_1-1} a_i 2^{k_1-1-i}$$

7.4.7 Mapping of Bits to Cell Words

CMs are granted transmission opportunities by minislots, and minislots are associated with subcarriers. All subcarriers of a specific type (data subcarriers, pilots, complementary pilots) within a minislot have the same modulation order, although different minislots may have different modulation orders; the modulation order to be used is determined by the Profile associated with the minislot.

The CM shall modulate the incoming serial binary bitstream from the data scrambler to constellation symbols using the constellation mapping described in clause 7.4.8.2.

The CM shall map the incoming bit stream {a0, a1, a2, …} to {y0, y1, …} for each QAM symbol such that the first incoming bit is the most-significant bit of the constellation symbol when bits are mapped into constellation symbols.

The CM shall have the same nominal average power for all constellations.

7.4.8 Mapping and De-mapping Bits to/from QAM Subcarriers

7.4.8.0 Minislots Structure

CMs are granted transmission opportunities by minislots, and minislots are associated with subcarriers. All subcarriers of a specific type (i.e. data subcarriers, pilots, complementary pilots or zero-valued subcarriers) within a minislot have the same modulation order, although different minislots may have different modulation orders; the modulation order to be used is determined by the Profile associated with the minislot.

Some minislots may be specified as zero-valued in some profiles. The CM shall not transmit anything in the minislots of these profiles. The CM shall set all subcarriers, including data subcarriers, pilots and complementary pilots to zero in these minislots of these profiles. A zero-valued minislot in one profile may not be zero-valued in another profile.

7.4.8.1 Modulation Formats

The CM modulator shall support zero valued subcarriers of upstream OFDMA channels.

The CM modulator shall support BPSK, QPSK, 8-QAM, 16-QAM, 32-QAM, 64-QAM, 128-QAM, 256-QAM, 512-QAM, 1024-QAM, 2048-QAM, and 4096-QAM for subcarriers of upstream OFDMA channels. BPSK is used for pilots and complimentary pilots only, and not used for data transmission.

The CMTS demodulator shall support zero valued subcarriers of upstream OFDMA channels.

The CMTS demodulator shall support BPSK, QPSK, 8-QAM, 16-QAM, 32-QAM, 64-QAM, 128-QAM, 256-QAM, 512-QAM, and 1024-QAM for subcarriers of upstream OFDMA channels. BPSK is used for pilots and complimentary pilots only, and not used for data transmission.

The CMTS demodulator should support 2048-QAM and 4096-QAM for subcarriers of upstream OFDMA channels.

7.4.8.2 Constellation Mapping

The CM shall encode the bit stream such that the first bit is the most-significant bit of the first QAM subcarrier constellation m-tuple.

$\longrightarrow$ Time

$x_0\ x_1\ \dots\ x_{m-1}\ x_m\ x_{m+1}\ \dots\ x_{2m-1}\ x_{2m}\ \dots$

	MSB	LSB
QAM m-tuple 1:	$x_0\ x_1\ \dots$	x_{m-1}
QAM m-tuple 2:	$x_m\ x_{m+1}\ \dots$	x_{2m-1}
QAM m-tuple 3:	$x_{2m}\ x_{2m+1}\ \dots$	x_{3m-1}

Figure 7.10: Bitstream to QAM M-Tuple Mapping

The CM shall modulate the interleaved m-tuples onto subcarriers using QAM constellation mappings described in Annex A.

The CM shall ensure that subcarriers of all QAM constellations have the same nominal average power using the scaling factors given in table A.1 of Annex A.

The CMTS receiver shall demodulate the incoming QAM constellation subcarriers of a minislot according to the Profile associated with the minislot, with the first demapped value associated with the most-significant bit of the constellation point.

7.4.9 REQ Messages

REQ messages are short messages used by the CM to request transmission opportunities from the CMTS. These messages have a different structure than the data messages: they are always 56 bits long, they always use QPSK modulation, do not apply any FEC, and are block interleaved.

REQ message processing is described in figure 7.11.

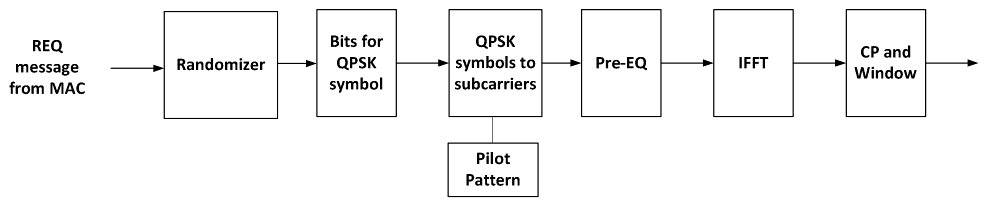

Figure 7.11: REQ Messages Processing

The CM shall randomize REQ messages using the randomizer described in clause 7.4.5.

The CM shall modulate REQ messages using QPSK.

The CM shall use the subslot minislots with the pilot patterns as specified in clauses 7.4.17.4 and 7.4.17.5 for 25 KHz and 50 kHz subcarrier spacing.

The CM shall write the REQ messages QPSK symbols into subslots as described in clause 8.2.4.2.

The CM shall use the same CP size and RP size used for the data transmission.

7.4.10 IDFT

The upstream OFDMA signal transmitted by the CM is described using the following IDFT equation:

$$x(i) = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} X(k) \exp\left(j \frac{2\pi i\left(k-\frac{N}{2}\right)}{N}\right), \quad for \;\; i = 0, 1, \ldots, (N-1)$$

Where N equals 2 048 with 50 KHz subcarrier spacing and 4 096 with 25 KHz subcarrier spacing. The resulting time domain discrete signal, $x(i)$, is a baseband complex-valued signal, sampled at 102,4 Msamples per second.

In this definition of the IDFT $X(0)$ is the lowest frequency component; and $X(N\text{-}1)$ is the highest frequency component.

The IDFT operation is illustrated in figure 7.12.

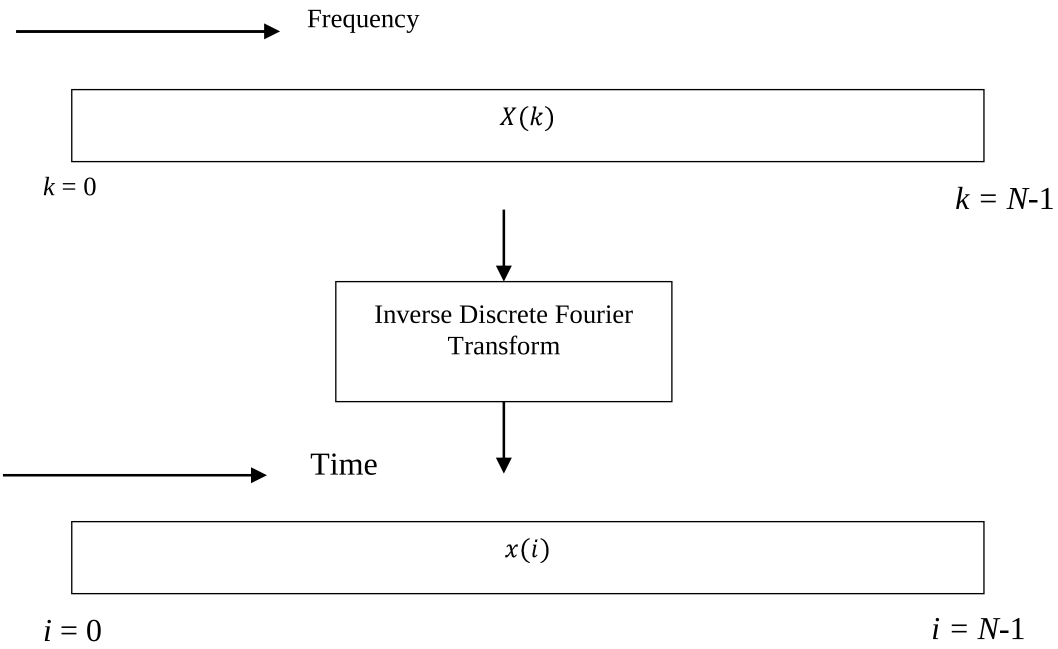

Figure 7.12: Inverse Discrete Fourier Transform

7.4.11 Cyclic Prefix and Windowing

7.4.11.0 Background and General Requirements

Cyclic prefix and windowing are applied in the upstream transmission. Cyclic prefix is added in order to enable the receiver to overcome the effects of inter-symbol interference (ISI) and caused by micro-reflections in the channel. Windowing is applied in order to maximize channel capacity by sharpening the edges of the spectrum of the OFDMA signal. Spectral edges occur at the two ends of the spectrum of the OFDM symbol, as well as at the ends of internal exclusion bands.

In the presence of a micro-reflection in the transmission medium, the received signal is the sum of the main signal and the delayed and attenuated micro-reflection. As long as this delay (τ) is less than the time duration of the cyclic prefix (T_{CP}), the CMTS receiver can trigger the FFT to avoid any inter-symbol or inter-carrier interference due to this micro reflection, as shown in figure 7.13.

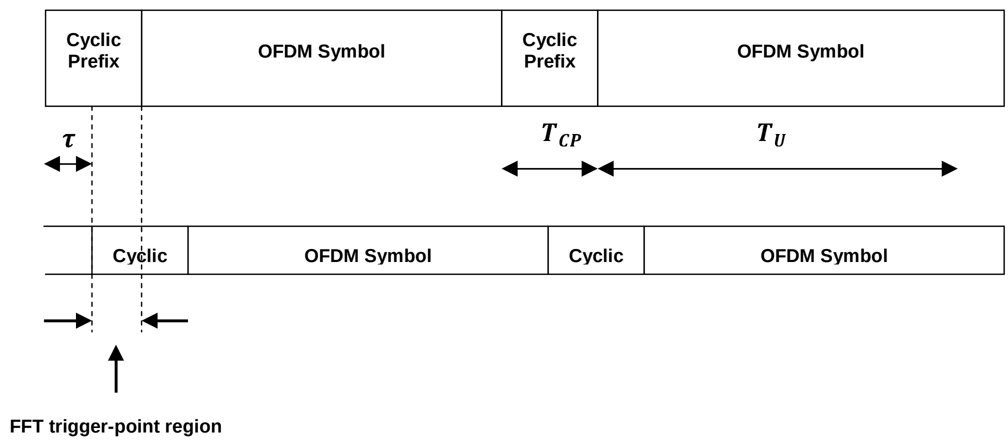

Figure 7.13: Signal with Micro-Reflection at the Receiver

If the delay of the micro-reflection exceeds the length of the cyclic prefix, the ISI resulting from this micro-reflection is:

$$ISI = \frac{(\tau - T_{CP})A^2}{T_U}$$

where:

τ is the delay introduced by the micro-reflection

T_{CP} is the cyclic prefix length in µs

A is the relative amplitude of the micro-reflection

T_U is FFT duration (20 or 40 µs)

The inter-carrier-interference introduced by this micro-reflection is of the same order as the ISI.

The CM transmitter shall apply the configured CP and Windowing as described in clause 7.4.11.1.

The CM transmitter shall support the cyclic prefix values defined in table 7.4.

The CMTS receiver shall support the cyclic prefix values defined in table 7.4.

Table 7.4: Cyclic Prefix (CP) Values

Cyclic Prefix (µs)	Cyclic Prefix Samples (N_{cp})
0,9375	96
1,25	128
1,5625	160
1,875	192
2,1875	224
2,5	256
2,8125	288
3,125	320
3,75	384
5,0	512
6,25	640

In table 7.4 the cyclic prefix (in µs) is converted into samples using the sample rate of 102,4 Msamples/s.

Windowing is applied in the time domain by tapering (or rolling off) the edges using a raised cosine function.

The CMTS shall support the eight roll-off period values listed in table 7.5.

The CM shall support the eight roll-off period values listed in table 7.5.

The CMTS shall not allow a configuration in which the Roll-Off Period value is ≥ the Cyclic Prefix value.

Table 7.5: Roll-Off Period (RP) Values

Roll-Off Period (µs)	Roll-Off Period Samples (N_{RP})
0	0
0,3125	32
0,625	64
0,9375	96
1,25	128
1,5625	160
1,875	192
2,1875	224

The Roll-Off Period is given in µs and in number of samples using the sample rate of 102,4 Msamples/s.

7.4.11.1 Cyclic Prefix and Windowing Algorithm

The algorithm for cyclic prefix extension and windowing is described here with reference to figure 7.14.

The CM shall support cyclic prefix extension and windowing as described in this clause.

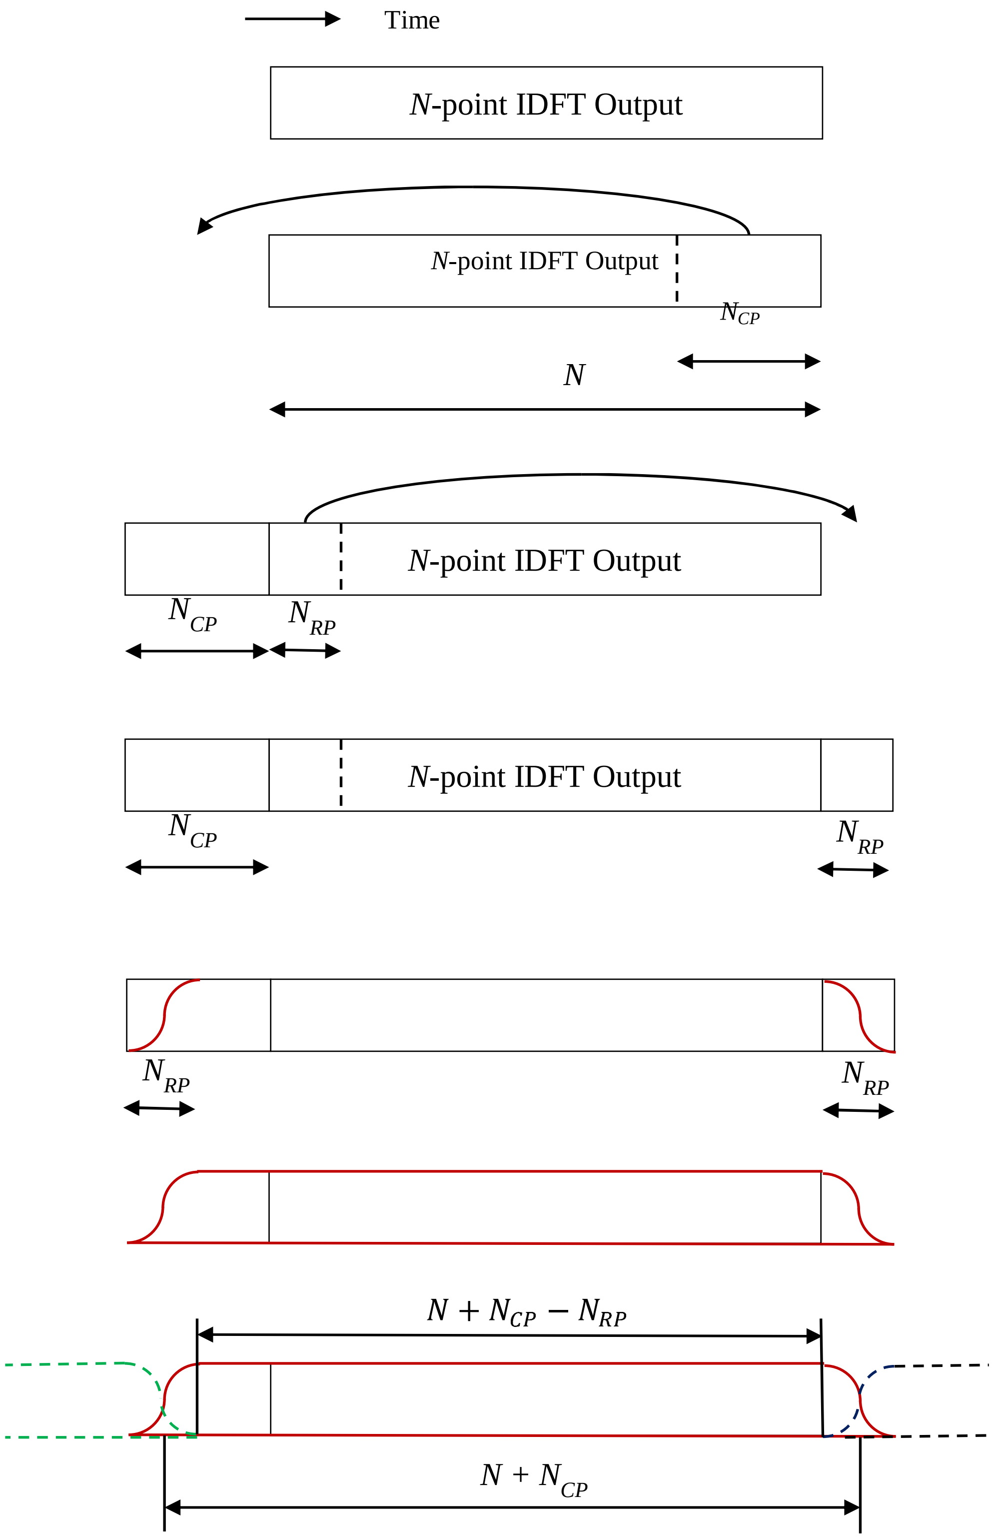

Figure 7.14: Cyclic Prefix and Windowing Algorithm

Processing begins with the *N*-point output of the IDFT. Let this be:

$$\{\, x(0), x(1), \dots\ , x(N-1)\}$$

The N_{CP} samples at the end of this N-point IDFT are copied and prepended to the beginning of the IDFT output to give a sequence of length $(N+N_{CP})$:

$$\{x(N-N_{CP}), x(N-N_{CP}+1), \ldots, x(N-1), x(0), x(1), \ldots, x(N-1)\}$$

The N_{RP} samples at the start of this N-point IDFT are copied and appended to the end of the IDFT output to give a sequence of length $(N+N_{CP}+N_{RP})$:

$$\{x(N-N_{CP}), x(N-N_{CP}+1), \ldots, x(N-1), x(0), x(1), \ldots, x(N-1), x(0), x(1), \ldots, x(N_{RP}-1)\}$$

Let this extended sequence of length $(N+N_{CP}+N_{RP})$ be defined as:

$$\{y(i),\ i = 0, 1, \ldots, (N+N_{CP}+N_{RP}-1)\}$$

N_{RP} samples at both ends of this extended sequence are subject to tapering. This tapering is achieved using a raised-cosine window function; a window is defined to be applied to this entire extended sequence. This window has a flat top and raised-cosine tapering at the edges, as shown in figure 7.15.

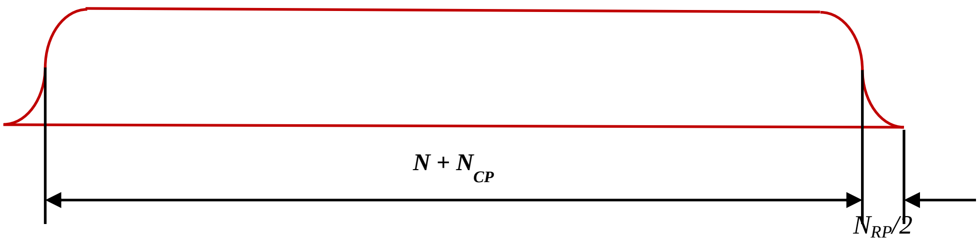

Figure 7.15: Tapering Window

The window function $w(i)$ is symmetric at the centre; therefore, only the right half of the window is defined in the following equation:

$$w\left(\frac{N+N_{CP}+N_{RP}}{2}+i\right) = 1.0,\ \ for\ \ i = 0, 1, \ldots, \left(\frac{N+N_{CP}-N_{RP}}{2}-1\right)$$

$$w\left(i+\frac{N+N_{CP}+N_{RP}}{2}\right) = \frac{1}{2}\left(1-\sin\left(\frac{\pi}{\alpha(N+N_{CP})}\left(i-\frac{N+N_{CP}}{2}+1/2\right)\right)\right),$$

$$\text{for } i = \left(\frac{N+N_{CP}-N_{RP}}{2}\right), \ldots, \left(\frac{N+N_{CP}+N_{RP}}{2}-1\right)$$

Here,

$$\alpha = \frac{N_{RP}}{N+N_{CP}}$$

defines the window function for $(N+N_{CP}+N_{RP})/2$ samples. The complete window function of length $(N+N_{CP}+N_{RP})$ is defined using the symmetry property as:

$$w\left(\frac{N+N_{CP}+N_{RP}}{2}-i-1\right) = w\left(\frac{N+N_{CP}+N_{RP}}{2}+i\right),$$

$$\text{for } i = 0, 1, \ldots, \frac{N+N_{CP}+N_{RP}}{2}-1$$

This yields a window function (or sequence): $\{w(i),\ i = 0, 1, \ldots, (N+N_{CP}+N_{RP}-1)\}$. The length of this sequence is an even-valued integer.

The above window function is applied to the sequence $\{y(i)\}$:

$$z(i) = y(i)\,w(i), for\ i = 0, 1, \ldots, (N+N_{CP}+N_{RP}-1)$$

Each successive set of N samples at the output of the IDFT yields a sequence $z(i)$ of length $(N+N_{CP}+N_{RP})$. Each of these sequences is overlapped at each edge by N_{RP} samples with the preceding and following sequences, as shown in the last stage of figure 7.14. Overlapping regions are added together.

To define this "overlap and add" function mathematically, consider two successive sequences $z_1(i)$ and $z_2(i)$. The overlap and addition operations of these sequences are defined using the following equation:

$$z_1(N + N_{CP} + i) + z_2(i), \text{for } i = 0, 1, \ldots, N_{RP} - 1$$

That is, the last N_{RP} samples of sequence $z_1(i)$ are overlapped and added to the first N_{RP} samples of sequence $z_2(i)$.

7.4.11.2 Parameter Optimization

7.4.11.2.1 Impacts of Cyclic Prefix and Windowing

The combination of cyclic prefix insertion and windowing can impact OFDM symbol duration: once the CP and RP additions have been made, the length of the extended OFDM symbol is ($N + N_{CP} + N_{RP}$) samples. Of this, ($N_{RP}/2$) samples are within the preceding symbol, and ($N_{RP}/2$) samples are within the following symbol. This yields a symbol period of ($N + N_{CP}$) samples.

In addition, successive symbols interfere with each other by ($N_{RP}/2$) samples. Therefore, the non-overlapping flat segment of the transmitted symbol = ($N + N_{CP} - N_{RP}$).

There are eleven possible values for N_{CP} and eight possible values for N_{RP}. This gives 88 possible values for α. However, combinations $N_{RP} \geq N_{CP}$ are not permitted. This limits the number of possible combinations for α.

The user would normally select the cyclic prefix length N_{CP} to meet a given delay spread requirement in the channel. Then the user would select the N_{RP} to meet the roll-off (i.e. the α parameter) requirement. Increasing α parameter leads to sharper spectral edges in the frequency domain.

However, increasing N_{RP} for a given N_{CP} reduces the non-overlapping flat region of the symbol, thereby reducing the ability of the receiver to overcome inter-symbol-interference. Similarly, increasing N_{CP} for a given N_{RP} does reduce the roll-off parameter α.

7.4.11.2.2 Joint Optimization of Cyclic Prefix and Windowing Parameters

It is clear from the clause above that the parameters N_{CP} and N_{RP} have to be jointly optimized for given channel, taking into account the following properties of the channel:

a) Bandwidth of the transmitted signal

b) Number of exclusion zones in the transmitted bandwidth

c) Channel micro-reflection profile

d) QAM constellation

The QAM constellation defines the tolerable inter-symbol and inter-carrier interference. This in turn defines the cyclic prefix for a given micro-reflection profile. The bandwidth of the transmitted signal and the number of exclusion zones define the sharpness of the spectral edges, and hence the amount of tapering. However, the amount of tapering and the flat region of the cyclic prefix are not independent variables. Therefore, an optimization program is needed to identify optimum values for N_{CP} and N_{RP} for the above parameters. This optimization is important because it does have significant impact on channel capacity, i.e. the bit rate.

The joint optimization of N_{CP} and N_{RP} is left to the network operator.

7.4.12 Burst Timing Convention

7.4.12.0 OFDMA Transmission

The start time of an OFDMA transmission by a CM is referenced to an OFDMA frame boundary that corresponds to the starting minislot of the transmission opportunity.

For all transmissions, except Fine Ranging and Requests in subslots not at the start of a frame, the CM transmits the first sample of the CP of the first symbol at the starting frame boundary. For fine ranging, the CM starts transmission one OFDMA symbol (including the CP) after the start of the first OFDMA frame of the ranging opportunity. Request opportunities in subslots not at the start of a frame are referenced to the symbol boundary at the start of the subslot.

7.4.12.1 Upstream Time Reference of an OFDMA Frame

The upstream time reference is defined as the first sample of the first symbol of an OFDMA frame, pointed to by the dashed arrow of figure 7.16. The parameter N_{FFT} refers to the length of the FFT duration which is either 2 048 or 4 096, and the parameter N_{CP} is the length of the configurable cyclic prefix. The sample rate is 102,4 Msamples per second.

7.4.12.2 Upstream Time Reference of an OFDMA Symbol

The upstream time reference for construction of each OFDMA symbol is defined as the first sample of each FFT duration of each OFDMA symbol, pointed to by the dotted arrow of figure 7.16.

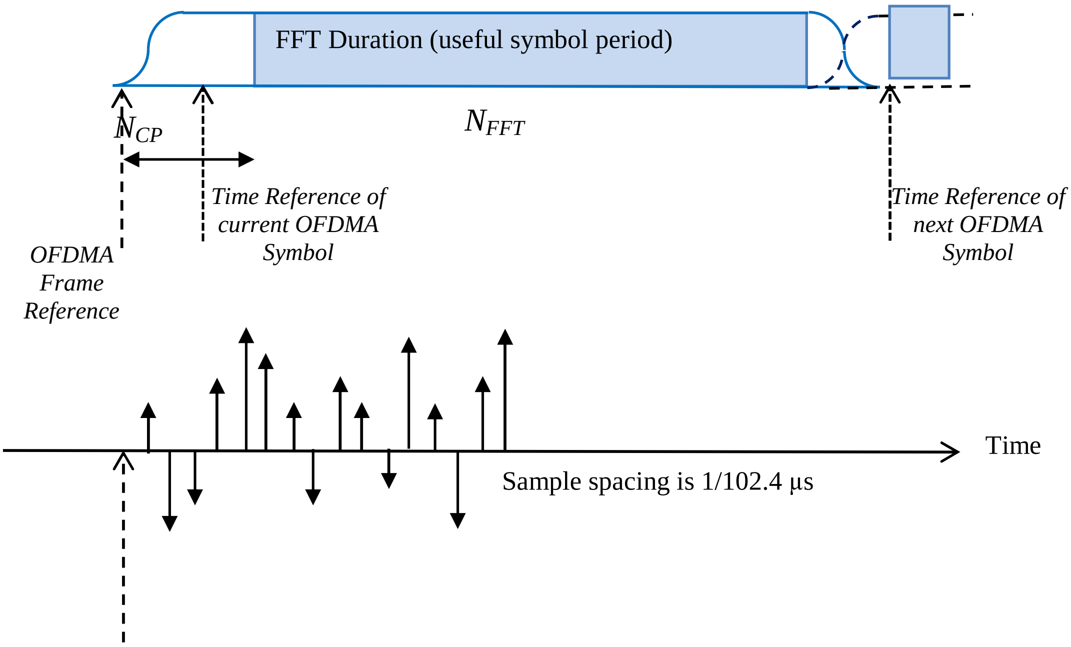

Figure 7.16: Time References for OFDMA Symbol and Frame

7.4.13 Fidelity Requirements

7.4.13.0 CM Transmit Channel Set

A DOCSIS 3.1 CM is required to generate up to 8 channels of legacy DOCSIS plus up to 2 OFDMA channels as defined in clauses 6.2.1 and 7.2.1.

A CM's Transmit Channel Set (TCS) is the combination of legacy channels and OFDMA channels being transmitted by the CM.

With BW_{legacy} being the combined occupied bandwidth of the legacy channel(s) in its TCS, and BW_{OFDMA} being the combined occupied bandwidth of the OFDMA channel(s) in its TCS, the CM is said to have N_{eq} = ceil(BW_{legacy} (MHz)/1,6 MHz) + ceil(BW_{OFDMA} (MHz)/1,6 MHz) "equivalent DOCSIS channels" in its TCS. BW_{OFDMA} (MHz) is the sum of the bandwidth of the maximum modulated spectrum of all the OFDMA channels that are active.

"Equivalent channel power" of a legacy DOCSIS channel refers to the power in 1,6 MHz of spectrum.

The "equivalent channel power" of an OFDMA channel is the average power of the OFDMA subcarriers of the channel normalized to 1,6 MHz bandwidth. This equivalent channel power of an OFDMA channel is denoted as $P_{1.6}$. The TCS has N legacy (N from zero to eight) plus zero, one, or two OFDMA channels, but also is described as having N_{eq} number of equivalent DOCSIS channels.

Each channel in the TCS is described by its reported power, which is the power of the channel when it is fully granted. Each channel is also characterized by its "equivalent channel power," which is the channel power normalized to 1,6 MHz (Power Spectral Density of the average power of the channel multiplied by 1,6 MHz).

7.4.13.1 Maximum Scheduled Minislots

While transmitting on the large upstream spectrum supported by DOCSIS 3.1, a CM can encounter large upstream attenuation and can have a power deficit when attempting to reach the CMTS receiver at the nominal OFDMA channel set power. A CMTS has several options in dealing with such CMs: it can limit the TCS to the channel set that will enable the CM to reach the CMTS receiver at the nominal set power; it can assign the CM a profile which includes reduced modulation depth enabling proper reception even at lower received power; or, it can operate that CM under Maximum Scheduled Minislots (MSM).

Complete control of MSM operation is under the CMTS. The CMTS does not inform the CM when it decides to assign it to MSM operation in a specific OFDMA channel. Instead, the CMTS instructs the CM to transmit with a higher power spectral density than the CM is capable of with a 100 % grant. In addition, the CMTS limits the number of minislots concurrently scheduled to the CM, such that the CM is not given transmit opportunities on that OFDMA channel that will result in overreaching its reported transmission power capability. The CMTS also optimizes the power used by the CM to probe an OFDMA channel, for which the CM is operating under MSM, by using the Power Control parameters in the Probe Information Element directed to that CM. Refer to clause 7.4.13.3 for details.

Note that when operating under MSM, it is expected that a CM that normally meets the fidelity and performance requirements will only exhibit graceful degradation. Refer to clause 7.4.13.2 for details. Also note that the CMTS is expected to discriminate between a CM whose fidelity degrades gracefully and a CM whose fidelity does not, and provide the capability to disallow a CM whose fidelity does not degrade gracefully from operating under MSM.

7.4.13.2 Transmit Power Requirements

7.4.13.2.0 Transmit Power Requirements Overview

The transmit power requirements are a function of the number and occupied bandwidth of the OFDMA and legacy channels in the TCS. The minimum highest value of the total power output of the CM P_{max} is 65 dBmV, although higher values are allowed. The total maximum power is distributed among the channels in the TCS, based on equal power spectral density (PSD) when the OFDMA and legacy channels are fully granted to the CM. Channels can then be reduced in power from their max power that was possible based on equal PSD allocated (with limits on the reduction). This ensures that each channel can be set to a power range (within the DRW) between its maximum power, $P_{1.6h}$, and minimum power, $P_{1.6low}$, and that any possible transmit grant combination can be accommodated without exceeding the transmit power capability of the CM.

Maximum equivalent channel power ($P_{1.6Max}$) is calculated as $P_{1.6Max} = P_{max}$ dBmV - $10\log_{10}(N_{eq})$.

NOTE: For DOCSIS 3.1 $P_{1.6hi} = P_{1.6Max.}$

For a CM operating in DOCSIS 3.1 mode, even on a SC-QAM channel, the CMTS shall limit the commanded $P_{1.6Max}$ to no more than 53,2 dBmV+ (P_{max} - 65) if the bandwidth of the modulated spectrum is ≤ 24 MHz. This enforces a maximum power spectral density of P_{max} dBmV per 24 MHz. This limit on power spectral density does not apply for a CM operating in DOCSIS 3.0 mode, where the fidelity requirements are the DOCSIS 3.0 fidelity requirements and not the fidelity requirements of the DOCSIS 3.1 mode.

SC-QAM channels that are 6,4 MHz in BW have a power of $P_{1.6r\_n}$ + 6 dB. The minimum equivalent channel power ($P_{1.6Min}$) for OFDMA channels with non-boosted pilots is $P_{1.6Min}$ = 17 dBmV. For OFDMA channels with boosted pilots, prior to pre-equalization, $P_{1.6Min}$ is 18 dBmV with 50 kHz subcarrier spacing and 17,5 dBmV with 25 kHz spacing. For Initial Ranging and before completion of Fine Ranging, transmissions may use power per subcarrier which is as much as 9 dB lower than indicated by $P_{1.6Min}$. These transmissions are prior to any data grant transmissions from the CM and as such the CM analog and digital gain balancing may be optimized for these transmissions. These transmissions, while possibly at very low power, are acceptable because, for example, they are not requiring severe underloading of a DAC.

The CMTS should not command the CM to set $P_{1.6r\_n}$ on any channel in the TCS to a value higher than the top of the DRW or lower than the bottom of the DRW, unless the CMTS is using MSM to accommodate a need to increase the PSD for the channel.

If the CMTS does issue such a command, fidelity and performance requirements on the CM do not apply. Note that when operating under MSM, it is expected that a CM that normally meets the fidelity and performance requirements, will only exhibit graceful degradation. Also note that the CMTS is expected to discriminate between a CM that does meet such expectations and a CM that does not, and provide the capability to disallow a CM that does not meet such expectations to operate under MSM.

If the CM is commanded to transmit on any channel in the TCS at a value higher than $P_{1.6hi}$ or lower than $P_{1.6low\_n}$, the cable modem indicates an error condition by setting the appropriate bit in the SID field of RNG-REQ messages for that channel until the error condition is cleared [4].

The CMTS sends transmit power level commands and pre-equalizer coefficients to the CM [4] to compensate for upstream plant conditions. The top edge of the DRW is set to a level, $P_{1.6load\_min\_set}$, close to the highest $P_{1.6}$ transmit channel to optimally load the DAC. In extreme tilt conditions, some of the channels will be sent commands to transmit at lower $P_{1.6}$ values that use up a significant portion of the DRW. Additionally, the pre-equalizer coefficients of the OFDMA channels will also compensate for plant tilts. The CMTS normally administers a DRW of 12 dB [4] which is sufficient to accommodate plant tilts of up to 10 dB from lower to upper edge of the upstream band. Since the fidelity requirements are specified in flat frequency conditions from the top of the DRW (Dynamic Range Window), it is desirable to maintain CM transmission power levels as close to the top of the DRW as possible. When conditions change sufficiently to warrant it, a global reconfiguration time should be granted and the top of the DRW adjusted to maintain the best transmission fidelity and optimize system performance.

7.4.13.2.1 Transmit Power Requirements with Multiple Transmit Channel Mode Enabled

The following requirements apply with Multiple Transmit Channel mode enabled. Requirements with Multiple Transmit Channel mode disabled are addressed in [5].

The CM shall support varying the amount of transmit power. Requirements are presented for

1) range of reported transmit power per channel;
2) step size of power commands;
3) step size accuracy (actual change in output power per channel compared to commanded change); and
4) absolute accuracy of CM output power per channel.

The protocol by which power adjustments are performed is defined in [4]. Such adjustments by the CM shall be within the ranges of tolerances described below. A CM shall confirm that the transmit power per channel limits are met after a RNG-RSP is received for each of the CM's active channels that is referenced and indicate that an error has occurred in the next RNG-REQ messages for the channel until the error condition is cleared [4].

An active channel for a CM is defined as any channel which has been assigned to the CM's Transmit Channel Set either in Registration Response Message or a DBC-REQ Message, and prior to registration the channel in use is an (the) active channel. After registration, the set of "active channels" is also called the Transmit Channel Set. If the CMTS needs to add, remove, or replace channels in the CM's Transmit Channel Set, it uses the Dynamic Bonding Request (DBC-REQ) Message with Transmit Channel Configuration encodings to define the new desired Transmit Channel Set. Note that the set of channels actually bursting upstream from a CM is a subset of the active channels on that CM; often one or all active channels on a CM will not be bursting, but such quiet channels are still "active channels" for that CM.

Transmit power per channel is defined as the average RF power in the occupied bandwidth (channel width), assuming equally likely QAM symbols, measured at the F-connector of the CM as detailed below. Reported power for a SC-QAM channel is expressed in terms of $P_{1.6r\_n}$, i.e. the actual channel power for a 6,4 MHz channel would be 6 dB higher than the reported power (neglecting reporting accuracy). For a 1,6 MHz channel, the actual channel power would be equal to the reported power (neglecting reporting accuracy). For SC-QAM signals, the reported power differs from the actual power in one respect for modulations other than 64-QAM, and that is the constellation gain as defined in tables 6-7, table 6-8 and table 6-9 of [5].

Reported transmit power for an OFDMA channel is also expressed as of $P_{1.6r\_n}$ and is defined as the average RF power of the CM transmission in the OFDMA channel, when transmitting in a grant comprised of 64 subcarriers at 25 kHz subcarriers or 32 subcarriers at 50 kHz subcarriers, for OFDMA channels which do not use boosted pilots. For OFDMA channels which have boosted pilots and 50 kHz subcarrier spacing, reported power is 1 dB higher than the average RF power of the CM transmission with a probe comprised of 32 subcarriers. For OFDMA channels which have boosted pilots and 25 kHz subcarrier spacing, reported power is 0,5 dB higher than the average RF power of the CM transmission with a probe comprised of 64 subcarriers. The additions to the probe power account for the maximum possible number of boosted pilots in each OFDMA symbol when the OFDMA channel uses boosted pilots. Equivalent channel power for an OFDMA channel is the reported transmit power normalized to 1,6 MHz bandwidth (four minislots). Total transmit power is defined as the sum of the transmit power per channel of each channel transmitting a burst at a given time.

The CM's actual transmitted power per equivalent channel shall be within ±2 dB of the target power, $P_{1.6r\_n}$, with Pre-Equalization off taking into account whether pilots are present and symbol constellation values. The CM's target transmit power per channel shall be variable over the range specified in clause 7.4.13.3. The CM's target transmit power per channel may be variable over a range that extends above the maximum levels specified in clause 7.4.13.3. Note that all fidelity requirements specified in clause 7.4.13.3 still apply when the CM is operating over its extended transmit power range, but the fidelity requirements do not apply when the CM is commanded to transmit at power levels which exceed the top of the DRW.

The CM communicates its ability to transmit above 65 dBmV to the CMTS via a modem capability encoding as defined in [4]. When the CM indicates that it supports the extended range and the CMTS disables this capability. The CM shall use the default value of 65 dBmV for the maximum transmit power. The CMTS shall use the default value of 65 dBmV for the maximum transmit power.

With Multiple Transmit Channel mode enabled, let $P_{1.6load} = P_{1.6max} - P_{1.6r}$, for each channel, using the definitions for $P_{1.6max}$ and $P_{1.6r}$ in the following clauses of 7.4.13. The channel corresponding to the minimum value of $P_{1.6load}$ is called the highest loaded channel, and its value is denoted as $P_{1.6load\_1}$, in the present document even if there is only one channel in the Transmit Channel Set. A channel with high loading has a low $P_{1.6load\_n}$ value; the value of $P_{1.6load\_n}$ is analogous to an amount of back-off for an amplifier from its max power output, except that it is normalized to 1,6 MHz of bandwidth. A channel has lower power output when that channel has a lower loading (more back-off) and thus a higher value of $P_{1.6load\_n}$. Note that the highest loaded channel is not necessarily the channel with the highest transmit power, since a channel's max power depends on the bandwidth of the channel. The channel with the second lowest value of $P_{1.6load}$ is denoted as the second highest loaded channel, and its loading value is denoted as $P_{1.6load\_2}$; the channel with the n^{th} lowest value of $P_{1.6load}$ is the n^{th} highest loaded channel, and its loading value is denoted as $P_{1.6load\_n}$.

$P_{1.6load\_min\_set}$ defines the upper end of the DRW for the CM with respect to $P_{1.6hi}$ for each channel. $P_{1.6load\_min\_set}$ will limit the maximum power possible for each active channel to a value less than $P_{1.6max}$ when $P_{1.6load\_min\_set}$ is greater than zero. $P_{1.6load\_min\_set}$ is a value commanded to the CM from the CMTS when the CM is given a TCC in registration and RNG-RSP messages [4]. $P_{1.6load\_min\_set}$, $P_{1.6load\_ n}$, $P_{1.6hi}$, $P_{1.6 r\_n,}$ etc. are defined only when Multiple Transmit Channel mode is enabled.

The CMTS should command the CM to use a non-negative value for $P_{1.6load\_min\_set}$ such that $P_{1.6hi} - P_{1.6load\_min\_set} \geq P_{1.6low\_n}$ for each active channel, or equivalently:

$$0 \leq P_{1.6load\_min\_set} \leq P_{1.6hi} - P_{1.6low\_n}.$$

A value is computed, P_{low_multi}, which sets the lower end of the transmit power DRW for that channel, given the upper end of the range which is determined by $P_{1.6load\_min\_set}$.

$$P_{1.6low\_multi} = P_{1.6hi} - P_{1.6load\_min\_set} - 12 \text{ dB}$$

The effect of $P_{1.6low\_multi}$ is to restrict the dynamic range required (or even allowed) by a CM across its multiple channels, when operating with multiple active channels.

Unless the CMTS is using MSM to accommodate a need to increase the PSD for the channel in which case the fidelity performance of the CM is potentially degraded, the CMTS should command a $P_{1.6r\ \_n}$ consistent with the $P_{1.6load\_min\_set}$ assigned to the CM and with the following limits:

$$P_{1.6load\_min\_set} \leq P_{1.6hi} - P_{1.6r\_n} \leq P_{1.6load\_min\_set} + 12 \text{ dB}$$

and the equivalent:

$$P_{1.6hi} - (P_{1.6load\_min\_set} + 12 \text{ dB}) \leq P_{1.6r\ \_n} \leq P_{1.6hi} - P_{1.6load\_min\_set}$$

When the CMTS sends a new value of $P_{1.6load\_min\_set}$ to the CM, there is a possibility that the CM will not be able to implement the change to the new value immediately, because the CM may be in the middle of bursting on one or more of its upstream channels at the instant the command to change $P_{1.6load\_min\_set}$ is received at the CM. Some amount of time may elapse before the CMTS grants global reconfiguration time to the CM. Similarly, commanded changes to $P_{1.6r\_n}$ may not be implemented immediately upon reception at the CM if the n^{th} channel is bursting.

Commanded changes to $P_{1.6r\_n}$ may occur simultaneously with the command to change $P_{1.6load\_min\_set}$. The CMTS should not issue a change in $P_{1.6load\_min\_set}$ after commanding a change in $P_{1.6r\_n}$ until after also providing a sufficient reconfiguration time on the n^{th} channel. The CMTS should not issue a change in $P_{1.6load\_min\_set}$ after commanding a prior change in $P_{1.6load\_min\_set}$ until after also providing a global reconfiguration time for the first command. Also, the CMTS should not issue a change in $P_{1.6r\_n}$ until after providing a global reconfiguration time following a command for a new value of $P_{1.6load\_min\_set}$ and until after providing a sufficient reconfiguration time on the n^{th} channel after issuing a previous change in $P_{1.6r\_n}$. In other words, the CMTS is to avoid sending consecutive changes in $P_{1.6r\_n}$ and/or $P_{1.6load\_min\_set}$ to the CM without a sufficient reconfiguration time for instituting the first command. When a concurrent new value of $P_{1.6load\_min\_set}$ and change in $P_{1.6r\_n}$ are commanded, the CM may wait to apply the change in $P_{1.6r\_n}$ at the next global reconfiguration time (i.e. concurrent with the institution of the new value of $P_{1.6load\_min\_set}$) rather than applying the change at the first sufficient reconfiguration time of the n^{th} channel. The value of $P_{1.6load\_min\_set}$ which applies to the new $P_{1.6r\_n}$ is the concurrently commanded $P_{1.6load\_min\_set}$ value. If the change to $P_{1.6r\_n}$ falls outside the DRW of the old $P_{1.6load\_min\_set}$, then the CM shall wait for the global reconfiguration time to apply the change in $P_{1.6r\_n}$.

Unless the CMTS is using MSM to accommodate a need to increase the PSD for the channel in which case the fidelity performance of the CM is potentially degraded, the CMTS should not command the CM to increase the per channel transmit power if such a command would cause $P_{1.6load\_\ n}$ for that channel to drop below $P_{1.6load\_min\_set}$. Note that the CMTS can allow small changes of power in the CM's highest loaded channel, without these fluctuations impacting the transmit power dynamic range with each such small change. This is accomplished by setting $P_{1.6load\_min\_set}$ to a smaller value than normal, and fluctuation of the power per channel in the highest loaded channel is expected to wander.

The CMTS should not command a change of per channel transmit power which would result in $P_{1.6r\_n}$ falling below the DRW, $P_{1.6r\_n} < P_{low\_multi}$. Unless the CMTS is using MSM to accommodate a need to increase the PSD for the channel in which case the fidelity performance of the CM is potentially degraded, the CMTS should not command a change in $P_{1.6load\_min\_set}$ such that existing values of $P_{1.6r\_n}$ would fall outside the new DRW.

The following paragraphs define the CM and CMTS behavior in cases where there are DRW violations due to indirect changes to $P_{1.6max,}$ or addition of a new channel with incompatible parameters without direct change of $P_{1.6r\_n}$ or $P_{1.6load\_min\_set}$.

Adding or removing a channel from the TCS changes $P_{1.6max}$ of the existing active channels (due to different $P_{1.6max}$ calculation tables for different number of active channels). Prior to changing the channels in the TCS, the CMTS should change $P_{1.6r\_n}$ of all current active channels, if necessary, to fit in the new expected DRW.

When adding a new active channel to the transmit channel set, the new channel's power is calculated according to the offset value defined in TLV 46.8.4 [4], if it is provided. The CMTS should not set an offset value that will result in a $P_{1.6r\_n}$ for the new channel outside the DRW. In the absence of the TLV, the new channel's power is initially set by the CM at the minimum allowable power, i.e. the bottom of the DRW.

If the CMTS changes the symbol rate for a SC-QAM channel, the CM maintains constant total power for the channel by adjusting $P_{1.6r\_n}$ for that channel. The CMTS should not send a UCD change for the n^{th} active channel that violates $P_{1.6load\_n} - P_{1.6load\_min\_set} \geq 0$. If a UCD changes the $P_{1.6r\_n}$ for the n^{th} active channel [4], the CM adjusts $P_{1.6r\_n}$ for that channel, when the change count of the MAP matches the change count in the new UCD. The CM adjusts $P_{1.6load\_n}$ at the time the MAP change count matches the new UCD change count, and calculates a new $P_{1.6r\_n}$ and target power for the channel, to be applied for bursts granted in the MAP with change count matching the change count in the new UCD. The spurious performance requirements of clauses 7.4.12 do not apply if $P_{1.6load\_n}$ becomes negative for a channel or if $P_{1.6load\_n} - P_{load\_min\_set}$ becomes negative for a channel.

The CM's actual transmitted power per channel, within a burst, shall be constant to within 0,1 dB peak to peak, even in the presence of power changes on other active channels. This excludes the amplitude variation, which is theoretically present due to QAM amplitude modulation, pulse shaping, pre-equalization, and varying number of allocated minislots with OFDMA or varying number of spreading codes in S-CDMA channels.

The CM shall support the transmit power calculations defined in clause 7.4.13.3.

7.4.13.3 OFDMA Transmit Power Calculations

In OFDMA mode the CM determines its target transmit power per channel $P_{1.6t\_n}$, as follows, for each channel which is active. Define for each active channel, for example, upstream channel n:

$P_{1.6r\_n}$ = reported power level (dBmV) of CM for channel n.

$P_{1.6hi} = P_{max}/N_{eq}$.

$P_{1.6low\_n} = P_{1.6low}$.

The CM updates its reported power per channel in each channel by the following steps:

1) $P_{1.6r\_n} = P_{1.6r\_n} + \Delta P$ //Add power level adjustment (for each channel) to reported power level for each channel.
   The CMTS should ensure the following:

2) $P_{1.6r\_n} \leq P_{1.6hi}$ //Clip at max power limit per channel unless the CMTS is using MSM to accommodate a need to increase the PSD for the channel in which case the fidelity performance of the CM is potentially degraded.

3) $P_{1.6r\_n} \geq P_{1.6low}$ //Clip at min power limit per channel.

4) $P_{1.6r\_n} \geq P_{1.6min\_multi}$ //Power per channel from this command would violate the set DRW.

5) $P_{1.6r} \leq P_{1.6hi} - P_{1.6load\_min\_set}$ //Power per channel from this command would violate the set DRW, unless the CMTS is using MSM to accommodate a need to increase the PSD for the channel in which case the fidelity performance of the CM is potentially degraded.

The CM then transmits each data subcarrier with target power:

$$P_{t\_sc\_i} = P_{1.6r\_n} - P_{1.6delta\_n} + \text{Pre-Eq}_i - 10\log(\text{number\_of subcarriers in 1,6 MHz } \{32 \text{ or } 64\})$$

where Pre-Eq_i is the magnitude of the i^{th} subcarrier pre-equalizer coefficient (dB), and $P_{1.6delta\_n}$ equals 0 dB for non-boosted channels, 0,5 dB for boosted channels with 25 kHz subcarrier spacing and 1 dB for boosted channels with 50 kHz subcarrier spacing.

That is, the reported power for channel n, normalized to 1,6 MHz, minus compensation for the presence of boosted pilots plus the pre-equalization for the subcarrier, less a factor taking into account the number of subcarriers in 1,6 MHz.

$Probe_{delta\_n}$ for the n^{th} OFDMA channel is the change in subcarrier power for probes compared to subcarrier power for data depending on the mode as defined in [4] in addition to Pre-Equalization on or off.

The CM transmits probes with the same target power as given above + $Probe_{delta\_n}$ when Pre-EQ is enabled for probes in the P-MAP which provides the probe opportunity:

$$P_{t\_sc\_i} = P_{1.6r\_n} - P_{1.6delta\_n} + Probe_{delta\_n} + Pre\text{-}Eq_i - 10\ \log(\text{number\_of subcarriers in 1,6 MHz } \{32 \text{ or } 64\})$$

When the Pre_EQ is disabled for the probe opportunity in the P-MAP, the CM then transmits probe subcarrier with target power:

$$P_{t\_sc\_i} = P_{1.6r\_n} - P_{1.6delta\_n} + Probe_{delta\_n} - 10\ \log_{10}(\text{number\_of subcarriers in 1,6 MHz } \{32 \text{ or } 64\})$$

where $P_{1.6delta\_n}$ equals 0 dB for non-boosted channels, 0,5 dB for boosted channels with 25 kHz subcarrier spacing and 1 dB for boosted channels with 50 kHz subcarrier spacing.

That is, the reported power for channel n, normalized to 1,6 MHz, minus compensation for the presence of boosted pilots less a factor taking into account the number of subcarriers in 1,6 MHz.

For Channels with boosted pilots, the CM then transmits each boosted pilot with target power:

$$P_{t\_pilot} = P_{1.6r\_n} - P_{1.6delta\_n} + Pre\text{-}Eq_i - 10\ \log_{10}(\text{number\_of subcarriers in 1,6 MHz } \{32 \text{ or } 64\}) + 10\log_{10}(3)$$

where $Pre\text{-}Eq_i$ is the magnitude of the i^{th} subcarrier pre-equalizer coefficient (dB), and $P_{1.6delta\_n}$ equals 0,5 dB for 25 kHz subcarrier spacing and 1 dB for 50 kHz subcarrier spacing.

That is, the reported power for channel n, normalized to 1,6 MHz, minus compensation for the presence of boosted pilots plus the pre-equalization for the subcarrier, less a factor taking into account the number of subcarriers in 1,6 MHz, plus the pilot boost in power by a factor of 3.

The total transmit power in channel n, P_{t_n}, in a frame is the sum of the individual transmit powers P_{t,sc_i} of each subcarrier in channel n, where the sum is performed using absolute power quantities [non-dB domain].

The transmitted power level in channel n varies dynamically as the number and type of allocated subcarriers varies.

7.4.13.4 Global Reconfiguration Time for OFDMA Channels

"Global reconfiguration time" is defined as the inactive time interval provided between active upstream transmissions, which simultaneously satisfies the requirement in clause 6.2.20 for all TDMA channels in the TCS and the requirement in clause 6.2.20 for all S-CDMA channels in the TCS and the requirement here for OFDMA.

Global "quiet" across all active channels requires the intersection of ungranted burst intervals across all active OFDMA channels to be at least 20 microseconds. Even with a change or re-command of $P_{1.6load\_min\_set}$, the CM shall be able to transmit consecutive bursts as long as the CMTS allocates at least one frame in between bursts, across all OFDMA channels in the Transmit Channel Set, where the quiet lapses in each channel contain an intersection of at least 20 microseconds. (From the end of a burst on one channel to the beginning of the next burst on any channel, there is to be at least 20 microseconds duration to provide a "global reconfiguration time" for OFDMA channels.)

With mixed channels operating in the upstream, the Global Reconfiguration times for DOCSIS 3.0 CMs remain the same as defined in [5]. For DOCSIS 3.1 CMs operating in a mixed upstream, the requirements for the intersection of quiet times for all channels in the TCS is that it be at least 10 microseconds plus 96 symbols on each of the SC-QAM channels.

With Multiple Transmit Channel mode enabled, the CMTS should provide global reconfiguration time to a CM before (or concurrently as) the CM has been commanded to change any upstream channel transmit power by ±3 dB cumulative since its last global reconfiguration time.

7.4.13.5 OFDMA Fidelity Requirements

7.4.13.5.0 TCS Fidelity Requirements

The following requirements assume that any pre-equalization is disabled, unless otherwise noted.

When channels in the TCS are commanded to the same equivalent powers, the reference signal power in the "dBc" definition is to be interpreted as the measured equivalent to the reported power as defined in clause 7.4.13.2.1. When channels in the TCS are commanded to different equivalent channel powers, the commanded total power of the transmission is computed, and a difference is derived compared to the commanded total power which would occur if all channels had the same $P_{1.6}$ as the highest equivalent channel power in the TCS, whether or not the channel with the largest equivalent channel power is included in the grant. Then this difference is added to the measured total transmit power to form the reference signal power for the "dBc" spurious emissions requirements.

For purposes of the OFDMA fidelity requirements, even if Maximum Scheduled Minislots (MSM) is enabled in a CM, the 100 % Grant Spectrum for spurious emissions calculations is unchanged by application of MSM.

7.4.13.5.1 Spurious Emissions

7.4.13.5.1.0 Spurious Emissions Requirements

The noise and spurious power generated by the CM shall not exceed the levels given in tables 7.6, 7.7 and 7.8. Up to five discrete spurs can be excluded from the emissions requirements listed in tables 7.6, 7.7 and 7.8 and have to be less than -42 dBc relative to a single subcarrier power level.

SpurFloor is defined as:

$$\text{SpurFloor} = \max\{ -57 + 10 * \log_{10}(100\ \%\ \text{Grant Spectrum}/192\ \text{MHZ}), -60\}\ \text{dBc}$$

Under-grant Hold Number of Users is defined as:

$$\text{Under-grant Hold Number of Users} = \text{Floor}\{ 0{,}2 + 10\text{\^{}}((-44 - \text{SpurFloor})/10) \}$$

Under-grant Hold Bandwidth is defined as:

$$\text{Under-grant Hold Bandwidth} = (100\ \%\ \text{Grant Spectrum})/(\text{Under-grant Hold Number of Users})$$

When Multiple Transmit Channel mode is enabled, these spurious performance requirements only apply when the CM is operating within certain ranges of values for P_{load_n}, for n = 1 to the number of upstream channels in the TCS, and for granted bandwidth of Under-grant Hold Bandwidth or larger; where P_{load_1} the highest loaded channel in the present document (i.e. its power is the one closest to P_{hi}).

When a modem is transmitting over a bandwidth of less than Under-grant Hold Bandwidth the spurious emissions requirement limit is the power value (in dBmV), corresponding to the specifications for the power level associated with a grant of bandwidth equal to Under-grant Hold Bandwidth. In addition, when a modem is transmitting over a bandwidth such that the total power of the modem is less than 17 dBmV, but other requirements are met, then the modem spurious emissions requirements limits is the power value (in dBmV) computed with all conditions and relaxations factored in, plus an amount X dB where:

$$X\ \text{dB} = 17\ \text{dBmV} - \text{modem transmit power}$$

When Multiple Transmit Channel mode is enabled and there are two or more channels in the TCS, the CM's spurious performance requirements shall be met only when the equivalent DOCSIS channel powers ($P_{1.6}$) are within 6 dB of $P_{load\_min\_set}$ ($P_{load\_min\_set} + 6 \geq P_{load\_n} \geq P_{load\_min\_set}$).

Further, the CM's spurious emissions requirements shall be met with two or more channels in the TCS only when $P_{load\_1} = P_{load\_min\_set}$. When $P_{load\_1} < P_{load\_min\_set}$, the spurious emissions requirements in absolute terms are relaxed by $P_{load\_1} - P_{load\_min\_set}$.

When a modem is transmitting with any equivalent DOCSIS channel power with loading $P_{load\_min\_set}$ + 16 dB or lower ($P_{load} > P_{load\_min\_set}$ + 16 dB), the spurious emissions requirement limits are not enforced.

With Multiple Transmit Channel mode enabled, the spurious performance requirements do not apply to any upstream channel from the time the output power on any active upstream channel has varied by more than ±3 dB since the last global reconfiguration time through the end of the next global reconfiguration time changes, excluding transmit power changes due to UCD-induced change in P_{hi} [4].

In table 7.6, inband spurious emissions includes noise, carrier leakage, clock lines, synthesizer spurious products, and other undesired transmitter products. It does not include ISI. The measurement bandwidth for inband spurious for OFDM is equal to the Subcarrier Clock Frequency (25 kHz or 50 kHz) and is not a synchronous measurement. The signal reference power for OFDMA inband spurious emissions is the total transmit power measured and adjusted (if applicable) as described in clause 7.4.13.5, and then apportioned to a single data subcarrier.

For S-CDMA and TDMA, the measurement bandwidth is the modulation rate (e.g. 1 280 to 5 120 kHz), and the requirement is ≤ -50 dBc. All requirements expressed in dBc are relative to the largest equivalent DOCSIS channel power in the TCS, whether it is being transmitted or not.

The measurement bandwidth is 160 kHz for the Between Bursts (none of the channels in the TCS is bursting) specs of table 7.6, except where called out as 4 MHz or 250 kHz. The signal reference power for Between Bursts transmissions is the total transmit power measured and adjusted (if applicable) as described in clause 7.4.13.5.

The Transmitting Burst specs apply during the minislots granted to the CM (when the CM uses all or a portion of the grant), and for 20 µs before and after the granted minislot for OFDMA. The Between Bursts specs apply except during a used grant of minislots on any active channel for the CM, and 20 us before and after the used grant for OFDMA. The signal reference power for Transmitting Burst transmissions, other than inband, is the total transmit power measured and adjusted (if applicable) as described in clause 7.4.13.5.

For the purpose of spurious emissions definitions, a granted burst refers to a burst of minislots to be transmitted at the same time from the same CM; these minislots are not necessarily contiguous in frequency.

For Initial Ranging and before completion of Fine Ranging, spurious emissions requirements use tables 7.6, 7.7 and 7.8, with 100 % Grant Spectrum equal to the bandwidth of the modulation spectrum of the transmission, and if transmissions use subcarrier power which is X dB lower than indicated by $P_{1.6low}$, then the spurious emissions requirements in absolute terms are relaxed by X dB.

Spurious emissions requirements for grants of 10 % or less of the TCS (100 % grant spectrum) may be relaxed by 2 dB in an amount of spectrum equal to:

- measurement BW * ceil(10 % of the TCS / measurement BW)
- anywhere in the whole upstream spectrum for emission requirements specified in tables 7.7 and 7.8

A 2 dB relief applies in the measurement bandwidth. This relief does not apply to between bursts emission requirements.

Table 7.6: Spurious Emissions

Parameter	Transmitting Burst	Between Bursts (see note 3)
Inband	-45 dBc OFDMA 100 % grant (see notes 4, 5 and 6) -51 dBc OFDMA 5 % grant (see notes 4, 5 and 6) -50 dBc S-CDMA/TDMA (see note 7)	-72 dBc
Adjacent Band	See table 7.8	-72 dBc
Within the upstream operating range 5 - 42 MHz or 5 - 85 MHz, or 5 - 204 MHz (excluding assigned channel, adjacent channels)	See table 7.7	-72 dBc
For the case where the upstream operating range is 5 - 42 MHz: CM Integrated Spurious Emissions Limits (all in 4 MHz, includes discretes) (see note 1)		

Parameter	Transmitting Burst	Between Bursts (see note 3)
42 to 54 MHz 54 to 60 MHz 60 to 88 MHz 88 to 1 218 MHz	-40 dBc -35 dBmV -40 dBmV -45 dBmV	-26 dBmV -40 dBmV -40 dBmV max(-45 dBmV, -40 dB ref downstream) (see note 2)
For the case where the upstream operating range is 5 - 42 MHz: CM Discrete Spurious Emissions Limits (see note 1) 42 to 54 MHz 54 to 88 MHz 88 to 1 218 MHz	-50 dBc -50 dBmV -50 dBmV	-36 dBmV -50 dBmV -50 dBmV
For the case where the upstream operating range is 5 - 85 MHz: CM Integrated Spurious Emissions Limits (all in 4 MHz, includes discrete spurs) (see note 1) 85 to 108 MHz 85 to 108 MHz (Should) 108 to 136 MHz 136 to 1 218 MHz	-45 dBc -50 dBc -40 dBmV -45 dBmV	-31 dBmV -36 dBmV -40 dBmV max(-45 dBmV, -40 dB ref downstream) (see note 2)
For the case where the upstream operating range is 5 - 85 MHz: CM Discrete Spurious Emissions Limits (see note 1) 85 to 108 MHz 108 to 1 218	-50 dBc -50 dBmV	-36 dBmV -50 dBmV
For the case where the upstream operating range is 5 - 204 MHz: CM Integrated Spurious Emissions Limits (all in 4 MHz, includes discrete spurs) (see note 1) 204 to 258 MHz 204 to 258 MHz (Should) 258 to 1 218 MHz	-50 dBc -60 dBc -45 dBmV	-72 dBc -72 dBc max(-45 dBmV, -40 dB ref downstream) (see note 2)
For the case where the upstream operating range is 5 - 204 MHz: CM Discrete Spurious Emissions Limits (see note 1) 204 to 258 MHz 258 to 1 218 MHz	-50 dBc -50 dBmV	-36 dBmV -50 dBmV
NOTE 1: These spec limits exclude a single discrete spur related to the tuned received channel; this single discrete spur is to be no greater than -40 dBmV. NOTE 2: "dB ref downstream" is relative to the received downstream signal level. Some spurious outputs are proportional to the receive signal level. NOTE 3: Relative to the previous transmission. NOTE 4: Up to 5 subcarriers within the entire upstream bandwidth with discrete spurs may be excluded from the MER calculation if they fall within transmitted minislots. These 5 spurs are the same spurs that may be excluded for spurious emissions and not an additional or different set. NOTE 5: This value is to be met when $P_{load} = P_{load\_min\_set}$. NOTE 6: Receive equalization is allowed if an MER test approach is used, to take ISI out of the measurement; measurements other than MER-based to find spurs or other unwanted power may be applied to this requirement. NOTE 7: Also see MER requirement.		

7.4.13.5.1.1 Spurious Emissions in the Upstream Frequency Range

Table 7.7 lists the required spurious level in a measurement interval. The initial measurement interval at which to start measuring the spurious emissions (from the transmitted burst's modulation edge) is 400 kHz from the edge of the transmission's modulation spectrum. Measurements should start at the initial distance and be repeated at increasing distance from the carrier until the upstream band edge or spectrum adjacent to other modulated spectrum is reached.

For OFDMA transmissions with non-zero transmit windowing, the CM shall meet the required performance measured within the 2,0 MHz adjacent to the modulated spectrum using slicer values from a CMTS burst receiver or equivalent, synchronized to the downstream transmission provided to the CM.

In the rest of the spectrum, the CM shall meet the required performance measured with a bandpass filter (e.g. an unsynchronized measurement).

For OFDMA transmissions with zero transmit windowing, CM shall meet the required performance using synchronized measurements across the complete upstream spectrum.

For legacy transmissions, the measurement is performed in the indicated bandwidth and distance from the transmitted legacy channel edge.

Spurious emissions allocation for far out spurious emissions =

Round{ SpurFloor + $10*\log_{10}$(Measurement bandwidth/Under-grant hold Bandwidth), 0,1}.

For transmission bursts with modulation spectrum less than the Under-grant Hold Bandwidth, the spurious power requirement is calculated as above, but increased by $10*\log_{10}$ (Under-grant Hold Bandwidth/Grant Bandwidth).

Table 7.7: Spurious Emissions Requirements in the Upstream Frequency Range for Grants of Under-grant Hold Bandwidth and Larger (see note 1)

100 % Grant Spectrum (MHz)	SpurFloor (dBc)	Under-grant Hold #Users	Under-grant Hold Bandwidth (MHz)	Measurement Bandwidth (MHz)[2]	Specification in the Interval (dBc)
Up to 64 [e.g. 22 MHz]	-60,0	40	100 % Grant Spectrum/40 [0,55 MHz]	1,6	Round{ SpurFloor + $10*\log_{10}$(Measurement Bandwidth/Under-grant Hold Bandwidth), 0,1} [-55,4]
[e.g. 46 MHz]			[1,15 MHz]		[-58,6]
Greater than 64, up to 96	-60,0	40	100 % Grant Spectrum/40	3,2	Round{ SpurFloor + $10*\log_{10}$(Measurement Bandwidth/Under-grant Hold Bandwidth), 0,1
[e.g. 94 MHz]			[2,35 MHz]		[-58,7]
Greater than 96, up to 192	max{ -57 + $10*\log_{10}$(100 % Grant Spectrum/192 MHZ), -60}	Floor{ 0,2 + 10^((-44 - SpurFloor)/10) }	100 % Grant Spectrum)/(Under-grant Hold Number of Users)	9,6	Round{ SpurFloor + $10*\log_{10}$(Measurement Bandwidth/Under-grant Hold Bandwidth), 0,1}
[e.g.142 MHz]	[-58,3]	[27]	[5,3]		[-55,7]
[e.g. 190 MHz]	[-57,0]	[20]	[9,5]		[-57,0]
Greater than 192	max{ -57 + $10*\log_{10}$(100 % Grant Spectrum/192 MHZ), -60}	Floor{ 0,2 + 10^((-44 - SpurFloor)/10) }	100 % Grant Spectrum)/(Under-grant Hold Number of Users)	12,8	Round{ SpurFloor + $10*\log_{10}$ (Measurement Bandwidth/Under-grant Hold Bandwidth), 0,1}
[e.g. 200 MHz]	[-56,8]	[19]	[10,5]		[-55,9]
NOTE 1: Spurious Emissions Requirements in the Upstream Frequency Range Relative to the Per Channel Transmitted Burst Power Level for Each Channel for Grants of Under-grant Hold Bandwidth and Larger. NOTE 2: The measurement bandwidth is a contiguous sliding measurement window.					

The CM shall control transmissions such that within the measurement bandwidth of table 7.7, spurious emissions measured for individual subcarriers contain no more than +3 dB power larger than the required average power of the spurious emissions in the full measurement bandwidth. When non synchronous measurements are made, only 25 kHz measurement bandwidth is used.

For legacy transmissions, and optionally for OFDMA transmissions, bandpass measurements rather than synchronous measurements may be applied.

As an example illustrating use of table 7.7 for legacy channels, consider a TCS with a single 1,6 MHz SC-QAM channel (1,28 Msym/s) and a single 6,4 MHz SC-QAM channel (5,12 Msym/s). The grant BW is then 8,0 MHz and the 100 % grant spectrum is 8,0 MHz. So the spurfloor is -60 dBc, and the emissions specification is -51 dBc or equivalently -44 dBr.

As an example illustrating the smaller measurement bandwidth requirements, consider 94 MHz 100 % grant spectrum, with -58,7 dBc spurious emissions allowed in 3,2 MHz measurement bandwidth, with the measurement bandwidth starting as close as 400 kHz from the modulation edge of the transmitted burst. If the subcarrier spacing is 25 kHz, there are 128 subcarriers in the 3,2 MHz measurement bandwidth. Each subcarrier has, on average, a requirement of -58,7 dBc - 21,1 dB = -79,8 dBc, but the requirement is relaxed to -79,8 dBc + 3 dB = -76,8 dBc (noting that -21,07 dB corresponds to $1/128^{th}$). The under-grant hold bandwidth is 2,35 MHz for this example. When a 100 % grant has 65 dBmV transmit power, a grant of 2,4 MHz has 49,1 dBmV power. With a single OFDMA channel and its 100 % grant power at 65 dBmV, the spurious emissions requirement with a grant of 2,4 MHz, measured in 25 kHz is 49,1 dBmV - 76,8 dBc = -27,7 dBmV. -76,8 dBc corresponds to -57,1 dBr for this example (since 2,35 MHz/25 kHz is a factor of 94, or 19,7 dB).

7.4.13.5.1.2 Adjacent Channel Spurious Emissions

Spurious emissions from a transmitted burst may occur in adjacent spectrum, which could be occupied by a legacy carrier of any allowed modulation rate or by OFDMA subcarriers.

Table 7.8 lists the required adjacent channel spurious emission levels when there is a transmitted burst with bandwidth at the Under-grant Hold Bandwidth. The measurement is performed in an adjacent channel interval of 400 kHz adjacent to the transmitted burst modulation spectrum. For OFDMA transmissions, the measurement is performed starting on an adjacent subcarrier of the transmitted spectrum (both above and below), using the slicer values from a CMTS burst receiver or equivalent synchronized to the downstream transmission provided to the CM. For legacy transmissions, the measurement is performed in an adjacent channel interval of 400 kHz bandwidth adjacent to the transmitted legacy channel edge.

Firstly, it should be noted that the measurement bandwidth for table 7.8 is less than the measurement bandwidths in table 7.7. Thus comparing the two tables in terms of the specification "dBc" values requires appropriate scaling. Secondly, table 7.8 provides specification "dBc" only for grants of a specific amount for each row, while table 7.7 provides "dBc" specification for grants of all sizes from the Under-grant Hold Bandwidth to 100 %.

For transmission bursts with modulation spectrum less than the Under-grant Hold Bandwidth, the spurious power requirement is calculated as above, but increased by $10*\log_{10}$(Under-grant Hold Bandwidth/Grant Bandwidth).

For transmission bursts with modulation spectrum greater than the Under-grant Hold Bandwidth, the spurious power requirement in the adjacent 400 kHz is calculated by converting the requirement to absolute power "dBmV" for a grant of precisely Under-grant Hold Bandwidth from table 7.8, and similarly computing the absolute power "dBmV" from table 7.7 for a grant equal to:

The Given Grant - The Under-grant Hold Bandwidth.

Then the absolute power calculated from table 7.7 is scaled back in exact proportion of 400 kHz compared to the measurement bandwidth in table 7.7. Then the power from table 7.8 is added to the scaled apportioned power from table 7.7 to produce the requirement for the adjacent 400 kHz measurement with a larger grant than the Under-grant Hold Bandwidth. The requirement for adjacent spurious power in adjacent 400 kHz is:

P1(Grant Bandwidth - Under-grant Hold Bandwidth) = absolute power derived from table 7.7. (dBmV)

P2(Under-grant Hold Bandwidth) = absolute power derived from table 7.8. (dBmV)

$P1_{scaled}$ = P1 * (0,4 MHz)/(Measurement Bandwidth (MHz) used in table 7.7). (dBmV)

$P_{spec\_limit} = P1_{scaled} + P2$ (dBmV)

The CM shall control transmissions such that within the measurement bandwidth of table 7.8, spurious emissions measured for individual subcarriers contain no more than +3 dB power larger than the required average power of the spurious emissions in the full measurement bandwidth. For legacy transmissions, and optionally for OFDMA transmissions, bandpass measurements rather than synchronous measurements may be applied.

Table 7.8: Adjacent Channel Spurious Emissions Requirements Relative to the Per Channel Transmitted Burst Power Level for Each Channel

100 % Grant Spectrum (MHz)	SpurFloor (dBc)	Under-grant Hold #Users	Under-grant Hold Bandwidth (MHz)	Measurement Bandwidth (MHz)	Specification in Adjacent 400 kHz With Grant of Under-grant Hold Bandwidth (dBc)
Up to 64 [e.g. 22 MHz]	-60,0	40	100 % Grant Spectrum/40 [0,55 MHz]	0.4 MHz	Round{10*log$_{10}$(((10^(SpurFloor/10)) + (10^(-57/10))) x(0,4 MHz/Under-grant Hold Bandwidth)),0,1} [-56,6]
[Ex: 46 MHz]			[1,15 MHz]		[-59,8]
Greater than 64, up to 96	-60,0	40	100 % Grant Spectrum/40	0,4 MHz	Round{10*log$_{10}$(((10^(SpurFloor/10)) + (10^(-57/10))) x(0,4 MHz/Under-grant Hold Bandwidth)), 0,1}
[Ex 94 MHz]			[2,35 MHz]		[-62.9]
Greater than 96	max{ -57 + 10*log$_{10}$(100 % Grant Spectrum/192 MHZ), -60} Round nearest 0,1 dB	Floor{ 0,2 + 10^((-44 - SpurFloor)/10) }	100 % Grant Spectrum)/Under-grant Hold Number of Users	0,4 MHz	Round{10*log$_{10}$(((10^(SpurFloor/10)) + (10^(-57/10))) x(0,4 MHz/Under-grant Hold Bandwidth)), 0,1}
[e.g. 142 MHz]	[-58,3]	[27]	[5,3]		[-65,8]
[e.g. 190 MHz]	[-57,0]	[20]	[9,5]		[-67,7]
[e.g. 200 MHz]	[-56,8]	[19]	[10,5]		[-68,1]

7.4.13.5.2 Spurious Emissions During Burst On/Off Transients

The CM shall control spurious emissions prior to and during ramp-up, during and following ramp-down, and before and after a burst.

The CM's on/off spurious emissions, such as the change in voltage at the upstream transmitter output, due to enabling or disabling transmission, shall be no more than 50 mV.

The CM's voltage step shall be dissipated no faster than 4 µs of constant slewing. This requirement applies when the CM is transmitting at +55 dBmV or more per channel on any channel.

At backed-off transmit levels, the CM's maximum change in voltage shall decrease by a factor of 2 for each 6 dB decrease of power level in the highest power active channel, from +55 dBmV per channel, down to a maximum change of 3,5 mV at 31 dBmV per channel and below. This requirement does not apply to CM power-on and power-off transients.

7.4.13.5.3 MER Requirements

7.4.13.5.3.0 Modulation Error Ratio Overview

Transmit modulation error ratio (TxMER or just MER) measures the cluster variance caused by the CM during upstream transmission due to transmitter imperfections. The terms "equalized MER" and "unequalized MER" refer to a measurement with linear distortions equalized or not equalized, respectively, by the test equipment receive equalizer. The requirements in this clause refer only to unequalized MER, as described for each requirement. MER is measured on each modulated data subcarrier and non-boosted pilot (MER is computed based on the unboosted pilot power) in a minislot of a granted burst and averaged for all the subcarriers in each minislot.

MER includes the effects of Inter-Carrier Interference (ICI), spurious emissions, phase noise, noise, distortion, and all other undesired transmitter degradations with an exception for a select number of discrete spurs impacting a select number of subcarriers. MER requirements are measured with a calibrated test instrument that synchronizes to the OFDMA signal, applies a receive equalizer in the test instrument that removes MER contributions from nominal channel imperfections related to the measurement equipment, and calculates the value. The equalizer in the test instrument is calculated, applied and frozen for the CM testing. Receiver equalization of CM linear distortion is not provided; hence this is considered to be a measurement of unequalized MER, even though the test equipment contains a fixed equalizer setting.

7.4.13.5.3.1 Definitions

MER is defined as follows for OFDMA. The transmitted RF waveform at the F connector of the CM (after appropriate down conversion) is filtered, converted to baseband, sampled, and processed using standard OFDMA receiver methods, with the exception that receiver equalization is not provided. The processed values are used in the following formula. No external noise (AWGN) is added to the signal.

The carrier frequency offset, carrier amplitude, carrier phase offset, and timing will be adjusted during each burst to maximize MER as follows:

- One carrier amplitude adjustment common for all subcarriers and OFDM symbols in burst.
- One carrier frequency offset common for all subcarriers resulting in phase offset ramping across OFDM symbols in bursts.
- One timing adjustment resulting in phase ramp across subcarriers.
- One carrier phase offset common to all subcarriers per OFDM symbol in addition to the phase ramp.

MER_i is computed as an average of all the subcarriers in a minislot for the i^{th} minislot in the OFDMA grant:

$$\mathrm{MER_i}\ (\mathrm{dB}) = 10 \cdot \log_{10}\left(\frac{E_{avg}}{\frac{1}{N}\sum_{j=1}^{N}\left(\frac{1}{M}\sum_{k=1}^{M}\left|e_{j,k}\right|^2\right)}\right)$$

where:

- E_{avg} is the average constellation energy for equally likely symbols,
- M is the number of symbols averaged,
- N is the number of subcarriers in a minislot,
- $e_{j,k}$ is the error vector from the jth subcarrier in the minislot and kth received symbol to the ideal transmitted QAM symbol of the appropriate modulation order.

A sufficient number of OFDMA symbols shall be included in the time average so that the measurement uncertainty from the number of symbols is less than other limitations of the test equipment.

MER with a 100 % grant is defined as the condition when all OFDMA minislots and any legacy channels in the transmit channel set are granted to the CM.

MER with a 5 % grant is defined as the condition when less than or equal to 5 % of the available OFDMA minislots and no legacy channels have been granted to the CM.

7.4.13.5.3.2 Requirements

Unless otherwise stated, the CM shall meet or exceed the following MER limits over the full transmit power range, all modulation orders, all grant configurations and over the full upstream frequency range.

The following flat channel measurements with no tilt (see table 7.9) are made after the pre-equalizer coefficients have been set to their optimum values. The receiver uses best effort synchronization to optimize the MER measurement.

Table 7.9: Upstream MER Requirements (with Pre-Equalization)

Parameter	Value
MER (100 % grant)	Each minislot MER ≥ 44 dB (Notes 1 and 2)
MER (5 % grant)	Each minislot MER ≥ 50 dB (Notes 1 and 2)
Pre-equalizer constraints	Coefficients set to their optimum values
NOTE 1: Up to 5 subcarriers within the entire upstream bandwidth with discrete spurs may be excluded from the MER calculation if they fall within transmitted minislots. These 5 spurs are the same spurs that may be excluded for spurious emissions and not an additional or different set. NOTE 2: This value is to be met when Pload = Pload_min_set.	

The following flat channel measurements (table 7.10) are made with the pre-equalizer coefficients set to unity and no tilt and the receiver implementing best effort synchronization. For this measurement, the receiver may also apply partial equalization. The partial equalizer is not to correct for the portion of the CM's time-domain impulse response greater than 200 ns or frequency-domain amplitude response greater than +1 dB or less than -3 dB from the average amplitude. An additional 1 dB attenuation in the amplitude response is allowed in the upper 10 % of the specified passband frequency. It is not expected that the partial equalizer is implemented on CMTS receiver. A partial equalizer could be implemented offline via commercial receivers or simulation tools.

Table 7.10: Upstream MER Requirements (no Pre-Equalization)

Parameter	Value
MER (100 % grant)	Each minislot MER ≥ 40 dB (notes 1 and 2)
MER (5 % grant)	Each minislot MER ≥ 40 dB (notes 1 and 2)
Pre-equalizer constraints	Pre-equalization not used
NOTE 1: Up to 5 subcarriers within the entire upstream bandwidth with discrete spurs may be excluded from the MER calculation if they fall within transmitted minislots. These 5 spurs are the same spurs that may be excluded for spurious emissions and not an additional or different set. NOTE 2: This value is to be met when Pload = Pload_min_set.	

7.4.14 Cable Modem Transmitter Output Requirements

The CM shall output an RF Modulated signal with characteristics delineated in table 7.11.

Table 7.11: CM Transmitter Output Signal Characteristics

Parameter	Value
Frequency	Support and be configurable to a permitted subset (see clause 7.2.3 for allowed combinations) of the following list of frequency ranges: 5 - 42 MHz 5 - 65 MHz 5 - 85 MHz 5 - 117 MHz 5 - 204 MHz NOT to cause harmful interference above these frequencies for any configured option may support > 204 MHz
Signal Type	OFDMA
Maximum OFDMA Channel Bandwidth	96 MHz
Minimum OFDMA Occupied Bandwidth	6,4 MHz for 25 kHz subcarrier spacing 10 MHz for 50 kHz subcarrier spacing
Number of Independently configurable OFDMA channels	Minimum of 2
Subcarrier Channel Spacing	25 kHz, 50 kHz
FFT Size	50 kHz: 2 048 (2K FFT); 1 900 Maximum active subcarriers 25 kHz: 4 096 (4K FFT); 3 800 Maximum active subcarriers
Sampling Rate	102,4 MHz (96 MHz Block Size)
FFT Time Duration	40 µs (25 kHz subcarriers) 20 µs (50 kHz subcarriers)
Modulation Type	BPSK, QPSK, 8-QAM, 16-QAM, 32-QAM, 64-QAM, 128-QAM, 256-QAM, 512-QAM, 1024-QAM, 2048-QAM, 4096-QAM
Bit Loading	Variable from minislot to minislot Constant for subcarriers of the same type in the minislot Support zero valued subcarriers per profile and minislot.
Pilot Tones	14 data patterns and 2 subslot patterns, minislot subcarrier size and length dependent - see clause 7.4.17.
Cyclic Prefix Options	Samples µs 96 0,9375 128 1,25 160 1,5625 192 1,875 224 2,1875 256 2,5 288 2,8125 320 3,125 384 3,75 512 5,0 640 6,25
Windowing Size Options	Samples µs 0 0 32 0,3125 64 0,625 96 0,9375 128 1,25 160 1,5625 192 1,875 224 2,1875 Raised cosine absorbed by CP
Level	CM shall be capable of transmitting a total average output power of 65 dBmV CM may be capable of transmitting a total average output power greater than 65 dBmV
Output Impedance	75 ohms
Output Return Loss	> 6 dB 5-f_{max} MHz (42/65/85/117/204 MHz) > 6 dB f_{max} - 1 218 MHz > 6 dB f_{max} - 1,794 GHz for CMs that can receive up to 1,794 GHz
Connector	F connector per [17] or [19]

7.4.15 CMTS Receiver Capabilities

7.4.15.1 CMTS Receiver Input Power Requirements

The CMTS Upstream Demodulator shall operate with an average input signal level, including ingress and noise to the upstream demodulator, up to 31 dBmV.

The CMTS shall be settable according to table 7.12 for intended received power normalized to 6,4 MHz of bandwidth.

The CMTS Upstream demodulator shall operate within its defined performance specifications with received bursts within the ranges defined in table 7.12 of the set power.

Table 7.12: Upstream Channel Demodulator Input Power Characteristics

Modulation	Minimum Set Point	Maximum Set Point	Range
QPSK	-4 dBmV	10 dBmV	-9 / +3
8-QAM	-4 dBmV	10 dBmV	-9 / +3
16-QAM	-4 dBmV	10 dBmV	-9 / +3
32-QAM	-4 dBmV	10 dBmV	-9 / +3
64-QAM	-4 dBmV	10 dBmV	-9 / +3
128-QAM	0 dBmV	10 dBmV	-9 / +3
256-QAM	0 dBmV	10 dBmV	-9 / +3
512-QAM	0 dBmV	10 dBmV	-3 / +3
1024-QAM	0 dBmV	10 dBmV	-3 / +3
2048-QAM	7 dBmV	10 dBmV	-3 / +3
4096-QAM	10 dBmV	10 dBmV	-3 / +3

7.4.15.2 CMTS Receiver Error Ratio Performance in AWGN Channel

The required level for CMTS upstream post-FEC error ratio is defined for AWGN as less than or equal to 10^{-6} PER (packet error ratio) with 1 500 byte Ethernet packets. This clause describes the conditions at which the CMTS is required to meet this error ratio.

Implementation loss of the CMTS receiver shall be such that the CMTS achieves the required error ratio when operating at a CNR as shown in table 7.13, under input load and channel conditions as follows:

- A single transmitter, pre-equalized and ranged
- A single OFDMA 96 MHz channel.
- Ranging with same CNR and input level to CMTS as with data bursts, and with 8-symbol probes.
- Any valid transmit combination (frequency, subcarrier clock frequency, transmit window, cyclic prefix, OFDMA frame length, interleaving depth, pilot patterns, etc.) as defined in the present document.
- Input power level per constellation is the minimum set point as defined in table 7.12.
- OFDMA phase noise and frequency offset are at the max limits as defined for the CM transmission specification.
- Ideal AWGN channel with no other artifacts (reflections, burst noise, tilt, etc.).
- Large grants consisting of several 1 500 Bytes.
- CMTS is allowed to construct MAPs according to its own scheduler implementation.

Table 7.13: CMTS Minimum CNR Performance in AWGN Channel

Constellation	CNR[1,2] (dB)	Set Point	Offset
QPSK	11,0	-4 dBmV	0 dB
8-QAM	14,0	-4 dBmV	0 dB
16-QAM	17,0	-4 dBmV	0 dB
32-QAM	20,0	-4 dBmV	0 dB
64-QAM	23,0	-4 dBmV	0 dB
128-QAM	26,0	0 dBmV	0 dB
256-QAM	29,0	0 dBmV	0 dB
512-QAM	32,5	0 dBmV	0 dB
1024-QAM	35,5	0 dBmV	0 dB
2048-QAM	39,0	7 dBmV	0 dB
4096-QAM	43,0	10 dBmV	0 dB

NOTE 1: CNR is defined here as the ratio of average signal power in occupied bandwidth to the average noise power in the occupied bandwidth given by the noise power spectral density integrated over the same occupied bandwidth.

NOTE 2: Channel CNR is adjusted to the required level by measuring the source inband noise including phase noise component and adding the required delta noise from an external AWGN generator.

7.4.16 Ranging

7.4.16.0 Ranging Procedure

Ranging in DOCSIS 3.1 is divided into three steps, as illustrated in figure 7.17:

- *Initial ranging* is used by the CMTS to identify a new admitting CM and for coarse power and timing ranging.
- *Fine ranging* is used after initial ranging has been completed, to fine-tune timing and power.
- *Probing* is used during admission and steady state for pre-equalization configuration and periodic TX power and time-shift ranging.

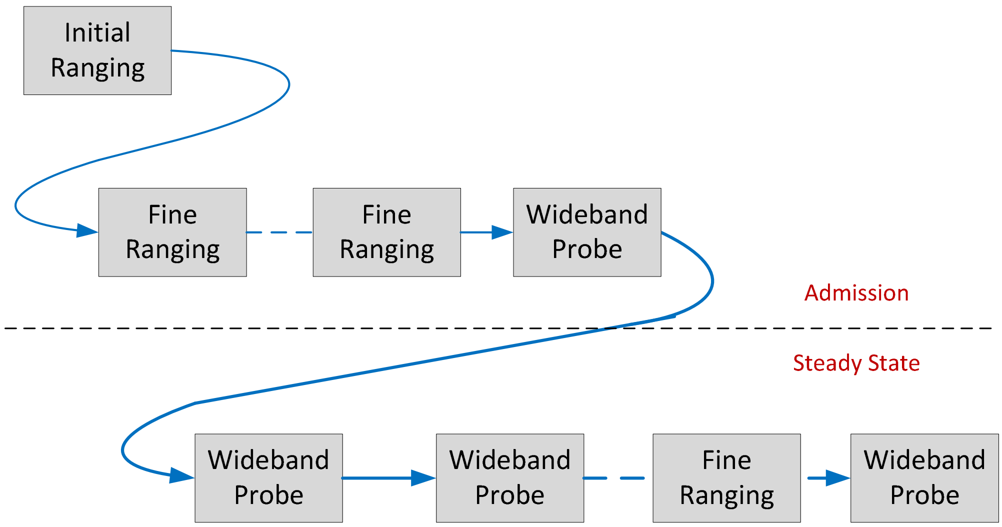

Figure 7.17: Ranging Steps

7.4.16.1 Initial Ranging

7.4.16.1.1 Initial Ranging Zone

The initial ranging zone consists of N by M contiguous minislots in the upstream frame. N and M are configured by the CMTS.

Minislots in the initial ranging zone do NOT carry pilots; all the FFT grid points in the initial ranging zone are used for the initial ranging signal, as illustrated in figure 7.18.

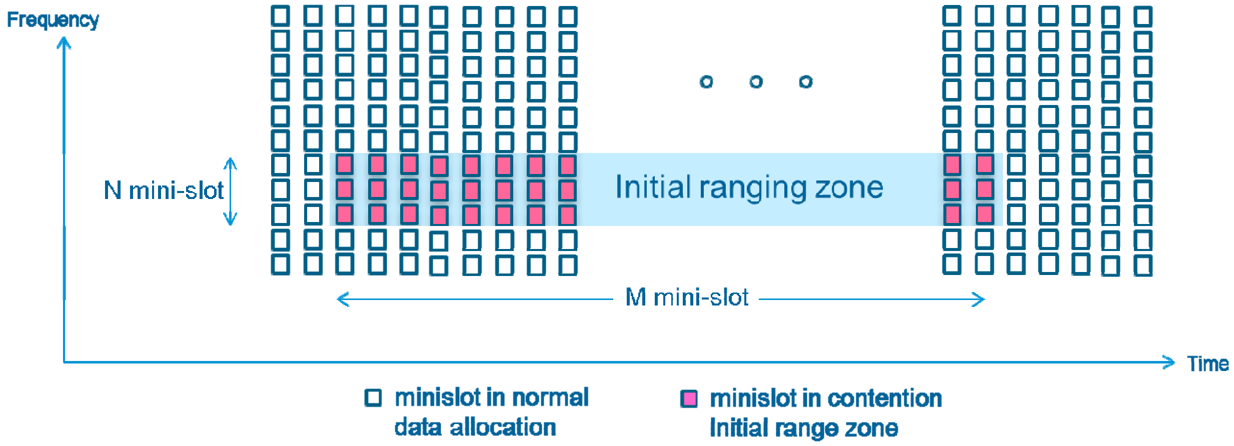

Figure 7.18: Initial Ranging Zone

7.4.16.1.2 Initial Ranging Signal

The initial ranging signal consists of a preamble sequence and a data part, as illustrated in figure 7.19. The data part is O-INIT-RNG-REQ as described in [i.8].

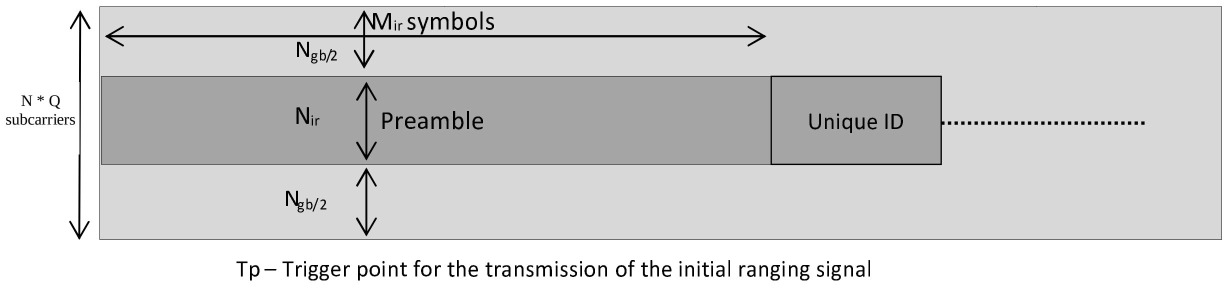

Figure 7.19: Initial Ranging Signal

When allocating an initial ranging opportunity, the CMTS shall allocate contiguous minislots within an OFDMA frame. See [4] for how an initial ranging opportunity that spans multiple OFDMA frames is specified in a MAP message. The preamble sequence is a BPSK binary sequence configured by the CMTS and sent by the CM. The length of the sequence is configured by the CMTS, and the bits contained in the sequence are configured by the CMTS.

The data portion of the initial ranging signal is the O-INIT-RNG-REQ message as described in [4]. It is composed of a 6-byte MAC address, plus a 1-byte downstream channel ID and 24 CRC bits. It is LDPC (128,80) encoded and randomized as described in the sections below.

To generate the 24-bit CRC the CM shall convert the 7 message bytes into a bit stream in MSB-first order. The CM shall use the first bit of the bit stream to be the MSB of the first byte of the 6-byte MAC address and the last bit of the bit stream shall be the LSB of the downstream channel ID. The 24 bits of CRC will be computed and appended to this bit stream as defined in Annex E to create the 80-bit sequence to be LDPC encoded.

The preamble sequence and the O-INIT-REG-REQ are duplicated and sent in a special structure of pair of symbols with identical BPSK content as described in figure 7.20.

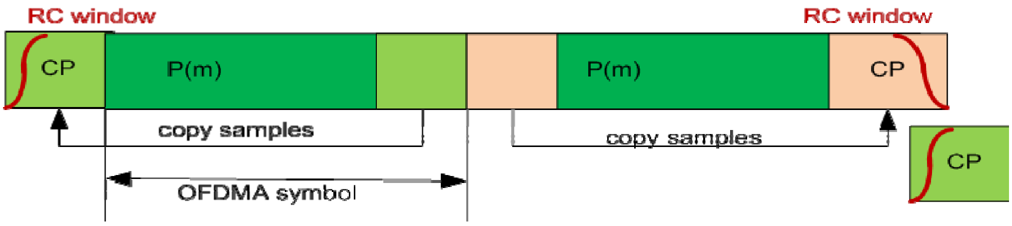

Figure 7.20: Initial Ranging Admission Slot Structure

A block diagram of the initial ranging signal processing in the transmitter is described in figure 7.21:

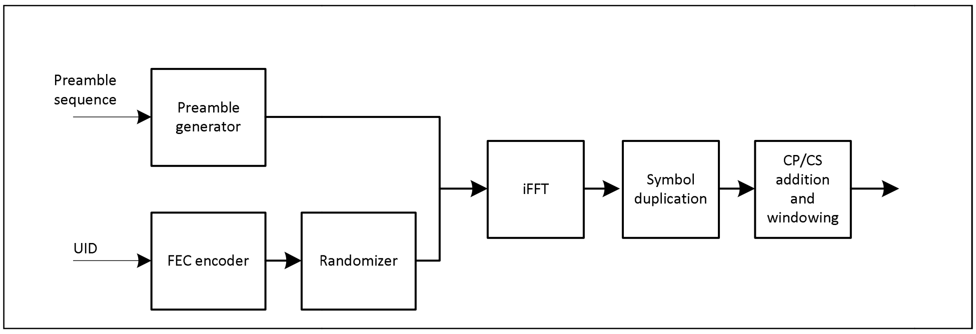

Figure 7.21: Block Diagram of Initial Ranging Transmitter Processing

7.4.16.1.3 Preamble Construction

The CMTS shall configure the BPSK Preamble sequence and its length L_p, with the limitations described in clause 7.4.16.2.6 and the number of subcarriers, N_{ir}, to be used for the transmission of the initial ranging signal.

The CMTS shall allocate minislots for the initial ranging signal, comprising of the number of subcarriers N_{ir} and an appropriate guard band.

The CM shall construct the preamble part of the initial ranging signal by converting the preamble sequence bits into BPSK symbols. The preamble is comprised of M_{ir} symbols each with N_{ir} subcarriers.

The CM shall convert the first N_{ir}*M_{ir} bits in the preamble sequence into N_{ir}*M_{ir} BPSK symbols in the following order: The first N_{ir} BPSK symbols are written to the N_{ir} subcarriers of the first preamble symbol starting from the lowest subcarrier, the next N_{ir} BPSK symbols to the N_{ir} subcarriers of the second preamble symbol and the last N_{ir} BPSK symbols to the N_{ir} subcarriers of the last (the M_{ir}) preamble symbol.

7.4.16.1.4 FEC for the Initial Ranging Data

The CM shall encode the 80 bit O-INIT-RNG-REQ message using the LDPC (128,80) encoder as described below.

A puncturing encoder consists of two steps. The first step encodes the input bit sequence with an encoder of the mother code. The second step, called puncturing step, deletes one or more bits from the encoded codeword.

The mother code is a rate ½ (160, 80) binary LDPC code. A parity check matrix of the mother code is represented by table 7.14, where sub-matrix size (lifting factor) L = 16, see clause 7.4.4.2 for the compact definition of parity check matrix.

Table 7.14: (160,80) LDPC code Parity Check Matrix

1	11	10	12	7	9	-	-	-	-
2	1	14	15	14	14	12	-	-	-
0	9	3	2	-	-	11	7	-	-
6	8	-	10	3	-	-	10	4	-
12	13	11	-	0	-	-	-	5	2

Let the information bits sent to the mother code encoder be denoted by $(a_0, \cdots, a_{79})$ and let the encoder output be denoted by $(a_0, \cdots, a_{79}, b_{80}, \cdots, b_{159})$, where $b_{80}, \cdots, b_{159}$ are parity-check bits. The bits to be deleted by the puncturing step are (also see figure 7.22)

- Period 1: 16 consecutive bits $a_0, \cdots, a_{15}$
- Period 2: 16 consecutive bits $b_{144}, \cdots, b_{159}$

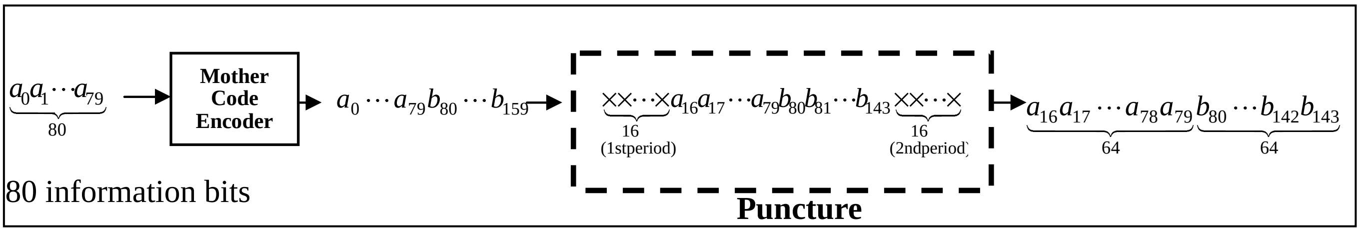

Figure 7.22: LDPC Two-Period Puncturing Encoder for Initial Ranging FEC

7.4.16.1.5 Padding and Randomizing

The CM shall pad and randomize the 128 encoded bits as described below.

The CM shall calculate the number of symbols required to transmit the INIT-RNG-REQ message as follows:

Nuid_sym = ceiling (128/N_{ir}) where N_{ir} is the number of subcarriers allocated for the INIT-RNG-REQ message.

CM shall pad the remaining bits with ones if the total number of bits (Nbits = Nuid_sym*N_{ir}) is greater than 128.

The CM shall randomize the 128 encoded bits and the padding bits as described in clause 7.4.5, with the randomizer initialized at the beginning of the 128 encoded bits. The randomized bits are converted to BPSK symbols as defined in (BPSK constellation) and are appended to the preamble sequence for transmission.

The CM shall add the BPSK symbols to the data part of the initial ranging signal in the following order: The first Nir BPSK symbols written to the N_{ir} subcarriers of the first symbol of the data part, the next N_{ir} BPSK symbols to the next data symbol, until all BPSK symbols are written vertically symbol by symbol. First BPSK symbol is written to the lowest indexed subcarrier of a data symbol.

7.4.16.1.6 Symbol Duplicating Cyclic Prefix and Windowing

The CM shall repeat each Initial Ranging OFDMA symbol twice. A cyclic prefix of N_{cp} samples is appended before the first repeated OFDMA symbol. A cyclic suffix of N_{rcp} ($N_{rcp} = N_{cp} + N_{rp}$) samples is appended after the second repeated OFDMA symbol.

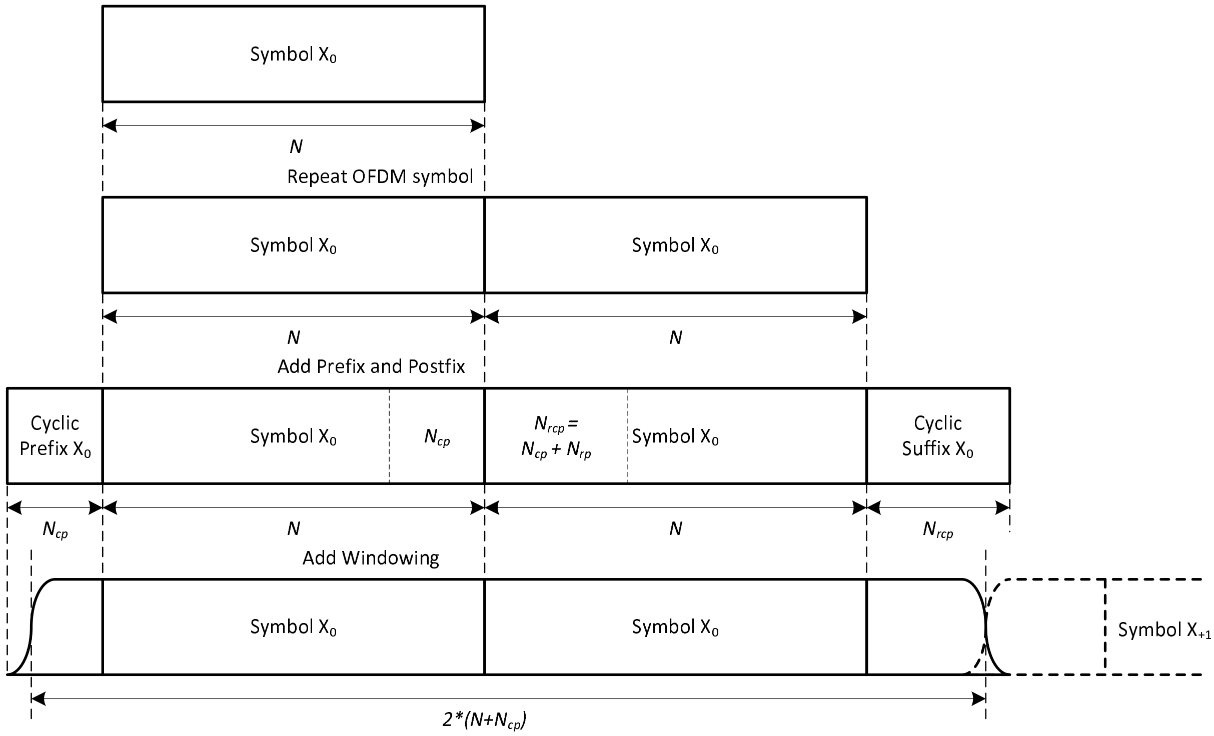

Figure 7.23: Initial Ranging Symbol Pair Structure

Table 7.15: Cyclic Prefix and Roll-off Samples for Initial Ranging

Cyclic Prefix Samples (N_{cp})	Roll-Off Samples (N_{rp})
96	96
128	128
160	160
192	192
224	224
256	224
288	224
320	224
384	224
512	224
640	224

7.4.16.2 Fine Ranging

7.4.16.2.0 Fine Ranging Operations

This clause describes fine ranging operations for the CM transmitter.

Fine ranging is used by the CMTS for the second step of the admission of a new CM process, following successful initial ranging. During this step, a fine ranging signal is transmitted by a new CM joining the network, according to transmission parameters provided by the CMTS. When it receives the fine ranging signal, the CMTS is able to fine-tune the joining CM's transmission power and transmission timing.

At the end of the fine ranging step, the CMTS can assign transmission opportunities to the new CM, using optimal transmission power, without interfering with coexisting transmitters on the same OFDMA frame.

7.4.16.2.1 Fine Ranging Signal

Figure 7.24 illustrates a fine ranging signal.

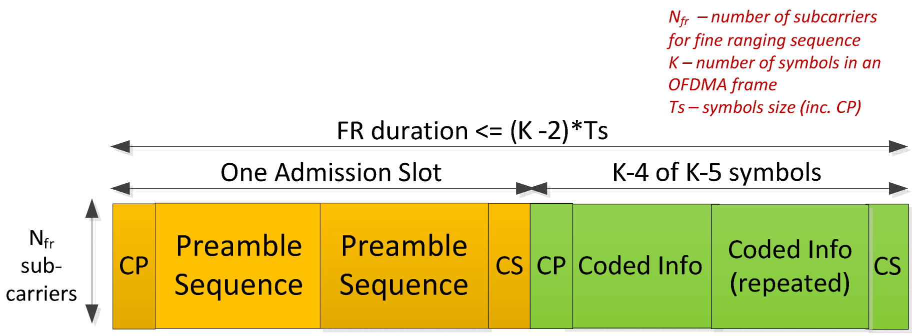

Figure 7.24: Fine Ranging Signal

Fine ranging is a narrowband signal integrated into a single data OFDMA frame. It is comprised of two parts: a BPSK preamble sequence of one pair of preamble symbols (as defined in clause 7.4.16.1.2 for the initial ranging), and 34 bytes of FEC-encoded data spread over two or more OFDMA symbols. The data part of the fine ranging signal is QPSK-modulated and FEC encoded. The data part has a similar structure to the duplicated pair of symbols (refer to clause 7.4.16.1.2 for the initial ranging data structure).

The CM shall transmit the fine ranging signal when allocated to it, with the following configurable parameters:

- Time shift
- TX power
- Number of minislots allocated to the fine ranging signal (number of minislots incorporate the fine ranging signal plus the required guardband as described in figure 7.25)
- Number of subcarriers for the fine ranging signal
- Preamble sequence

The CM shall use the first portion of the preamble sequence defined for the Initial Ranging signal for the BPSK PRBS sequence of the fine ranging.

7.4.16.2.2 Transmission of the Fine Ranging Signal

The CM shall duplicate the OFDMA symbols at the output of the IFFT as described in clause 7.4.16.1.6, adding a Cyclic Prefix to symbols 2n, and a Cyclic Suffix to symbols 2n+1, for n=0, 1, 2, …

The CM shall apply windowing as described in clause 7.4.16.1.6.

The CM shall add the Cyclic Prefix as described in clause 7.4.11, using the same CP value used for all other symbols.

The CM shall add a Cyclic Suffix as described in clause 7.4.16.1.6.

The CM shall use the Roll-off value specified in clause 7.4.11; the Roll-off value shall be the same as that for all other symbols except Initial Ranging Symbols.

NOTE: The Roll-off value used for fine ranging may be different from the corresponding value used for Initial Ranging.

The CM shall start to transmit the fine ranging signal one symbol (including the cyclic prefix) after the start time of the OFDMA frame.

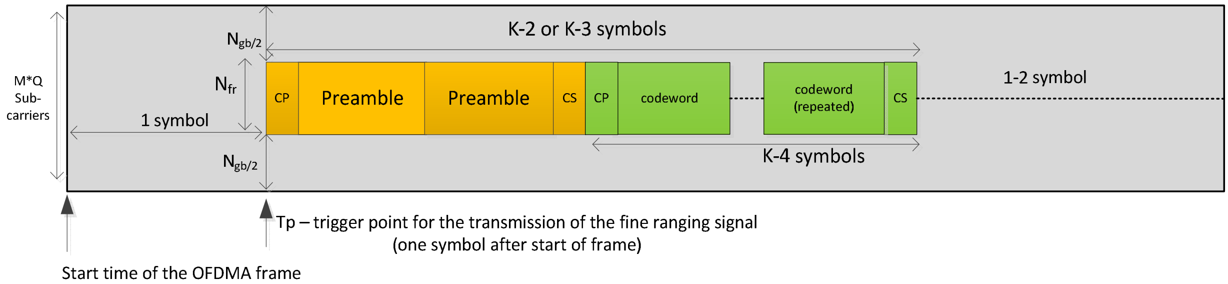

Figure 7.25: Fine Ranging Signal Transmission

The CM shall transmit the fine ranging signal with guardband of $N_{gb}/2$ subcarriers from each side of the allocated subcarriers. The CM calculates the number of subcarriers required for the guardband (N_{gb}) as follows:

$$N_{gb} = M*Q - N_{fr},$$

where:

M: is the number of minislots allocated for the fine ranging

N_{fr} : is the number of subcarriers as configured by the CMTS.

The CM shall transmit zero valued subcarriers in the guardband.

Figure 7.25 describes the fine ranging signal with M minislots of Q sub-carriers and K symbols and N_{gb} sub-carriers for the guardband.

A block diagram of the fine ranging signal processing in the transmitter is described in figure 7.26.

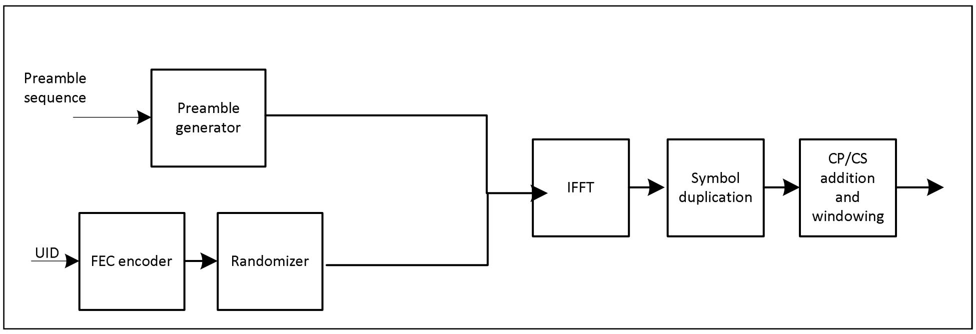

Figure 7.26: Fine Ranging Transmitter Processing

7.4.16.2.3 Fine Ranging FEC

The CM shall encode the 34 bytes of fine ranging information data using (362,272) shortening and puncturing LDPC encoder.

Shortening and puncturing encoder consists of three steps. In this step, the shortening step, one or more information bits are filled with 0 and the rest are filled with input bits. Then all information bits are encoded using the mother code matrix. After mother code encoding, the zero filled bits are deleted. The puncturing step is as described below.

The mother code is a rate 3/5 (480,288) binary LDPC code. A parity check matrix of the mother code is represented by table 7.16, where sub-matrix size (lifting factor) L = 48, see clause 7.4.4.2 for the compact definition of parity check matrix.

Table 7.16: (480, 288) LDPC Code Parity Check Matrix

16	1	28	9	40	38	16	-	-	-
28	42	36	11	39	9	8	38	-	-
5	2	18	16	25	47	-	2	19	-
18	18	40	18	0	34	-	-	7	32

Denote the information bits sent to the mother code encoder by $(a_0, \cdots, a_{287})$ and let the encoder output being $(a_0, \cdots, a_{287}, b_{288}, \cdots, b_{479})$, where $b_{288}, \cdots, b_{479}$ are parity-check bits. Then the shortening and puncturing steps can be described as follows:

The shortening step fills 0 to 16 consecutive bits starting at position 272, i.e. let $a_{272} = a_{273} = \cdots = a_{287} = 0$. The rest 272 bits i.e. $a_0, \cdots, a_{271}$, are fine ranging information data.

The bits to be deleted by the puncturing step are:

- Period 1: 54 consecutive bits $a_0, a_1, \cdots, a_{53}$
- Period 2: 48 consecutive bits $b_{432}, b_{433}, \cdots, b_{479}$

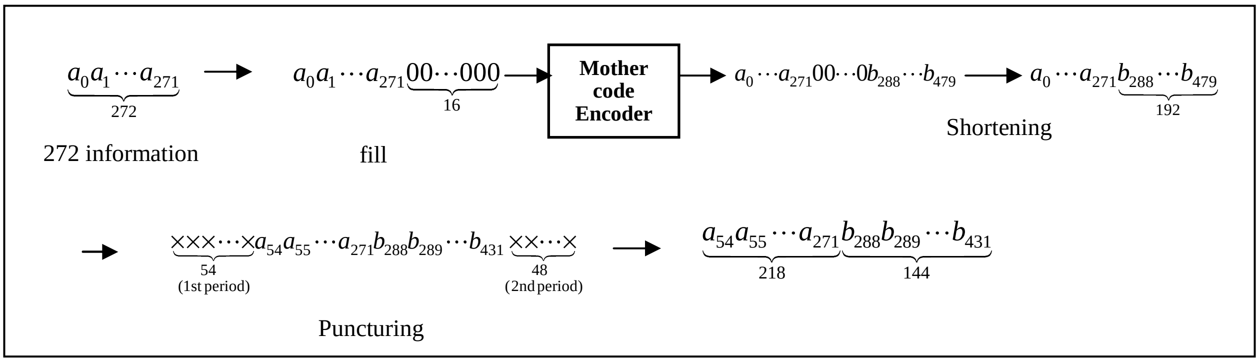

Figure 7.27: Shortening and Puncturing Encoder for the Fine Ranging FEC

7.4.16.2.4 Padding and Randomizing

The CM shall calculate the total number of data bits that can be transmitted in the fine ranging signal as follows:

Number_of_allocated_bits = N_{fr}* floor((K-4)/2)*2.

If the number of allocated bits is greater than 362, the CM shall pad the 362 bits output from the LDPC encoder with ones so that the encoded data and the pad bits equal the Number_of_allocated_bits.

The CM shall randomize the data and padding bits as described in clause 7.4.5.

The CM shall add the QPSK symbols to the data part of the fine ranging signal in the following order: The first N_{fr} QPSK symbols written to the N_{fr} subcarriers of the first symbol of the data part, the next N_{fr} QPSK symbols to the next data symbol, until all QPSK symbols are written vertically symbol by symbol. The first QPSK symbol is written to the lowest indexed subcarrier of a data symbol. Unfilled subcarriers in the last symbols are padded with 1s.

The CM shall transmit zero valued subcarriers in all symbol times not used for the preamble, data and pad bits.

NOTE: If K is an even number, the CM transmits K-2 symbols in the fine ranging signal (including the preamble), if K is an odd number, the CM transmits K-3 symbols (including the preamble).

7.4.16.2.5 Exclusion Bandwidths

Transmission of the fine ranging signal around exclusion bands is allowed, using same processing as explained in clause 7.4.16.2.2 with the same values of N_{fr} and N_{gb}.

Transmission with exclusion bands is illustrated in figure 7.28.

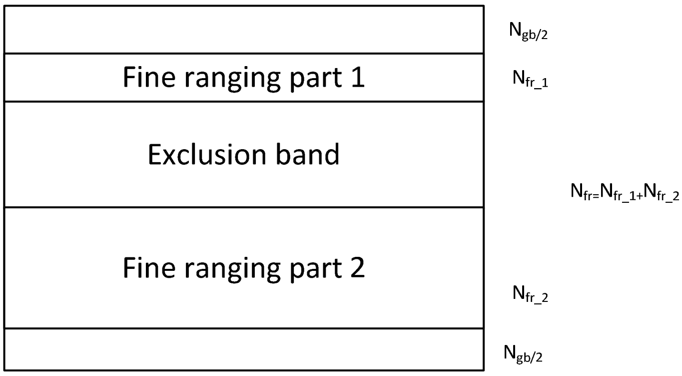

Figure 7.28: Fine Ranging and Exclusion Bands

When the fine ranging signal is transmitted around exclusion bands, the preamble sequence skips the exclusion band subcarriers. Figure 7.29 depicts an example of a fine ranging preamble and an exclusion band of K subcarriers.

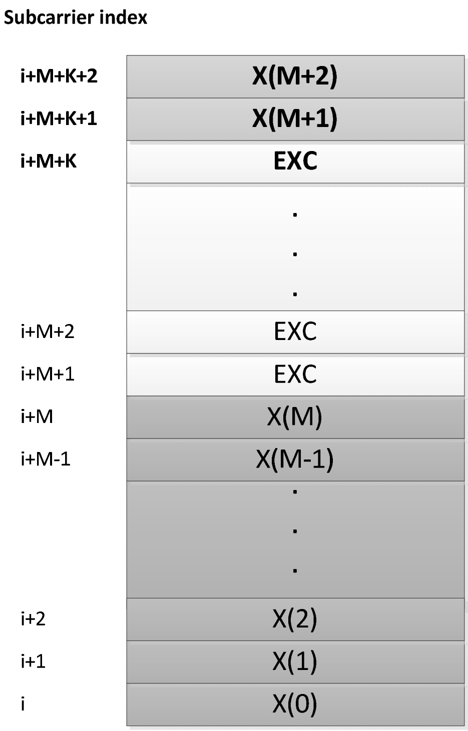

Figure 7.29: Fine Ranging Preamble and an Exclusion Band

7.4.16.2.6 Allowed Values and Ranges for Configuration Parameters

The CMTS shall configure the initial and fine ranging signal with the following limitations:

- The maximum number of subcarriers for initial ranging, not including the guardband, shall not exceed 64 subcarriers with 50 kHz subcarrier spacing
- The maximum number of subcarriers for initial ranging, not including the guardband, shall not exceed 128 subcarriers with 25 kHz subcarrier spacing.
- The maximum number of subcarriers for fine ranging, including subcarriers in the exclusion zones but excluding the guardband, shall not exceed 512 subcarriers with either 25 kHz or 50 kHz subcarrier spacing.
- Maximum number of subcarriers for fine ranging, excluding the subcarriers in the guardband and subcarriers in the exclusion bands, shall not exceed 256 subcarriers with 50 kHz subcarrier spacing.

- Maximum number of subcarriers for fine ranging, excluding the subcarriers in the guardband and subcarriers in the exclusion bands, shall not exceed 512 subcarriers with 25 kHz subcarrier spacing.
- Maximum preamble sequence size for initial and fine ranging shall not exceed 512 bits (64 Bytes) for 50 kHz and for the 25 kHz subcarrier spacing.
- Maximum number of preamble symbols (before duplication) for initial ranging shall not exceed 8.
- Maximum number of preamble symbols (before duplication) for fine ranging shall be 1.

7.4.16.2.7 Power and Time Adjustments

Algorithms for power and time adjustments (such as number of fine ranging trials, frequency allocations, etc.) are vendor-specific implementation.

7.4.16.3 Probing

7.4.16.3.0 Probing Usage

Probing is used during admission and steady state for pre-equalization configuration and periodic transmission power and time-shift ranging.

7.4.16.3.1 Probing Frame

A probing frame consists of K contiguous probing symbols (OFDM symbols), where K is the number of symbols in the minislot. The probing frame is aligned with the minislot boundaries in the time domain.

7.4.16.3.2 Probing Symbol Pilots

Probing symbol pilots are BPSK subcarriers, generated from the PRBS generation scheme described in clause 7.4.16.3.3.

The CM shall use the generation scheme detailed in clause 7.4.16.3.3 to generate 2 048/4 096 subcarriers for 2K/4K FFT.

The CM shall use the same BPSK modulation for a specific subcarrier in all probing symbols.

The CM shall transmit zero valued subcarriers in exclusion subcarriers.

Probing symbol pilot i is always associated with the i-th subcarrier number, where:

$$i = 0, 1, \ldots , 2\ 047 \text{ for 2K FFT}$$

and

$$i = 0, 1, \ldots , 4\ 095 \text{ for 4K FFT}$$

(Subcarriers are numbered in ascending order starting from 0.)

7.4.16.3.3 PRBS Generation Scheme

The polynomial definition for the PRBS scheme is $X^{12} + X^9 + X^{8+} X^{5+} 1$, where the seed is 3 071. The period of the PRBS is 2^{12}-1 bits, which is sufficient to create one probe symbol without repetitions. The sequence is illustrated in figure 7.30.

The CM's linear feedback shift register shall be clocked after every subcarrier starting at subcarrier 0, i.e. subcarrier with k=0 in the IDFT equation of clause 7.4.10.

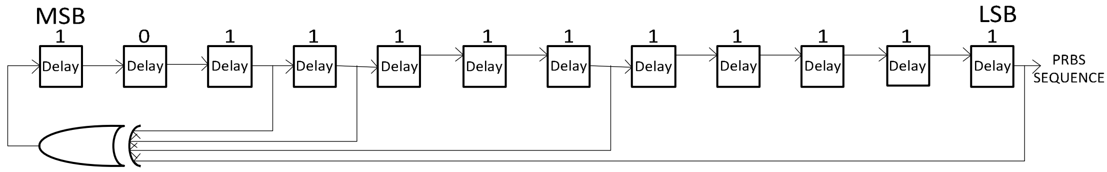

Figure 7.30: Polynomial Sequence for Pseudorandom Binary Sequence Generation

The PRBS sequence for 4K FFT is:

1 1 1 1 1 1 1 1 1 1 0 1 0 1 1 1 … 1 1 1 0 0 1 0 1 0 0 0 1 0 1

The PRBS sequence for 2K FFT is:

1 1 1 1 1 1 1 1 1 1 0 1 0 1 1 1 … 0 0 1 0 0 1 1 0 0 0 0 0 1 1

The PRBS sequence is mapped to the BPSK pilots as follows:

0 is mapped to a BPSK pilot of 1

1 is mapped to a BPSK pilot of -1

7.4.16.3.4 Probing Information

The CMTS shall allocate a specific probing symbol within the probing frame and instruct the CM to transmit the probing sequence in that symbol.

The CMTS shall specify the probing symbol within the probing frame through the parameter "Symbol in Frame".

The CMTS shall send three parameters to the CM: "st", "Start Subcarrier", and "Subcarrier Skipping".

The CM shall support staggering pattern [4] for probing, when the staggering bit "st" is set to one, when "st" is set to zero, the staggering is off.

The CMTS shall define a probing pattern consisting of either the pilots from all the subcarriers of the assigned probing symbol, or a set of pilots from scattered subcarriers of the assigned probing symbol. Please refer to clause 6.4.4 in [4] for detailed probe mapping.

The range of "start subcarrier" is from 0 to 7. The range of "subcarrier skipping" is from 0 to 7. Figure 7.31 and figure 7.32 illustrate the use of these parameters.

The CM shall use the *start subcarrier* and *subcarrier skipping* parameters to determine which subcarriers are to be used for probing transmission, as follows:

- The "start subcarrier" parameter is the starting subcarrier number.
- The "subcarrier skipping" parameter is the number of subcarriers to be skipped between successive pilots. "Subcarrier skipping" = 0 implies no skipping of subcarriers (i.e. all subcarriers in a single symbol belong to a single transmitter).

The CM shall not transmit the probing sequence using excluded subcarriers. Excluded subcarriers are those subcarriers in which no CM is allowed to transmit, generally because they are frequencies used by other systems (including guard-band subcarriers). The CM shall transmit the probing sequence using both used and unused subcarriers.

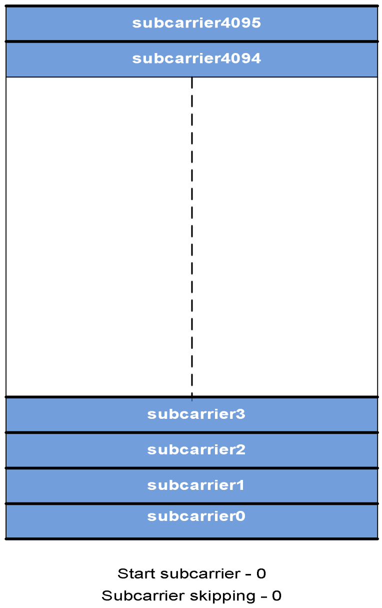

Figure 7.31: 4K FFT Example, All Subcarriers Used for Probing, no Skipping

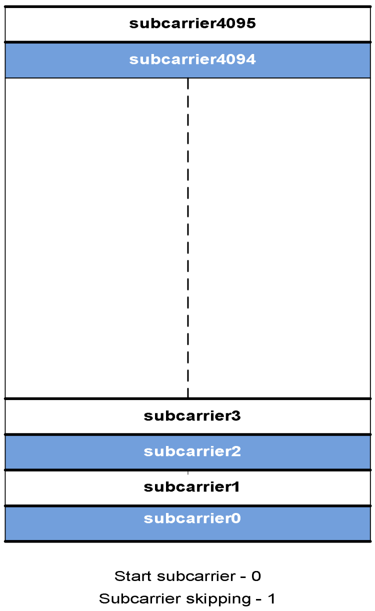

Figure 7.32: 4K FFT Example, Alternate Subcarriers Used for Probing

The CMTS shall not configure more than a single type of probe ("st", "Start Subcarrier", "Subcarrier Skipping" and PW value) on the same OFDMA frame per CM.

The CMTS shall have the ability to scale the transmission power per subcarrier by configuring the PW bit in the P-IE [4].

The CM shall scale its transmission power per subcarrier when transmitting the probe as required by the CMTS in the P-IE [4]. The range of the scaling values is Probedelta_n = -2 to -9 dB. See clause 7.4.13.3.

7.4.17 Upstream Pilot Structure

7.4.17.0 Edge and Body Minislots

Pilots are used by the CMTS receiver to adapt to channel conditions and frequency offset.

DOCSIS 3.1 specifies two minislot types, differing in the number of subcarriers per minislot, 8- and 16-subcarrier minislots.

Two types of minislots are defined for each minislot size: edge minislots and body minislots.

The CM shall use an edge minislot as the first minislot in a transmission burst.

The CM shall use body minislots for all other minislots in a transmission burst with the following two exceptions:

1) The CM shall use an edge minislot for the first minislot of an OFDMA frame that is not a zero valued minislot.

2) The CM shall use an edge minislot for the first minislot after an exclusion band or after one or more contiguous skipped subcarriers or after a zero valued minislot.

Figure 7.33 describes the usage of edge and body minislots. Note that TX-2 is a one minislot burst comprising of a single edge minislot.

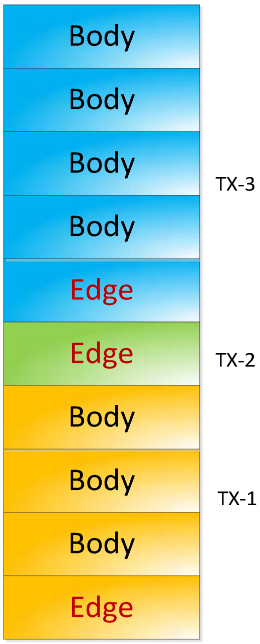

Figure 7.33: Edge and Body Minislots in a Transmission Burst

Pilots are subcarriers that do not carry data. Instead, a pilot subcarrier encodes a pre-defined BPSK symbol known to the receiver (see clause 7.4.16.3.3). DOCSIS 3.1 also specifies complementary pilots. Complementary pilots are subcarriers that carry data, but with a lower modulation order than other data subcarriers in the minislot. If the modulation order used for data in the minislot is *M*, the CM shall use complementary pilots with modulation order equal to the maximum between *M*-4 and 1 (BPSK). For example if the bit loading in a minislot is 12, Complementary Pilots use 8 bits. If the bit loading is 4, Complementary Pilots will use BPSK. The CMTS receiver may use complementary pilots to enhance its signal processing, such as to improve the accuracy of the centre frequency offset acquisition.

For each minislot size, seven pilot patterns are defined.

Pilot patterns differ by the number of pilots in a minislot, and by their arrangement within the minislot. The different pilot patterns enable the CMTS to optimize its performance (physical layer rate and pilot overhead) according to different loop conditions and variations of SNR with frequency. Each pilot pattern defines edge and body minislots.

Two additional pilot patterns are specified for subslots (see clauses 7.4.17.4 and 7.4.17.5); these are required for both the CM and the CMTS.

The following clauses describe the seven pilot patterns for each minislot size, and the pilot patterns for subslots.

7.4.17.1 Pilot Patterns for 8-Subcarrier Minislots

Figure 7.34 and figure 7.35 define the pilot patterns for edge and body minislots with 8 subcarriers.

The CM shall support pilot patterns 1-7.

The CMTS shall support pilot patterns 1-4.

The CMTS SHOULD support pilot patterns 5-7.

The CMTS shall use either pilots pattern 1-4 or pilot patterns 5-7 on the same OFDMA channel.

The CMTS shall not use a mixture of pilot patterns 1-4 and 5-7 on the same OFDMA channel.

In each figure, the horizontal axis represents OFDMA symbols, and the vertical axis represents the subcarriers. Each square in a figure represents a subcarrier at a specific symbol time. Pilots are designated by "P" and complementary pilots by "CP". All other subcarriers carry data with the modulation order of the minislot.

The figures show patterns for K between 6 and 16. For K>16 the complementary pilots are always located in the 14th and 16th symbols, all symbols from the 17th symbol to the end of the frame carry data only. Pilot locations are the same for any K.

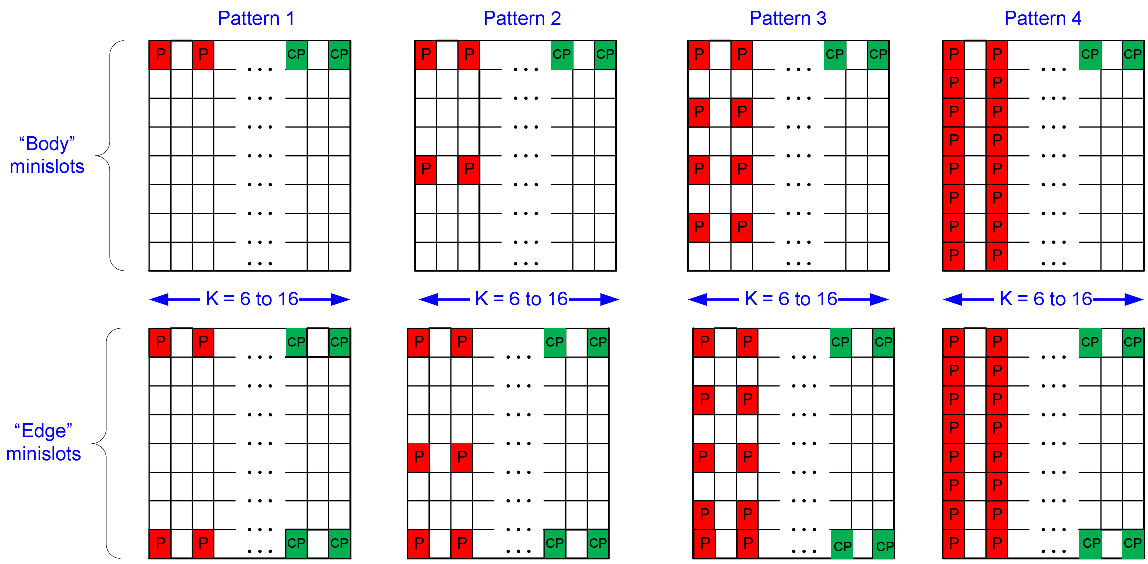

Figure 7.34: Pilot Patterns 1-4 for Minislots with 8 Subcarriers

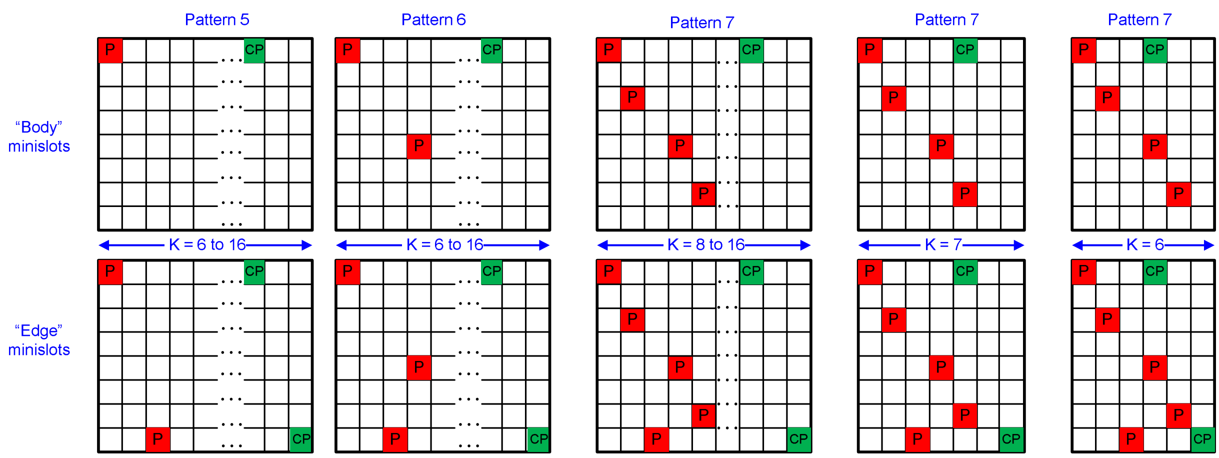

Figure 7.35: Pilot Patterns 5 - 7 for Minislots with 8 Subcarriers

7.4.17.2 Pilot patterns for 16-Subcarrier Minislots

Figure 7.36 and figure 7.37 define the pilot patterns for minislots with 16 subcarriers.

The CM shall support pilot patterns 8-14.

The CMTS shall support pilot patterns 8-11.

The CMTS should support pilot patterns 12-14.

The CMTS shall use either pilots pattern 8-11 or pilot patterns 12-14 on the same OFDMA channel.

The CMTS shall not use a mixture of pilot patterns 8-11 and 12-14 on the same OFDMA channel.

The CMTS shall configure minislots with 16 subcarriers to be used with 25 kHz subcarrier spacing with the exception for RFoG.

The figures show patterns for K between 6 and 9. For K > 9, the complementary pilots are always located in the 7th and 9th symbols, all symbols from the 10th symbol to end of frame carry data only. Pilot locations are the same for any K.

The horizontal axis in the figure represents OFDMA symbols, and the vertical axis represents subcarriers. Each square in a figure represents a subcarrier at a specific symbol time. Pilots are designated by "P" and complementary pilots by "CP". All other subcarriers carry data with the modulation order of the minislot.

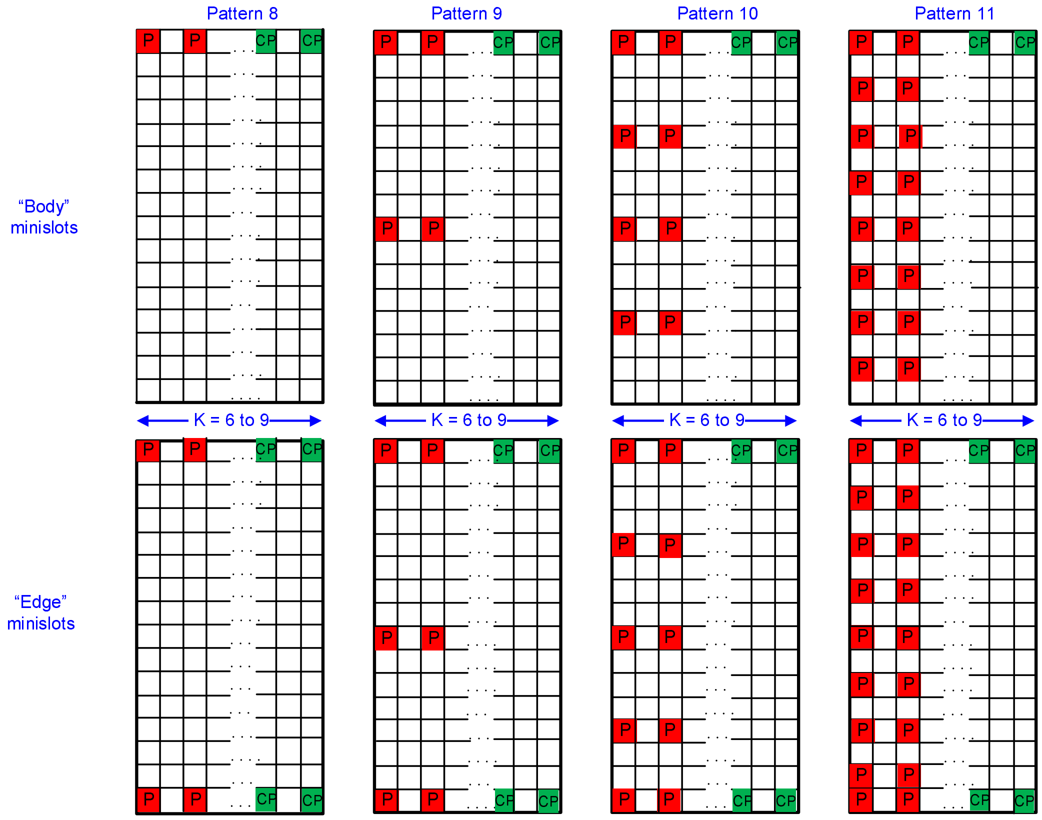

Figure 7.36: Pilot Patterns 8-11 for Minislots with 16 Subcarriers

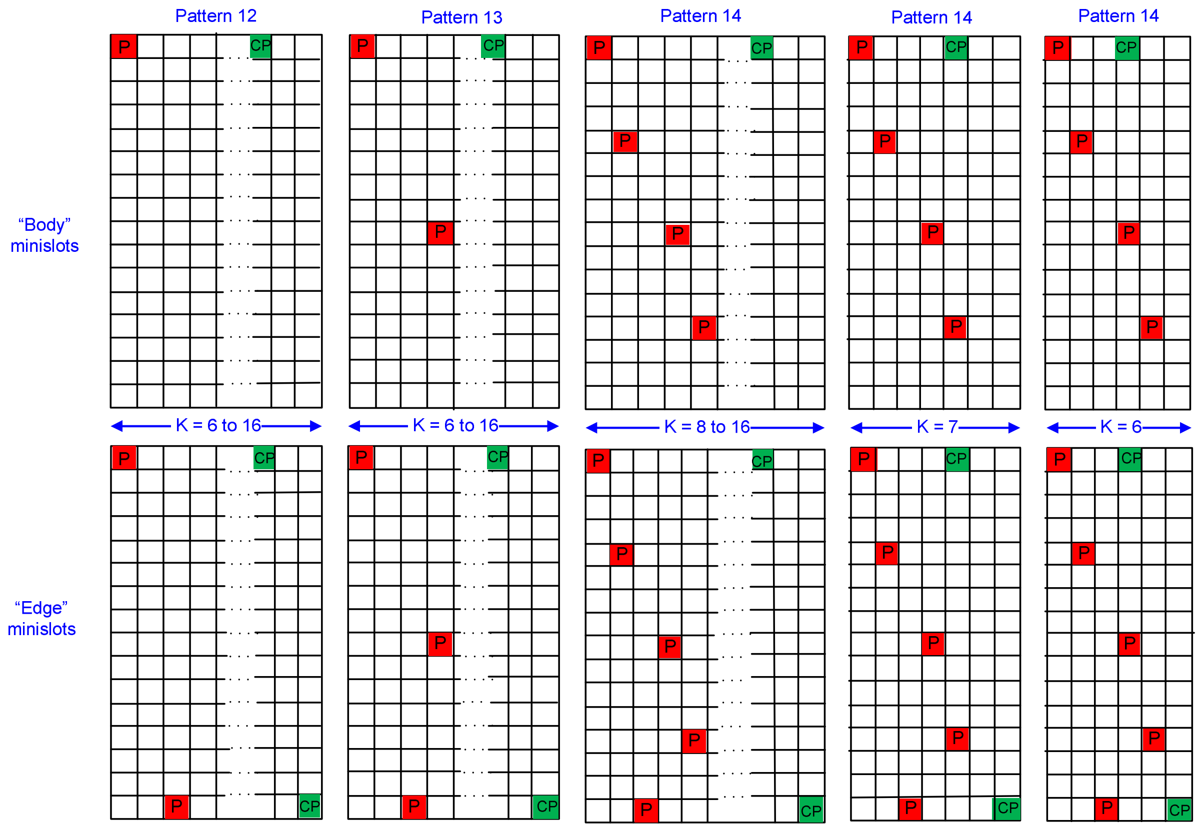

Figure 7.37: Pilot Patterns 12 - 14 for Minislots with 16 Subcarriers

7.4.17.3 Pilot Boosting

The CM shall use higher power (pilot boost) when transmitting pilots and complementary pilots with pilot patterns 5-7 and patterns 12-14, with the following exception:

The CM shall use boosted power for the pilot and normal power for the complementary pilot when both are used in the same symbol and in the same minislot.

The CM shall boost pilots and complementary pilots by a factor of 3 in power (about 4,7 dB).

7.4.17.4 Pilot Patterns for 8-Subcarrier Subslots

Subslots are used to carry REQ messages which are always 7 bytes or 56 bits long. Data subcarriers are always QPSK-modulated, and are not encoded by any FEC but are randomized using the randomizer described in clause 7.4.5.

Figure 7.38 depicts the pilot pattern for a subslot with 8 subcarriers.

The CM shall support the pilot pattern for 8-subcarrier subslots.

The CMTS shall support the pilot pattern for 8-subcarrier subslots.

Pilots are designated by "P", and no complementary pilots are used; all other subcarriers carry data with the modulation order of the subslot.

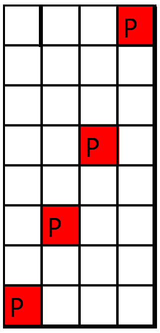

Figure 7.38: Pilot Pattern for Subslots with 8 Subcarriers

7.4.17.5 Pilot Patterns for 16-Subcarrier Subslots

Figure 7.39 depicts the pilot pattern for a subslot with 16 subcarriers.

The CM shall support the pilot pattern for 16-subcarrier subslots.

The CMTS shall support the pilot pattern for 16-subcarrier subslots.

Pilots are designated by "P", and no complementary pilots are used: all other subcarriers carry data with the modulation order of the subslot.

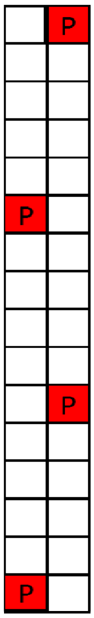

Figure 7.39: Pilot Pattern for Subslots with 16 Subcarriers

7.4.17.6 Pilot Modulation

The CM shall BPSK modulate the pilots using the PRBS defined in clause 7.4.16.3.3 using the feedback shift register illustrated in figure 7.30. This feedback shift register is initialized for the subcarrier with index k=0 of the IDFT equation of clause 7.4.10. It is then clocked once for every subcarrier of the IDFT. If the subcarrier happens to be a pilot this is BPSK modulated with the output of the feedback shift register, with a value of 0 mapping to (1 + j0) and a value of 1 mapping to (-1 + j0).

7.4.18 Upstream Pre-Equalization

A CM shall implement a linear pre-equalizer with a single complex coefficient per subcarrier.

The CMTS shall be able to direct a CM to pre-equalize its upstream transmission using CMTS-assigned pre-equalization coefficients as a step in the ranging process.

The CMTS uses the CM's probe signal for pre-equalizer coefficient updates. The probes are described in clause 7.4.16.3. The message used to send information required for updating the pre-equalizer coefficients is described in the Ranging Response (RNG-RSP) section of [4].

The CMTS may specify the subcarriers (i.e. frequency range) over which coefficient updates is to be performed.

The CMTS shall have the ability to scale the transmission power per subcarrier when configuring the probe transmission using the Range Response message.

The CM shall scale its transmission power per channel when transmitting the probe as required by the CMTS in the RNG-RSP message. The range of the scaling values is: 0 to [10log(skip+1)] dB. Skip is defined in clause 7.4.16.3.4.

The CM shall use a default value of 1+j0 for all pre-equalizer coefficients of the used and unused subcarriers. The CM shall set to zero all pre-equalizer coefficients that correspond to the excluded subcarriers.

The CM shall set the pre-equalizer coefficient to one for any subcarrier whose status is changed from excluded to non-excluded. At the next probe opportunity the CM shall use a pre-equalization coefficient of 1+j0 on the subcarriers whose status has changed.

The CM shall update the pre-equalizer coefficients according to the RNG_RSP message as described below.

The RNG-RSP MAC message carries the pre-equalization adjustment information.

The RNG_RSP message sent by the CMTS specifies whether the pre-equalization coefficients sent by the CMTS are for coefficient initialization or for coefficient adjustment. If coefficient initialization is specified, the CM shall replace the pre-equalizer coefficients with the coefficients sent by the CMTS. In the case of an adjustment, the CM shall multiply the coefficients values sent by the CMTS with the current pre-equalization coefficient values, to get the new coefficients, as follows:

$$Ck(i+1)=Ck(i) * Ak(i)$$

where:

> Ck(i) is the pre-equalizer coefficient of the k-th subcarrier, as used in the last probe transmission, Ck(i+1) is the updated pre-equalizer coefficient of the k-th subcarrier and Ak(i) is the update coefficient information received in the RNG-RSP as a response to the corresponding probe transmission. "*" indicates a complex multiplication.

The CMTS shall use complex numbers for the update coefficients values in the form of I+j*Q where I and Q are both using 16-bit fractional two's complement notation -"s1.14" (sign bit, integer bit and 14 fractional bits).

The CM shall normalize the new calculated coefficients as follows:

> mean (abs (Ck)^2) = 1 (mean value computed over all pre-equalizer coefficients corresponding to the used and unused subcarriers).

The CM shall pre-equalize all transmissions other than probe signals, as defined by the CMTS via the RNG_RSP message. The CM shall pre-equalize all probe transmissions unless the bit in the P-MAP message that defines the presence or absence of pre-equalization, is set to "equalizer disabled".

The CM shall be able to transmit a probe signal with or without pre-equalization (all coefficients are reset to 1+j*0) as instructed by the CMTS using the P-MAP message described in [4].

7.5 Downstream Transmit and Receive

7.5.1 Overview

This clause specifies the downstream electrical and signal processing requirements for the transmission of OFDM modulated RF signals from the CMTS to the CM.

7.5.2 Signal Processing

7.5.2.0 Downstream PHY Processing

Serial data signals received from the PHY-MAC Convergence Layer are received and processed by the PHY as illustrated in figure 7.40. This process yields OFDM symbols with 4 096 subcarriers for the 4K FFT mode and 8 192 subcarriers for the 8K FFT mode, with each symbol consisting of:

- Data subcarriers
- Scattered pilots
- Continuous pilots
- PLC subcarriers
- Excluded subcarriers that are set to zero

This clause briefly describes that process and provides links to the specific requirements for each process described in the present document.

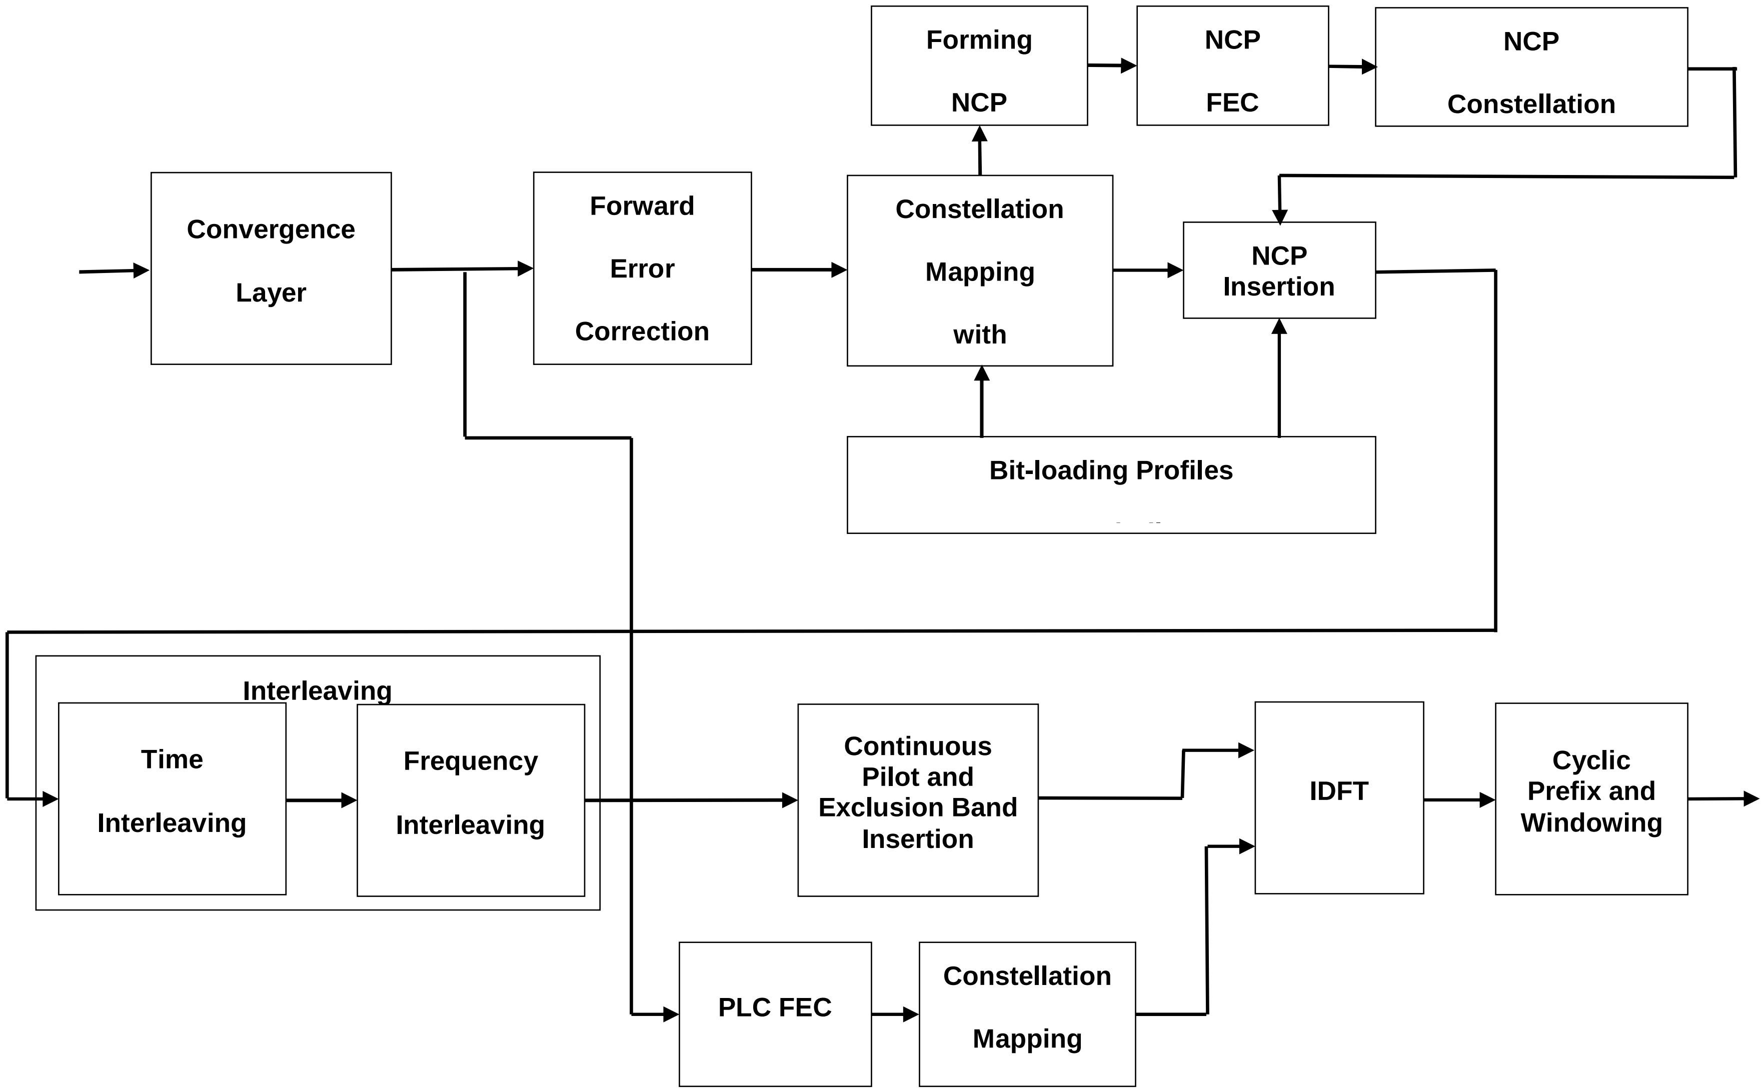

Figure 7.40: Downstream PHY Processing

7.5.2.1 Forward Error Correction (FEC) Encoding

The PHY begins processing incoming data by FEC encoding data bits to form encoded codewords. Forward error correction adds redundancy to the transmitted data; these redundant bits can be used by the receiver to detect and correct errors in the transmission. For DOCSIS 3.1, FEC encoding applies a concatenated BCH-LDPC encoder, based on [6], and then shuffling the bits in a codeword via bit interleaving. Downstream forward error correction is described in detail in clause 7.5.4.

7.5.2.2 Symbol Mapping to QAM Constellations

Once FEC encoded codewords have been created, the codewords are placed into OFDM symbols. Because each subcarrier in an OFDM symbol can have a different QAM modulation, the codewords are to be first demultiplexed into parallel cell words; these cell words are then mapped into constellations based on the corresponding bit loading pattern of the subcarrier's QAM constellation. In DOCSIS 3.1, QAM constellations extend from 16-QAM to 4096-QAM (excluding 32-QAM), with both square and non-square constellations. This process is described in clause 7.5.4.1.1.

7.5.2.3 Scattered Pilot Placeholder Insertion

OFDM transmission requires the insertion of scattered pilots to enable channel estimation and equalization in the receiver. While the insertion happens after time and frequency interleaving, since these pilots are not in the same spectral location in every symbol, insertion of these scattered pilots disrupt the spectral location of the QAM data subcarriers. To overcome this problem, place-holders for scattered pilot locations are inserted during the symbol mapping process.

7.5.2.4 Next Codeword Pointer Insertion

Detecting where the next codeword begins in an OFDM symbol can be difficult: more than one codeword may map into one OFDM symbol, the number of codewords per OFDM symbol may not be an integer, a codeword can overflow from one OFDM symbol to another, and the codeword could be shortened. Therefore, the transmitter is to convey to the receiver all of the locations where a new codeword begins within an OFDM symbol. These Next Codeword Pointers (NCPs) are encoded using another forward error correction method and are appended to OFDM symbols. The process of encoding and inserting the NCP for DOCSIS 3.1 is discussed in clauses 7.5.5.5.2 and 8.3.4.

7.5.2.5 Interleaving

These OFDM symbols, comprised of data subcarriers, scattered pilot placeholders, and NCPs, are then subjected to time and frequency interleaving. Time interleaving mitigates the impact of burst noise, while frequency interleaving mitigates the effect of ingress.

Time interleaving disperses the subcarriers of an input symbol over a set of output symbols, based on the depth of interleaving. Therefore, if an OFDM symbol is corrupted by a noise burst, this burst is spread over the symbols when it is de-interleaved, thereby reducing the error correction burden on the decoder. The time interleaving process is described in clause 7.5.6.

Frequency interleaving occurs after time interleaving. Frequency interleaving disperses subcarriers of the symbol along the frequency axis; therefore, OFDM subcarriers impacted by narrowband ingress are distributed between several codewords, reducing the number of errors in each codeword. The frequency interleaving process is described in clause 7.5.6.

7.5.2.6 Insertion of Pilots and Exclusion Sub-Bands

When interleaving is complete, placeholders for continuous pilots are inserted. These will be subject to modulation later, together with the placeholders already inserted for scattered pilots. Continuous pilots are pilots that occur at the same subcarrier location in every symbol. These are needed for receiver synchronization.

Exclusion bands and excluded subcarriers are inserted next. Nothing is transmitted at these subcarrier locations. The contiguous block of subcarriers allocated to the PHY Link Channel is also treated as an exclusion band at this stage; this is a placeholder for the PLC that is filled later. The regions outside the bandwidth of the OFDM signal may also be treated as exclusion bands.

Finally, the placeholders for scattered and continuous pilots are replaced with a BPSK pseudo-random sequence. This process is described in clause 7.5.15.

7.5.2.7 Encoding and Insertion of the PLC

The PLC is constructed within the convergence layer in parallel with the functions already discussed, relating to the main data channel. The PLC occupies the same contiguous set of subcarriers in every OFDM symbol. The PLC data is encoded for error correction, and then mapped into 16-QAM PLC data subcarriers. The PLC data is not subjected to the same time or frequency interleaving as the data; however they are block interleaved. The PLC is then inserted in place of its placeholder in each symbol. This process is described in clause 7.5.6.3.

7.5.2.8 IDFT Transformation and Cyclic Prefix Insertion

In this stage each OFDM symbol is transformed into the time domain using a 4 096-point or 8 192-point inverse discrete Fourier transform (IDFT). This 4 096 or 8 192 sample sequence is referred to below as the IDFT output. This process is described in clause 7.5.7.

7.5.2.9 Cyclic Prefix and Windowing

A segment at the end of the IDFT output is prepended to the IDFT, and this is referred to as the Cyclic Prefix (CP) of the OFDM symbol. There are five possible values for the length of the CP and the choice depends on the delay spread of the channel - a longer delay spread requires a longer cyclic prefix.

For windowing purposes another segment at the start of the IDFT output is appended to the end of the IDFT output - the roll-off period (RP). There are five possible values for the RP, and the choice depends on the bandwidth of the channel and the number of exclusion bands within the channel. A larger RP provides sharper edges in the spectrum of the OFDM signal; however, there is a time vs. frequency trade-off. Larger RP values reduce the efficiency of transmission in the time domain, but because the spectral edges are sharper, more useful subcarriers appear in the frequency domain. There is an optimum value for the RP that maximizes capacity for a given bandwidth and/or exclusion band scenario.

These topics are discussed in detail in clause 7.5.8.

7.5.3 Time and Frequency Synchronization

7.5.3.0 Timing and Frequency Synchronization Requirements

This clause specifies the timing and frequency synchronization requirements for DOCSIS 3.1 CMTS transmitters and CM receivers.

The purpose of this clause is to ensure that the CMTS transmitter can provide proper timing and frequency references for DOCSIS 3.1 downstream OFDM operation and that the CM receiver can acquire the system timing and subcarrier from the downstream for proper DOCSIS 3.1 operation.

The CMTS downstream OFDM symbol and subcarrier frequency and timing relationship is defined in clause 7.3.3.

7.5.3.1 Downstream sampling rate

The CMTS shall lock the 204,8 MHz Downstream OFDM Clock to the 10,24 MHz CMTS Master Clock (see table 7.1).

7.5.3.2 OFDM RF Transmission Synchronization

The CMTS shall lock the Downstream OFDM RF transmissions to the 10,24 MHz CMTS Master Clock (see table 7.1).

7.5.3.3 Downstream OFDM Symbol Clock Jitter

The CMTS shall adhere to the following clock jitter requirements for the downstream OFDM symbol clock over the specified frequency ranges:

- < [-21 + 20*log (f_{DS} /204,8)] dBc (i.e. < 0,07 ns RMS) 10 Hz to 100 Hz

- < [-21 + 20*log (f_{DS} /204,8)] dBc (i.e. < 0,07 ns RMS) 100 Hz to 1 kHz
- < [-21 + 20*log (f_{DS} /204,8)] dBc (i.e. < 0,07 ns RMS) 1 kHz to 10 kHz
- < [-4 + 20*log (f_{DS} /204,8)] dBc (i.e. < 0,5 ns RMS) 10 kHz to 100 kHz
- < [2 + 20*log (f_{DS} /204,8)] dBc (i.e. < 1 ns RMS) 100 kHz to (f_{DS} /2)

  where f_{DS} is the frequency of the measured downstream clock in MHz.

The CMTS shall use a value of f_{DS} that is an integral multiple or divisor of the downstream symbol clock. For example, an f_{DS} = 409,6 MHz clock may be measured if there is no explicit 204,8 MHz clock available.

In addition to meeting the clock jitter requirements given above, the CMTS is required to meet the phase noise specifications defined in table 7.36 of clause 7.5.9.1. In the event of a conflict between the clock jitter and the phase noise requirement, the CMTS shall meet the more stringent requirement.

7.5.3.4 Downstream Timing Acquisition Accuracy

The downstream clock timing is defined with respect to downstream PLC frame.

The CM shall be able to adjust its clock to synchronize its own clock timing with PLC frame for proper operation.

The CM shall be able to acquire downstream clock timing from downstream traffic (pilots, preambles, or mixed pilots, preambles, and data).

The CM shall have a timing acquisition accuracy better than 1 sample (4,8828125 ns).

7.5.3.5 Downstream Carrier Frequency Acquisition

The CM shall be able to acquire the carrier frequency from downstream (pilots, preambles, or mixed pilots, preambles and data).

7.5.3.6 Downstream Acquisition Time

The CM shall achieve downstream signal acquisition (frequency and time lock) in less than 60 s for a device with no previous network frequency plan knowledge.

Nonetheless, it is expected that the CM would be able to achieve downstream acquisition in less than 30 s.

7.5.4 Downstream Forward Error Correction

7.5.4.0 FEC Scheme

This clause describes the downstream forward error correction scheme used for DOCSIS 3.1. It is based on [6] clause 6.1, FEC Encoding; it is used here with the following modifications:

- A codeword will be the size of the short FEC Frame (16 200 bits); the "normal" FEC Frame (64 800 bits) is not used.
- Only the code rate 8/9 is used.
- Support for non-square constellations (128-QAM, 512-QAM, and 2048-QAM) is introduced.
- Support for mixed modulation codewords is introduced.
- Support for codeword shortening is introduced.
- Bit Interleaving for non-square constellations (128-QAM, 512-QAM, and 2048-QAM), mixed modulation mode constellations and for shortened codewords is introduced.

- Demultiplexing for non-square constellations (128-QAM, 512-QAM, and 2048-QAM), mixed modulation mode constellations, and for shortened codewords is introduced.
- Support for QPSK modulation is not required.

These changes are described in the following clauses.

7.5.4.1 Definitions

7.5.4.1.1 Mixed-Modulation Codewords

Before downstream FEC can be defined, it is important to understand what a mixed-modulation codeword is, as these codewords are handled differently. A mixed-modulation codeword belongs to a profile that does not use the same modulation constellation for all sub-carriers of the OFDM symbol. Note that subcarrier zero bit loading is not taken into account when determining if a codeword is a mixed-modulation codeword. In other words, if a profile has the same modulation constellation (i.e. same bit loading profile) for all non-zero bit-loaded subcarriers of the OFDM symbol, then the codewords of that profile are not considered to be mixed modulation.

As an example consider a profile in which even numbered subcarriers, excluding zero bit-loaded subcarriers (i.e. non-zero bit-loaded subcarriers 0, 2, 4, 6, etc.), are modulated with modulation A, and odd numbered subcarriers (i.e. non-zero bit-loaded subcarriers 1, 3, 5, 7, etc.) are modulated with modulation B. This provides a bits/s/Hz value that is the mean of the bits/s/Hz values of modulations A and B. Any codeword that belongs to that profile is a mixed-modulation codeword.

In this example, if these subcarriers are modulated as shown in the following table, the modulation combinations provide approximately 1,5 dB SNR granularity of additional spectral efficiency.

Table 7.17: Mixed Modulation with 1.5 dB SNR Granularity

Modulation A	Modulation B
128-QAM	256-QAM
256-QAM	512-QAM
512-QAM	1024-QAM
1024-QAM	2048-QAM
2048-QAM	4096-QAM

Another example of the use of mixed-modulation codewords can be applied to the case of an OFDM channel at the high frequency end with a significant tilt in SNR. In this case, a modulation profile for this part of the spectrum could use four different QAM constellations covering the OFDM symbol: 1024-QAM, 512-QAM, 256-QAM and 64-QAM. All of the codewords to which this profile is applied would be of the mixed-modulation type.

It is important to note that a codeword is treated as being a mixed-modulation type even if all of the subcarriers have the same modulation order; being of the mixed-modulation type is determined by the profile. For example, consider a codeword of the above profile in which all the subcarriers happen to be 256-QAM. Despite the fact that all subcarriers of this codeword have the same modulation, this codeword is treated as the mixed-modulation type since it belongs to a mixed-modulation profile. It is necessary to do this because the FEC encoder has no knowledge as to which subcarriers the codeword is going to be mapped while the encoding is being performed. Therefore, FEC encoder operations are determined by the profile applied to the codeword only.

As a final example, consider a profile that consists of 75 % 1024-QAM subcarriers and 25 % zero-bit-loaded subcarriers. In this case the codewords of that profile are not of the mixed-modulation type, since zero-bit-loaded subcarriers are ignored when determining mixed-modulation type.

7.5.4.1.2 Codeword vs. FECFrame

[6] uses the term FECFrame to refer to the bits of one LDPC encoding operation. In the present specification, the term codeword is used for the same concept.

7.5.4.2 FEC Encoding

7.5.4.2.0 FEC Encoding Requirements and Parameters

[6] clause 6.1, FEC Encoding, describes the FEC encoding requirements for the CMTS transmitter. The CMTS shall meet the portion of [6] clause 6.1, FEC Encoding, as described below:

The CMTS shall support the 8/9 code rate for the short codeword (N_{ldpc} = 16,200 bits) only. Support for other code rates and codeword sizes is not required.

The CMTS shall support the FEC coding parameters specified in table 7.18. This table is based on table 3(b), from [6].

Table 7.18: Coding Parameters (for Short Codewords N_{ldpc} = 16,200 and Code Rate 8/9)

LDPC Code Rate	BCH Uncoded Block Size K_{bch}	BCH Coded Block N_{bch}	LDPC Uncoded Block Size K_{ldpc}	LDPC Coded Block Size N_{ldpc}
8/9	14,232	14,400	14,400	16,200

7.5.4.2.1 Outer Encoding (BCH)

[6] clause 6.1.1, Outer Encoding (BCH), details the outer encoding requirements for normal and short codewords (FECFrames). For the CMTS, only short codewords are required. The CMTS shall meet the outer encoding requirements for short FECFrames specified in [6] clause 6.1.1, Outer Encoding (BCH).

7.5.4.2.2 Inner Coding (LDPC)

[6] clause 6.1.2, Inner Encoding, and clause 6.1.2.2, Inner Coding for Short FECFrame, detail the inner coding requirements for short codewords. For DOCSIS 3.1 codewords, the CMTS shall meet the inner coding requirements for short codewords and code rate 8/9 specified in [6] clause 6.1.2, Inner Encoding, and clause 6.1.2.2, Inner Coding for Short FECFrame.

7.5.4.2.3 Support for Codeword Shortening

7.5.4.2.3.0 Codeword Shortening for Insufficient Data

Codeword shortening is used for two purposes:

- Create shortened codewords when there is insufficient data to fill complete codewords.
- Achieve strong burst noise protection.

The full FEC block size for the FEC code rate of 8/9 is provided in table 7.18.

Codeword shortening is accomplished by shortening the uncoded block size in the *BCH Uncoded Block Size* column of table 7.18. Note that the number of parity bits remains the same; there is no shortening of the parity bits either in the BCH or in the LDPC.

When a shortened codeword is needed, the CMTS shall complete the codeword shortening process described here.

There are six overall steps to the codeword shortening process:

1) Prepending zero bits (BCH) to the data.
2) BCH encoding.
3) Removing the prepended zero bits.
4) Appending zero bits (LDPC) to the data.
5) LDPC encoding.
6) Removing the appended zero bits.

This is done in both the BCH encoder and the LDPC encoder, as shown in figure 7.41.

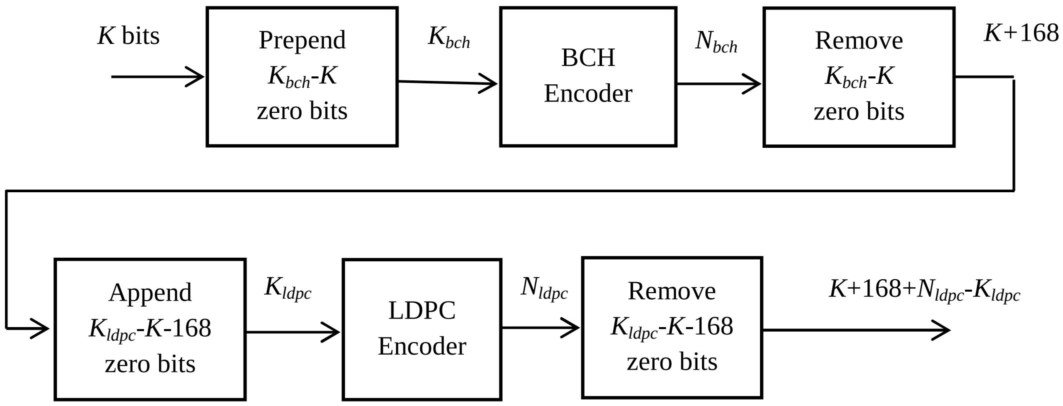

Figure 7.41: Codeword Shortening Process

The zero bit padding process is shown in more detail in figure 7.42.

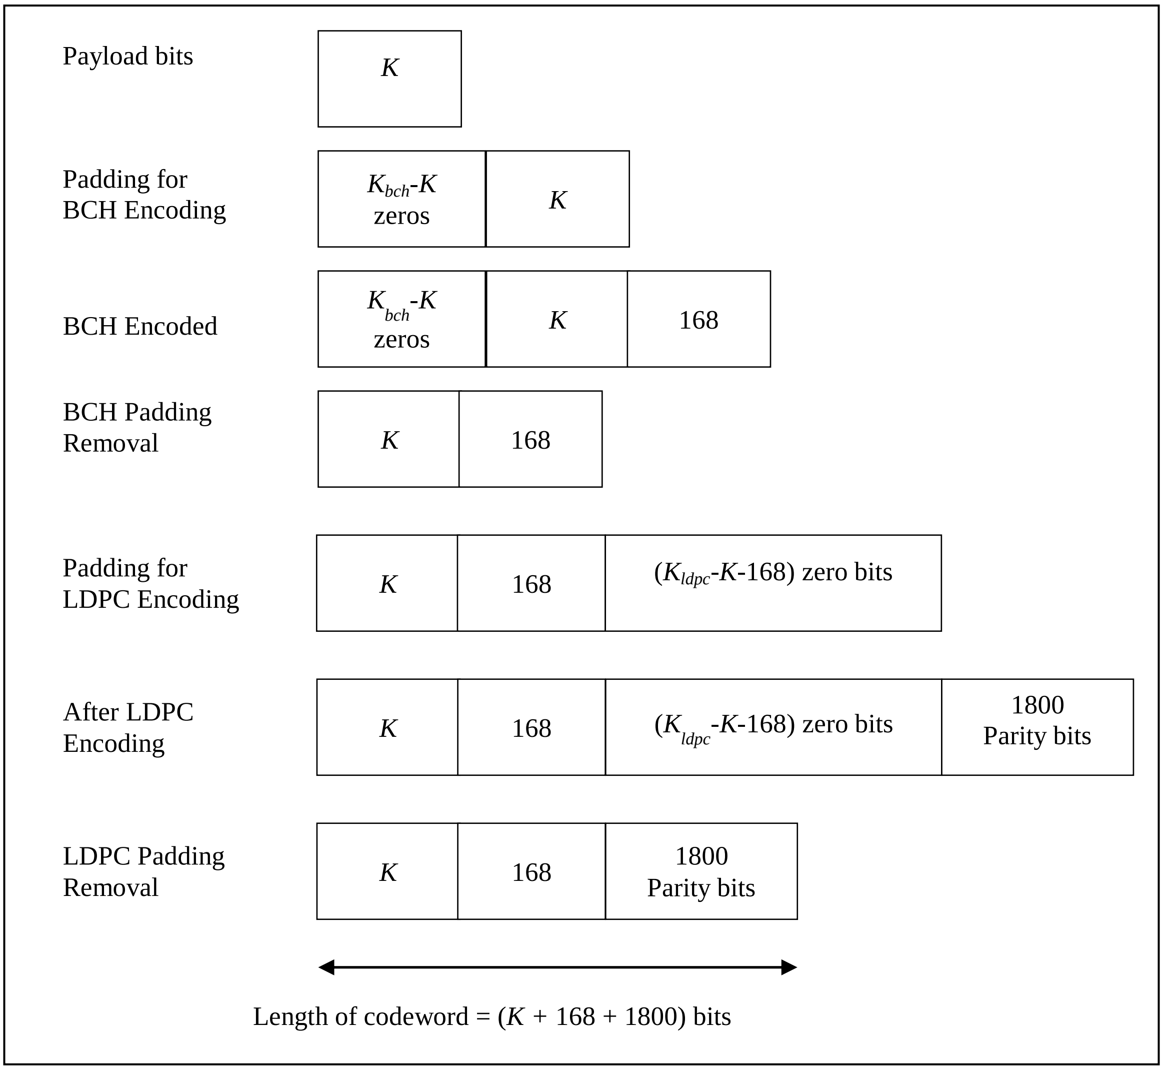

Figure 7.42: Padding Process

7.5.4.2.3.1 Codeword Shortening for Strong Burst Noise Protection

Although the primary purpose of codeword shortening is to support scenarios in which there is insufficient data to fill complete codewords, codeword shortening can also be used to provide signal protection in strong burst noise conditions. A lower code rate such as 7/9 has better burst noise capabilities than the 8/9 code rate. Through codeword shortening, it is possible to achieve the equivalent of a 7/9 code rate.

For example, the 16 200-bit block can be shortened by 8 096 bits. The number of parity bits remains unchanged at 1 800. Hence, this shortened codeword will have a block size of 8 104 with 6 304 information bits and 1 800 parity bits; this produces an effective code rate of approximately 7/9 (6 304/8 104 = 0,777887463). When the receiver receives this shortened codeword, it will pad the shortened 8 096 bits with zeros to create a 16 200-bit rate 8/9 codeword and decode it using the rate 8/9 decoder.

7.5.4.3 Bit Interleaving

7.5.4.3.1 Bit Interleaving for Non-Shortened Codewords

For non-shortened codewords that are not of the mixed-modulation type the CMTS shall apply parity interleaving, followed by column-twist interleaving as detailed in [6] clause 6.1.3, Bit Interleaver, with the number of rows, columns and column twisting parameters specified in this clause.

The number of rows and columns of the Bit Interleaver are specified by table 7.19.

Table 7.19: Bit Interleaver Structure

Modulation	Rows N_r	Columns N_c
16-QAM	2 025	8
64-QAM	1 350	12
128-QAM	2 315	7
256-QAM	2 025	8
512-QAM	1 800	9
1024-QAM	810	20
2048-QAM	1 473	11
4096-QAM	675	24
8192-QAM	1 247	13
16384-QAM	1 158	14

Since 16 200 is not divisible by 7, 11, 13 and 14, for 128-QAM, 2048-QAM, 8192-QAM and 16384-QAM constellations, the CMTS shall append zeros after parity interleaving and prior to column-twist interleaving at the end of the block: 5 zero bits for 128-QAM, 3 zero bits for 2048-QAM, 11 zero bits for 8192-QAM and 12 zero bits are added after the 16,200th bit. Thus, an extended block of 16 205 bits, 16 203 bits, 16 211 bits, and 16 212 bits will be interleaved by the column-twist interleaver for 128-QAM, 2048-QAM, 8192-QAM and 16384-QAM, respectively.

For non-shortened codewords that are not of the mixed-modulation type, the CMTS shall serially write the data bits into the column-twist interleaver column-wise, and serially read out row-wise, where the write start position of each column is twisted by t_c, as specified in tables 7.20 and 7.21.

Table 7.20: Column Twisting Parameter t_c (Columns 0 - 11)

Codeword Modulation Type	Columns N_c	Twisting Parameter t_c											
		Col. 0	1	2	3	4	5	6	7	8	9	10	11
16-QAM	8	0	0	0	1	7	20	20	21	-	-	-	-
64-QAM	12	0	0	0	2	2	2	3	3	3	6	7	7
128-QAM	7	0	1	2	2	2	3	3	-	-	-	-	-
256-QAM	8	0	0	0	1	7	20	20	21	-	-	-	-
512-QAM	9	0	1	2	3	5	6	7	9	11	-	-	-
1024-QAM	20	0	0	0	2	2	2	2	2	5	5	5	5
2048-QAM	11	0	0	0	0	0	3	3	4	4	4	4	-
4096-QAM	24	0	0	0	0	0	0	0	1	1	1	2	2
8192-QAM	13	0	0	0	0	0	0	2	3	5	8	8	9
16384-QAM	14	0	0	2	2	2	2	2	4	4	6	6	8

Table 7.21: Column Twisting Parameter t_c (Columns 12 - 23)

Codeword Modulation Type	Columns N_c	Twisting Parameter t_c											
		12	13	14	15	16	17	18	19	20	21	22	23
16-QAM	8	-	-	-	-	-	-	-	-	-	-	-	-
64-QAM	12	-	-	-	-	-	-	-	-	-	-	-	-
128-QAM	7	-	-	-	-	-	-	-	-	-	-	-	-
256-QAM	8	-	-	-	-	-	-	-	-	-	-	-	-
512-QAM	9	-	-	-	-	-	-	-	-	-	-	-	-
1024-QAM	20	5	7	7	7	7	8	8	10	-	-	-	-
2048-QAM	11	-	-	-	-	-	-	-	-	-	-	-	-
4096-QAM	24	2	3	7	9	9	9	10	10	10	10	10	11
8192-QAM	13	9	-	-	-	-	-	-	-	-	-	-	-
16384-QAM	14	9	9	-	-	-	-	-	-	-	-	-	-

7.5.4.3.2 Bit Interleaving for Non-Shortened Mixed Modulation Codewords

To support non-shortened mixed-modulation codewords, the bit interleaver specified in [6] has been modified. For non-shortened mixed-modulation codewords, the CMTS shall apply parity interleaving followed by column-twist interleaving as detailed in [6] clause 6.1.3, Bit Interleaver, and in the following discussion.

Because specific columns of the bit de-interleaver cannot be mapped to specific bits of the QAM constellation, column twisting interleaving is used over all 24 columns.

For non-shortened mixed-modulation codewords, the CMTS shall serially write the data bits into the column-twist interleaver column-wise, and serially read out row-wise, where the write start position of each column is twisted by t_c, as specified in tables 7.22 and 7.23.

Table 7.22: Column Twisting Parameter t_c (Columns 0 - 11)

Codeword Modulation Type	Columns N_c	Twisting Parameter t_c											
		Col. 0	1	2	3	4	5	6	7	8	9	10	11
Mixed-Modulation	24	0	0	0	0	0	0	0	1	1	1	2	2

Table 7.23: Column Twisting Parameter t_c (Columns 12 - 23)

Codeword Modulation Type	Columns N_c	Twisting Parameter t_c											
		Col. 12	13	14	15	16	17	18	19	20	21	22	23
Mixed-Modulation	24	2	3	7	9	9	9	10	10	10	10	10	11

7.5.4.3.3 Bit Interleaving for Shortened Codewords

Shortened codewords fall into one of the following three types:

- Square modulation
- Non-square modulation
- Mixed modulation

The CMTS shall interleave all types of shortened codewords as described in this clause.

The CMTS shall interleave the 1 800 parity bits as described in [6] clause 6.1.3, Bit Interleaver.

Because the shortened codeword can be quite small, it is possible that the entire codeword could map to one column of the interleaver and hence not get interleaved. To avoid this, the CMTS shall use the maximum number of columns in the bit interleaver (24) on shortened codewords. The number of rows that would be occupied by the shortened codeword is given by the following equation:

$$\boldsymbol{Row\_Count = ceil\left(\frac{K+1968}{24}\right)}$$

If (K+1 968) is not divisible by 24, then the last column will only be partially filled by the encoded shortened codeword, as illustrated in figure 7.43. For shortened codewords, the CMTS shall fill the unfilled part of the last column with bits that are labelled as "unused". For shortened codewords, the CMTS shall discard these "unused" bits in the memory read operation described below.

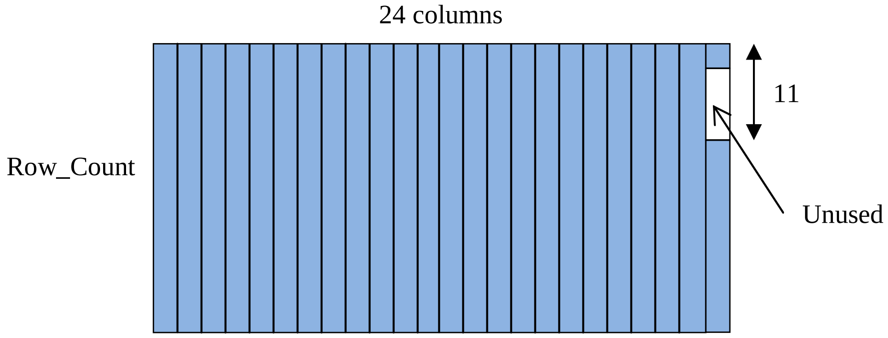

Figure 7.43: Bit De-interleaver Block for a Shortened Codeword

For shortened codewords, the CMTS shall serially write the data bits into the column-twist interleaver column-wise, and serially read out row-wise, where the write start position of each column is twisted by t_c, as specified in tables 7.24 and 7.25. For shortened codewords, the CMTS is to fill any unfilled bits of the last column with bits marked "unused". The CMTS is to write the first bit of the last column beginning from the 12th location since $t_c=11$ for column 24. If there are any bits left over after writing in the last location of the column, these bits are to be written beginning from the top of the column. For shortened codewords, the CMTS is to discard any bits labelled as "unused" during the process of reading along the rows of the two-dimensional array.

Table 7.24: Column Twisting Parameter t_c (Columns 0 - 11)

Codeword Type	Columns N_c	Twisting Parameter t_c											
		Col. 0	1	2	3	4	5	6	7	8	9	10	11
Shortened	24	0	0	0	0	0	0	0	1	1	1	2	2

Table 7.25: Column Twisting Parameter t_c (Columns 12 - 23)

Codeword Type	Columns N_c	Twisting Parameter t_c											
		Col. 12	13	14	15	16	17	18	19	20	21	22	23
Shortened	24	2	3	7	9	9	9	10	10	10	10	10	11

7.5.4.4 Downstream Receiver FEC Processing

Downstream data is encoded for FEC by the CMTS. The CM shall decode the FEC-applied codeword to correct for any bit errors introduced by noise and interference in the transmission medium. This process is discussed in this clause.

The FEC decoder at the CM operates on the QAM subcarriers of OFDM symbols to generate an error corrected bit stream. In addition, the decoder generates error statistics such as codeword error ratios.

The FEC decoding process is shown in figure 7.44.

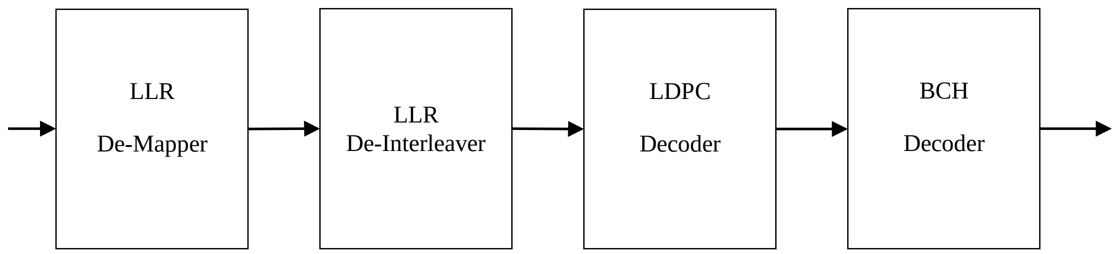

Figure 7.44: FEC Decoding Process

The receiver FEC decoder consists of the following components:

- The log-likelihood ratio (LLR) de-mapper processes one OFDM subcarrier at a time from the OFDM symbol and generates the LLRs for all bits of the QAM constellation, as defined by the bit-loading profile for the specific subcarrier. For example, if the subcarrier is 1024-QAM, the LLR de-mapper will generate 10 LLRs for the subcarrier and the values of the LLRs are implementation specific. The LLR de-interleaver operates on the LLRs. This is the inverse of the bit-interleaver that has been applied by the CMTS transmitter, described in clause 7.5.2.5. Note that the receiver operates on LLRs and not on bits.

- The LDPC decoder decodes the 16 200-bit LDPC codeword (or a shortened codeword). LDPC decoding is implemented using an iterative algorithm that uses message passing between the bit nodes and the check nodes of the Tanner graph of the LDPC code. If the CMTS has transmitted a shortened codeword (e.g. when the payload is not large enough to fill a complete codeword), the receiver augments the shortened codeword to full size with LLRs corresponding to zero-valued bits. The receiver then decodes the codeword using the 16 200-bit LDPC decoder, discarding the augmented bits.

- The BCH decoder generates an error corrected bit stream. The BCH decoder is also required to operate on shortened codewords.

- An error monitor determines codeword error ratios for reporting and troubleshooting.

7.5.5 Mapping Bits to QAM Constellations

7.5.5.0 Bits to QAM Mapping Methods

This clause describes the method used in DOCSIS 3.1 to map bits onto QAM constellations. It is based on [6] clause 6.2, Mapping Bits onto Constellations, and is used here with the following modifications:

- Parameters for mapping bits onto non-square constellations have been added
- Parameters for mapping bits of shortened codewords onto all constellation types have been added
- Parameters for mapping bits of mixed-modulation codewords onto all constellation types have been added

As described in [6] clause 6.2, Mapping Bits onto Constellations, the CMTS shall map each codeword to a sequence of QAM constellation values by:

- Demultiplexing the input bits into parallel cell words
- Mapping these cell words into constellation values

The mapping of bits to QAM constellation is carried out using the three sequential operations depicted in figure 7.45.

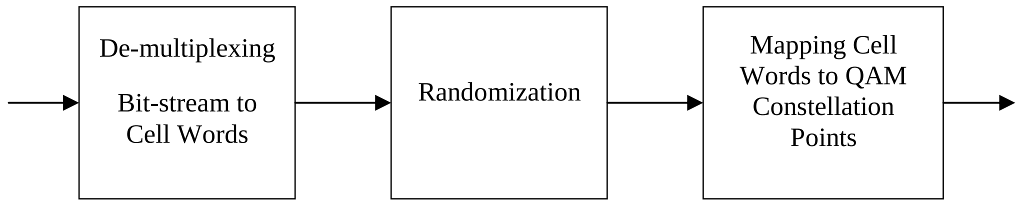

Figure 7.45: Bits to QAM Constellation Mapping

The CMTS shall use the number of bits per cell η_{MOD}, as defined in table 7.26, when bit mapping codewords to constellations.

For non-shortened codewords that are not of the mixed-modulation type, the CMTS shall use the Number of Output Data Cells defined in table 7.26 when bit mapping codewords to constellations.

Table 7.26: Parameters for Bit-Mapping onto Constellations

Modulation Mode	η_{MOD}	Number of Output Data Cells
16384-QAM	14	1 158
8192-QAM	13	1 247
4096-QAM	12	1 350
2048-QAM	11	1 473
1024-QAM	10	1 620
512-QAM	9	1 800
256-QAM	8	2 025
128-QAM	7	2 315
64-QAM	6	2 700
16-QAM	4	4 050

For the cases of mixed-modulation codewords and shortened codewords, the number of output symbols per LDPC block remains an integer. For both mixed-modulation codewords and shortened codewords, the CMTS shall pad the end of the LDPC block with zero bits to produce an integer number of bits in the final QAM symbol. The CM shall discard zero pad bits in the received symbol. This is described in further detail in the following clauses.

7.5.5.1 Modulation Formats

The CMTS modulator shall support 16-QAM, 64-QAM, 128-QAM, 256-QAM, 512-QAM, 1024-QAM, 2048-QAM, and 4096-QAM.

The CMTS modulator may support 8192-QAM and 16384-QAM.

The CM demodulator shall support 16-QAM, 64-QAM, 128-QAM, 256-QAM, 512-QAM, 1024-QAM, 2048-QAM, and 4096-QAM.

The CM demodulator may support 8192-QAM and 16384-QAM.

7.5.5.2 Bit-to-Cell Word Demultiplexer

7.5.5.2.1 Non-shortened Codewords

For non-shortened codewords that are not of the mixed-modulation type, the CMTS shall demultiplex the bit stream v_i from the bit interleaver into $N_{substreams}$ sub-streams, using the value of $N_{substreams}$ as defined in table 7.27 and the description following that table.

Table 7.27: Number of Sub-Streams in Demultiplexer

Modulation Mode	Number of Sub-Streams, $N_{substreams}$
16-QAM	8
64-QAM	12
128-QAM	7
256-QAM	8
512-QAM	9
1024-QAM	20
2048-QAM	11
4096-QAM	24
8192-QAM	13
16384-QAM	14

Bit-to-cell word demultiplexing is illustrated in figure 7.46.

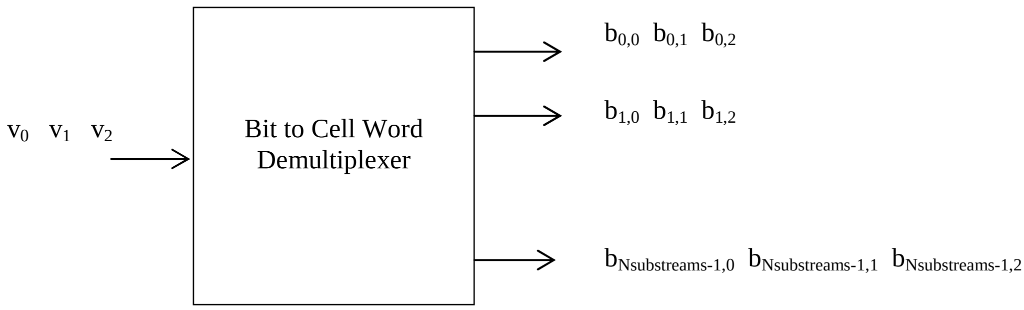

Figure 7.46: Bit-to-Cell Word Demultiplexer

For 16-QAM, 64-QAM, 256-QAM, 1024-QAM and 4096-QAM bit-to-cell word demultiplexing has to be carried out as described in [6] clause 6.2.1, Bit to Cell Word Demultiplexer.
Bit-to-cell word demultiplexing is defined as a mapping of the bit-interleaved input bits, v_{di}, onto the output bits $b_{e,do}$, where:

- v_{di} is the input to the demultiplexer;
- di is the input bit number;

- e is the demultiplexer sub-stream index ($0 \leq e < N_{substreams}$), which depends on ($di$ modulo $N_{substreams}$), as defined in table 7.28 through table 7.32;
- $do = floor\left(\frac{di}{N_{substreams}}\right)$ is the output cell number from the demultiplexer;
- $b_{e,do}$ is the output from the demultiplexer.

Table 7.28: Parameters for Demultiplexing of Bits to Sub-Streams for 8/9 Code Rate with 128-QAM

Input bit number, *di* mod $N_{substreams}$	0	1	2	3	4	5	6
Output bit number, *e*	6	5	4	1	2	3	0

Table 7.29: Parameters for Demultiplexing of Bits to Sub-Streams for 8/9 Code Rate with 512-QAM

Input bit number, *di* mod $N_{substreams}$	0	1	2	3	4	5	6	7	8
Output bit number, *e*	8	7	6	1	2	3	4	5	0

Table 7.30: Parameters for Demultiplexing of Bits to Sub-Streams for 8/9 Code Rate with 2048-QAM

Input bit number, *di* mod $N_{substreams}$	0	1	2	3	4	5	6	7	8	9	10
Output bit number, *e*	10	9	8	7	2	3	4	5	6	1	0

Table 7.31: Parameters for Demultiplexing of Bits to Sub-Streams for 8/9 Code Rate with 8192-QAM

Input bit number, *di* mod $N_{substreams}$	0	1	2	3	4	5	6	7	8	9	10	11	12
Output bit number, *e*	12	11	2	3	4	5	6	7	8	9	10	1	0

Table 7.32: Parameters for Demultiplexing of Bits to Sub-Streams for 8/9 Code Rate with 16384-QAM

Input bit number, *di* mod $N_{substreams}$	0	1	2	3	4	5	6	7	8	9	10	11	12	13
Output bit number, *e*	13	12	2	3	4	5	6	7	8	9	10	11	1	0

For example, in the case of 128-QAM there will be 7 substreams at the output of the bit-to-cell word demultiplexer. The first 7 bits at the input to the demultiplexer are sent to sub-streams 6, 5, 4, 1, 2, 3 and 0, in that order. The next 7 input bits are also mapped in that order. The cell words are defined from the demultiplexer output as:

$$\left[y_{0,do} \cdot\cdot y_{\eta_{mod}-1,do}\right] = \left[b_{0,do} \cdot\cdot b_{N_{substreams}-1,do}\right]$$

Note that the non-shortened LDPC codeword size is not divisible by 7. However, with reference to the section on bit interleaving, it is seen that for 128-QAM the size of the non-shortened codeword has been extended to become a multiple of 7 through zero-padding. The same comments are applicable to non-shortened 2048-QAM, 8192-QAM and 16384-QAM codewords.

7.5.5.2.2 Shortened Codewords and Mixed-Modulation Codewords

It is important to emphasize that shortened codewords can have square modulation, non-square modulation or may be of the mixed-modulation type. The CMTS shall bypass the bit-to-cell demultiplexer and apply the bit-to-cell mapping described in this clause for all types of shortened codewords, as well as for non-shortened mixed modulation codewords.

When the bit-to-cell word demultiplexer is bypassed, the bit-to-cell mapping becomes:

Cell 0: $\left[y_{0,0}..y_{\eta_{mod0}-1,0}\right] = \left[v_0 \ \ldots \ v_{\eta_{mod0}-1}\right]$

Cell 1: $\left[y_{0,1}..y_{\eta_{mod1}-1,1}\right] = \left[v_{\eta_{mod0}} \ \ldots \ v_{\eta_{mod0}+\eta_{mod1}-1}\right]$

etc.

The modulation assigned to cells 0 and 1 in the previous equations correspond to the η_{mod} values given by table 7.26. The first cell has the modulation corresponding to η_{mod0} and the second cell has the modulation corresponding to η_{mod1}. This modulation is defined by the bit loading pattern assigned to the profile to which this codeword belongs.

This mapping is simply a case of partitioning the interleaved bit stream to blocks of bits of size η_{mod0}, η_{mod1}, … , $\eta_{modLAST}$, where the sequence $\{\eta_{mod0}, \eta_{mod1}, \ldots, \eta_{modLAST}\}$ is given by the bit loading pattern of the profile to which this codeword belongs.

Let $\eta_{modLAST}$ correspond to the bit loading of the last cell of the sequence. It is possible that the shortened and/or mixed-modulation codeword at the output of the bit interleaver might not have sufficient bits to complete this cell. In this case zero-padding of the input bit stream has to be used for cell completion.

7.5.5.3 Randomization

The CMTS shall randomize cell words of data subcarriers, NCP subcarriers and PLC subcarriers, just before mapping these onto QAM constellations, as described in this clause.

The CMTS shall also introduce BPSK-modulated subcarriers for the following subcarriers during the randomization process, as described in this clause.

a) Zero-bit-loaded subcarriers of the codewords of individual profiles

b) Zero-bit-loaded subcarriers in the NCP segment

c) Zero-bit-loaded subcarriers that are introduced to complete the symbol

NCP and zero bit-loading are described in clause 7.5.5.5.

The wordlength (η_{MOD}) of a cell word ranges from 4 bits for 16-QAM to 14 bits for 16384-QAM.

For 16-QAM to 4096-QAM the CMTS shall randomize each cell word through a bit-wise exclusive-OR operation with the η_{MOD} least significant bits (LSBs) of the 12-bit register D0 of the linear feedback shift register (LFSR) shown in figure 7.47.

$$\left(z_0 \,..\, z_{\eta_{MOD}-1}\right) = \left(y_0..y_{\eta_{MOD}-1}\right) \mathit{bitwiseXOR} \ \ (D_0[0]..D_0[\eta_{MOD}-1])$$

For 8192-QAM the CMTS shall randomize the 13 bits of the cell word through a bit-wise exclusive-OR operation with the 12 bits of register D0 and the LSB of register D1 of figure 7.47, as given below:

$$(z_0 \,..\, z_{12}) = (y_0..y_{12}) \ \mathit{bitwiseXOR} \ \ (D_0[0]..D_0[\eta_{MOD}-1] \ D_1[0])$$

For 16384-QAM the CMTS shall randomize the 14 bits of the cell word through a bit-wise exclusive-OR operation with the 12 bits of register D0 and the 2 LSBs of register D1 of figure 7.47, as given below:

$$(z_0 \,..\, z_{13}) = (y_0..y_{13}) \ \mathit{bitwiseXOR} \ \ (D_0[0]..D_0[\eta_{MOD}-1] \ D_1[0] \ D_1[1])$$

NCP subcarrier cell words are 2-bit for QPSK, 4-bit for 16-QAM or 6-bit for 64-QAM. The CMTS shall randomize these through bit-wise exclusive-OR operation with the 2, 4 or 6 LSBs of the 12-bit register D0.

The CMTS shall set the zero-bit-loaded subcarriers in the data segment and NCP segment to the BPSK modulation given by LSB of register D0.

$$z_0 = D_0[0]$$

The CMTS shall clock the LFSR once, after each of the previous operations.

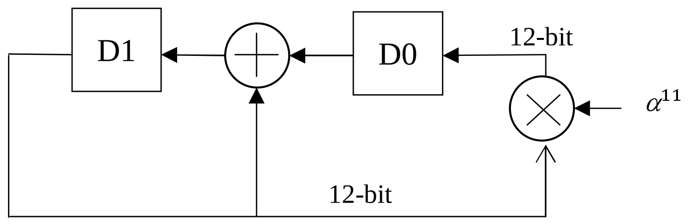

Figure 7.47: Linear Feedback Shift Register for Randomization Sequence

The LFSR is defined by the following polynomial in $GF[2^{12}]$.

$$x^2 + x + \alpha^{11}$$

The $GF[2^{12}]$ is defined through polynomial algebra modulo the polynomial:

$$\alpha^{12} + \alpha^6 + \alpha^4 + \alpha + 1$$

Each 12-bit $GF[2^{12}]$ element is a polynomial of α with a maximum degree of 11. The coefficient of α^0 is referred to as the LSB and the coefficient of α^{11} is referred to as the MSB.

This LFSR is initialized to the hexadecimal numbers given below:

D0 = "555"

D1 = "AAA"

This initialization is carried out at the beginning of an OFDM symbol, synchronized to the preamble of the PLC. Since the PLC subcarriers are inserted after time and frequency interleaving and data subcarriers are randomized before time and frequency interleaving, the following explanation is provided about how randomization is synchronized to the PLC.

Note that the first subcarrier of an OFDM symbol passes through the time interleaver arm with zero delay. Therefore the LFSR is initialized when this subcarrier is part of the OFDM symbol following the last OFDM symbol carrying the PLC preamble. Hence LFSR is initialized once for every 128 OFDM symbols.

The first subcarrier referred to previously can be a data subcarrier or a scattered pilot placeholder because both of these are time interleaved. If it is a data subcarrier then the cell word of that data subcarrier is randomized with the initialized values of D0 and D1, namely hexadecimal "555" and "AAA". After that the LFSR is clocked once. If the first subcarrier mentioned previously is a scattered pilot placeholder the LFSR is initialized but it is not clocked. This is because the LFSR is clocked only after each data or NCP subcarrier (including zero-bit-loaded subcarriers).

7.5.5.4 Cell Word Mapping into I/Q Constellations

The CMTS shall modulate each randomized cell word $(z_0 .. z_{\eta_{MOD}-1})$ from the randomizer described in clause 7.5.5.3 using a BPSK, QPSK, 16-QAM, 64-QAM, 128-QAM, 256-QAM, 512-QAM, 1024-QAM, 2048-QAM, 4096-QAM, 8192-QAM or 16384-QAM constellation as described in Annex A.

7.5.5.5 Transmitter Bit Loading for Symbol Mapping

7.5.5.5.0 Overview of Transmitter Bit Loading for Symbol Mapping

All subcarriers of an OFDM symbol may not have the same constellation; the constellation for each subcarrier is given in a table that details the bit loading pattern. This bit-loading pattern may change from profile to profile. This clause describes how the bits to symbol mapping is performed, with reference to a bit-loading pattern, in the presence of interleaving, continuous pilots, scattered pilots and excluded subcarriers.

Excluded subcarriers are subcarriers that are forced to zero-valued modulation at the transmitter. Subcarriers are excluded to prevent interference to other transmissions that occupy the same spectrum as the DOCSIS 3.1 OFDM transmission, for example, to accommodate legacy channels. Subcarriers are also excluded outside of the active OFDM bandwidth.

Excluded subcarriers are common to all profiles. The non-excluded subcarriers are referred to as active subcarriers. Active subcarriers are never zero-valued. The notation $S^{(E)}$ is used here to define the set of excluded subcarriers. This set will never be empty because there are always excluded subcarriers at the edges of the OFDM channel.

Continuous pilots are pilots that occur at the same frequency location in every OFDM symbol. The notation $S^{(C)}$ is used here to define the set of continuous pilots.

The PLC resides in a contiguous set of subcarriers in the OFDM channel. The CMTS adds the PLC to the OFDM channel after time and frequency interleaving; the CM extracts the PLC subcarriers before frequency and time de-interleaving. These subcarriers occupy the same spectral locations in every symbol. The notation $S^{(P)}$ is used here to define the set of PLC subcarriers.

For bit loading, continuous pilots and the PLC are treated in the same manner as excluded subcarriers; hence, the set of subcarriers that includes the PLC, continuous pilots and excluded subcarriers is defined as:

$$S^{(PCE)} = S^{(P)} \cup S^{(C)} \cup S^{(E)}$$

The subcarriers in the set $S^{(PCE)}$ do not carry data (PLC carry signalling information). The other subcarriers that do not carry data are the scattered pilots. However, scattered pilots are not included in the set $S^{(PCE)}$ because they do not occupy the same spectral locations in every OFDM symbol.

The modulation order of the data subcarriers is defined using bit-loading profiles. These profiles include the option for zero bit-loading. Such subcarriers are referred to as zero-bit-loaded subcarriers and are BPSK modulated using the randomizer LSB, as described in clause 7.5.5.3.

All active subcarriers with the exception of pilots are transmitted with the same average power. Pilots are transmitted boosted by a factor of 2 in amplitude (approximately 6 dB).

Scattered pilots do not occur at the same frequency in every symbol; in some cases scattered pilots will overlap with continuous pilots. If a scattered pilot overlaps with a continuous pilot, then that pilot is no longer considered to be a scattered pilot. It is treated as a continuous pilot.

Because the locations of scattered pilots change from one OFDM symbol to another, the number of overlapping continuous and scattered pilots changes from symbol to symbol. Since overlapping pilots are treated as continuous pilots, the number of scattered pilots changes from symbol to symbol.

The following notation is used here:

N: The total number of subcarriers in the OFDM symbol, equalling either 4 096 or 8 192

N_C: The number of continuous pilots in an OFDM symbol

N_S: The number of scattered pilots in an OFDM symbol

N_E: The number of excluded subcarriers in an OFDM symbol

N_P: The number of PLC subcarriers in an OFDM symbol

N_D: The number of data subcarriers in an OFDM symbol

The values of N, N_C, N_E and N_P do not change from symbol to symbol for a given OFDM template; the values of N_S and N_D change from symbol to symbol.

The following equation holds for all symbols:

$$N = N_C + N_S + N_E + N_P + N_D$$

The value of N is 4 096 for 50 kHz subcarrier spacing and 8 192 for 25 kHz subcarrier spacing. From this equation it is clear that $(N_S + N_D)$ is a constant for a given OFDM template. Therefore, although the number of data subcarriers (N_D) and the number of scattered pilots (N_S) in an OFDM symbol changes from symbol to symbol, the sum of these two numbers is invariant over all symbols. Interleaving and de-interleaving are applied to the set of data subcarriers and scattered pilots of size $N_I = N_D + N_S$.

7.5.5.5.1 Bit Loading

The bit loading pattern defines the QAM constellations assigned to each of the 4 096 or 8 192 subcarriers of the OFDM transmission. This bit loading pattern can change from profile to profile. Continuous pilot locations, PLC locations and exclusion bands are defined separately, and override the values defined in the bit-loading profile. Let the bit loading pattern for profile i be defined as $A_i(k)$, where:

k is the subcarrier index that goes from 0 to (N-1)

N is either 4 096 or 8 192

$A_i(k) \in \{0, 4, 6, 7, 8, 9, 10, 11, 12, 13, 14\}$. A value of 0 indicates that the subcarrier k is zero-bit-loaded. Other values indicate that the modulation of subcarrier k is QAM with order $2^{A_i(k)}$.

Let the sequence $\{A_i(k),\ \ k = 0, 1, \dots\ ,\ (N-1),\ \ k \notin S^{PCE}\}$ be arranged as N_I consecutive values of another sequence:

$$B_i(k),\ \ k = 0, 1, \dots\ ,\ (N_I - 1)$$

Given the locations of the excluded subcarriers, continuous pilots and the PLC in the OFDM template, it is possible to obtain the bit-loading pattern $B_i(k)$ that is applicable only to spectral locations excluding excluded subcarriers, continuous pilots, and PLC subcarriers. However, note that $B_i(k)$ does contain the spectral locations occupied by scattered pilots; these locations change from symbol to symbol.

It is more convenient to define bit loading profiles in the domain in which subcarriers are transmitted. It is in this domain that signal-to-noise-ratios of subcarriers are calculated. Furthermore, defining the bit-loading patterns in the transmission domain allows significant data compression to be achieved, because a relatively large number of contiguous spectral locations can share the same QAM constellation.

Although the bit loading pattern is defined in the domain in which subcarriers are transmitted, the bit loading is not applied in that domain. Bit loading is applied prior to interleaving, as shown in figure 7.48. Hence there is a permutation mapping of subcarriers, defined by the interleaving function, between the domain in which bit loading is applied to subcarriers and the domain in which subcarriers are transmitted.

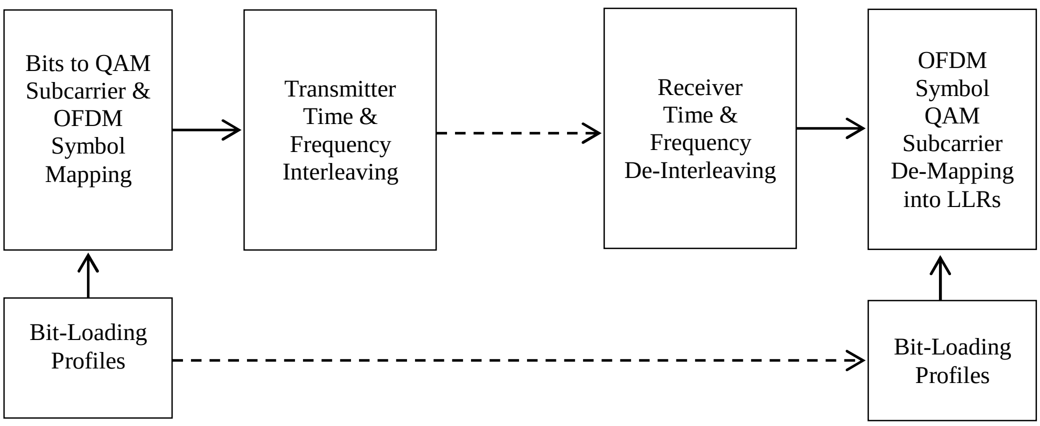

Figure 7.48: Bit Loading, Symbol Mapping, and Interleaving

The excluded subcarriers, PLC subcarriers, and continuous pilots are excluded from the processes of interleaving and de-interleaving; scattered pilots and data subcarriers are subject to interleaving and de-interleaving. Hence, the total number of subcarriers that pass through the interleaver and de-interleaver is $N_I = (N_D + N_S)$ and this number does not change from symbol to symbol.

The interleaver introduces a 1-1 permutation mapping P on the N_I subcarriers. Although interleaving consists of a cascade of two components, namely time and frequency interleaving, it is only frequency interleaving that defines the mapping P. This is because time interleaving does not disturb the frequency locations of subcarriers.

The corresponding permutation mapping applied at the receiver de-interleaver is P^{-1}.

In order to perform bit-loading, it is necessary to work out the bit loading pattern at the node at which it is applied, i.e. at the input to the interleavers. This is given by:

$$C_i(k) = P^{-1}(B_i(k))$$

Since the time interleaver does not change the frequency locations of subcarriers, the sequence $C_i(k)$ is obtained by sending $\{B_i(k), k = 1, 2, \ldots, N_I - 1\}$ through the frequency de-interleaver.

Note that $C_i(k)$ gives the bit-loading pattern for N_I subcarriers. Yet, some of these subcarriers are scattered pilots that have to be avoided in the bit-loading process. Hence, a two-dimensional binary pattern $D(k, j)$ is used to identify subcarriers to be avoided during the process of bit-loading. Because the scattered pilot pattern has a periodicity of 128 in the time dimension, this binary pattern also has periodicity 128 in the column dimension *j*.

$$D(k, j) \text{ is defined for } k = 0, 1, \ldots, (N_I - 1) \text{ and for } j = 0, 1, \ldots, 127$$

The process to create the binary pattern $D(k, j)$ begins with the transmitted scattered pilot pattern defined in clause 7.5.6. There are two scattered pilot patterns, one for 4K FFTs and the other for 8K FFTs; both patterns are defined in reference to the preamble of the PLC and have a periodicity of 128 symbols.

The CMTS executes the following steps to obtain the pattern $D(k, j)$:

1) Define a two-dimensional binary array $P(k, j)$ in the subcarrier transmitted domain that contains a one for each scattered pilot location and zero otherwise:

$$P(k, j), \text{ for } k = 0, 1, \ldots, N - 1 \text{ and for } j = 0, 1, \ldots, 127$$

Here, the value of *N* is either 4 096 or 8 192. The first column of this binary sequence corresponds to the first OFDM symbol following the preamble of the PLC.

1) Exclude the rows corresponding to excluded subcarriers, continuous pilots, and PLC from the two-dimensional array $P(k, j)$ to give an array $Q(k, j)$. The number of rows of the resulting array is N_I and the number of columns is 128.

2) Pass this two-dimensional binary array $Q(k, j)$ through the frequency de-interleaver and then the time de-interleaver, with each column treated as an OFDM symbol. After the 128 columns of the pattern have been input into the interleaver, re-insert the first *M* columns, where *M* is the depth of the time interleaver. This is equivalent to periodically extending $Q(k, j)$ along the dimension j and passing (128+*M*) columns of this extended sequence through the frequency de-interleaver and the time de-interleaver.

3) Discard the first *M* symbols coming out of the time de-interleaver and collect the remaining 128 columns into an array to give the binary two-dimensional array $D(k, j)$ of size $(N_I \times 128)$.

For bit loading the CMTS accesses the appropriate column j of the binary pattern bit $D(k, j)$ together with the appropriate bit loading profile $C_i(k)$. If the value of the bit $D(k, j)$ is 1, the CMTS shall skip this subcarrier k and move to the next subcarrier. This subcarrier is included as a placeholder for a scattered pilot that will be inserted in this subcarrier location after interleaving. After each symbol the column index *j* has to be incremented modulo 128.

The CMTS shall use this binary two-dimensional array $D(k, j)$ of size $(N_I \times 128)$ in order to do bit-loading of OFDM subcarriers, as described earlier in this clause.

The corresponding operation in the CM is de-mapping the QAM subcarriers to get Log-Likelihood-Ratios (LLRs) corresponding to the transmitted bits. This operation, described below, is much simpler than the mapping operation in the transmitter.

The scattered pilots and data subcarriers of every received symbol are subjected to frequency and time de-interleaving. The scattered pilots have to be tagged so that these can be discarded at the output of the time and frequency de-interleavers. This gives N_I subcarriers for every OFDM symbol. The CM accesses theses N_I de-interleaved subcarriers together with the bit-loading pattern $C_i(k)$ to implement the de-mapping of the QAM subcarriers into LLRs. If the subcarrier k happens to be a scattered pilot, then this subcarrier, as well as the corresponding value $C_i(k)$, is skipped and the CM moves to the next subcarrier $(k + 1)$.

7.5.5.5.2 NCP Insertion

Next Codeword Pointers (NCPs) point to the beginning of codewords in a symbol, counting only data subcarriers of that symbol, not including the locations reserved for the scattered pilots. The format of an NCP is described in clause 8.3.4, which also describes the FEC applied to the NCP. Each FEC encoded NCP is 48 bits wide. NCPs may be modulated using QPSK, 16-QAM or 64-QAM and this modulation is signalled by PLC. In addition to the NCPs carrying next codeword pointers, there will also be a NCP carrying the CRC for all the NCPs of the symbol. The CRC is generated as described in Annex E. As the NCPs are constructed while the OFDM symbols are being constructed, the NCPs are inserted in the opposite direction to data and beginning from the opposite end. Data is inserted beginning from the low frequency towards the high frequency end. The NCPs are inserted from the high frequency end towards the low frequency end.

Note that N_I subcarriers in each symbol are subjected to the data and NCP mapping operation. These subcarriers consist of data subcarriers and scattered pilot place-holder subcarriers as described in the preceding section. During the course of mapping data or NCP subcarriers, if a scattered pilot placeholder is encountered, this is skipped.

The figure given below shows an OFDM symbol comprising a Data segment, an NCP segment and a "Filler" segment. "Filler" subcarriers have to be inserted into the OFDM symbol when the number of codewords in the OFDM symbol has exceeded the upper limit or when it is not possible to begin a new codeword because of insufficient space to include a NCP. These filler subcarriers are zero-bit-loaded.

The CMTS shall use zero-bit-loaded filler subcarriers when the number of codewords has exceeded the upper limit or when it is not possible to begin a new codeword because of insufficient space to include a NCP. The CMTS shall define the location of a segment of zero-bit-loaded subcarriers using an NCP with Z-bit set to one as described in clause 8.3.4.1.

If the CMTS has no data to transmit, the CMTS shall adopt one of the following two options:

1) Insert zero-bit-loaded filler sub-carriers into OFDM symbols as described in this clause; or

2) Insert stuffing pattern of 0xFF bytes into codewords as described in clause 8.3.2.

Data segment contains codewords belonging to several profiles. Some of the subcarriers may be zero-bit-loaded in some of the profiles. The NCP also has a profile. This profile allows some of the subcarriers in the NCP segment to be zero-bit-loaded. Note that the NCP modulation is a constant for given OFDM transmission. It does not change from subcarrier to subcarrier.

Note that throughout the symbol there can be scattered pilot placeholders. These have to be skipped during the insertion of data subcarriers, NCP subcarriers or filler subcarriers. Moreover, these have to be tagged before sending the N_I subcarriers through the time and frequency interleavers. Scattered pilots will be inserted in their place with the appropriate BPSK modulation before the data is transmitted.

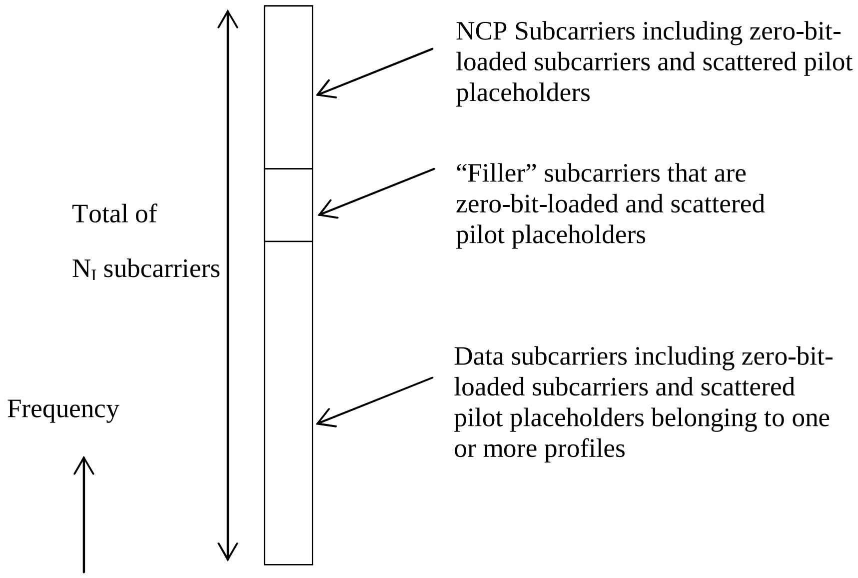

Figure 7.49: NCP Insertion

7.5.6 Interleaving and De-interleaving

7.5.6.0 Interleaving and De-interleaving Overview

To minimize the impacts of burst noise and ingress on the DOCSIS signals, time and frequency interleaving are applied to OFDM symbols in the following order: time interleaving, then frequency interleaving. These interleaving methods are discussed in this clause.

The time interleaver is a convolutional interleaver that operates in the time dimension on individual subcarriers of a sequence of OFDM symbols. The time interleaver does not change the frequency location of any OFDM subcarrier. A burst event can reduce the SNR of all the subcarriers of one or two consecutive OFDM symbols; the purpose of the time interleaver is to disperse these burst-affected OFDM subcarriers between M successive OFDM symbols, where M is the interleaver depth. This dispersion distributes the burst-affected subcarriers uniformly over a number of LDPC codewords.

The frequency interleaver works along the frequency dimension. The frequency interleaver changes the frequency locations of individual OFDM subcarriers; latency is not introduced, except for the data store and read latency. The aim of frequency interleaving is to disperse ingress, e.g. LTE that affects a number of consecutive subcarriers over the entire OFDM symbol. Frequency interleaving distributes the burst-affected subcarriers over a number of LDPC codewords.

The CMTS first applies a time interleaver to an OFDM symbol worth of N_I subcarriers to get a new set of N_I subcarriers. These N_I subcarriers are made up of N_D data subcarriers and N_S scattered pilots.

$$N_I = N_D + N_S$$

It is important to note that although N_D and N_S are not the same for every OFDM symbol, the value of N_I is a constant for all OFDM symbols in a given system configuration. The value of N_I is a function of the channel bandwidth, number of excluded subcarriers, number of PLC subcarriers and the number of continuous pilots. The CMTS then subjects these N_I subcarriers to frequency interleaving. The value of N_I does not exceed 7 537 for 8K FFT mode and 3 745 for the 4K FFT mode.

Note that both time and frequency interleaving are applied only to data subcarriers and scattered pilots. Continuous pilot, subcarriers that have been excluded (used to support legacy channels in spectral regions, for example) and the subcarriers of the physical layer link channel (PLC) are not interleaved. The CMTS shall not interleave continuous pilots, excluded subcarriers or the subcarriers of the PLC.

7.5.6.1 Time Interleaving

The CMTS shall time interleave as described in this clause. The CMTS shall time interleave after OFDM symbols have been mapped to QAM constellations and before they are frequency interleaved.

The time interleaver is a convolutional interleaver that operates at the OFDM subcarrier level. If the depth of the interleaver is *M*, then there are *M* branches, as shown in figure 7.50.

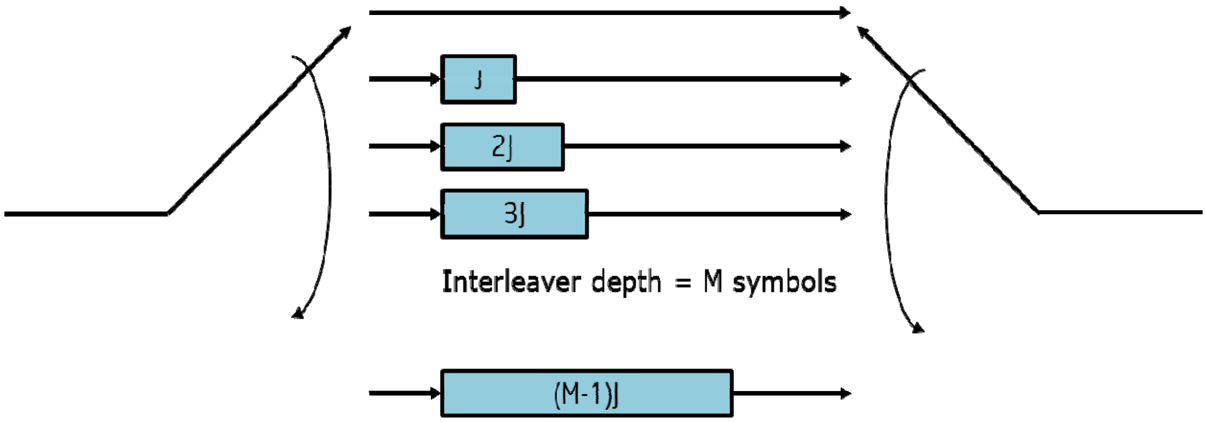

Figure 7.50: Time Interleaver Structure

The CMTS shall support a maximum value of M equal to 32 for 20 µs symbol duration (50 kHz subcarrier spacing) and 16 for 40 µs symbol duration (25 kHz subcarrier spacing).

The CMTS shall support all values of M from 1 to the maximum value of M (inclusive of both limits).

Each branch is a delay line; the input and output will always be connected to the same delay line. This delay line will be clocked to insert a new subcarrier into the delay line and to extract a subcarrier from the delay line. Next, the commutator switches at the input, and the output will move to the next delay line in the direction shown by the arrows in figure 7.40. After the delay line with the largest delay, the switch will move to the delay line with zero delay.

The lowest frequency subcarrier of an OFDM symbol always goes through the branch with zero delay. Then the commutator switch at input and the corresponding commutator switch at output are rotated by one position for every new subcarrier.

The value of *J* is given by the following equation:

$$J = ceil(\frac{N_I}{M})$$

Here, N_I is the number of data subcarriers and scattered pilots in an OFDM symbol. See clause 7.5.6.3 for details on interleaving scattered pilots.

If N_I were not divisible by *M*, all of the branches would not be filled. Therefore, "dummy subcarriers" are added to the symbol to make the number of subcarriers equal to a multiple of *M*. The number of dummy subcarriers is given by:

$$J * M - N_I$$

The dummy subcarriers are added for definition purposes only; at the output of the interleaver these dummy subcarriers are discarded. An implementation will use a single linear address space for all the delay lines in figure 7.40. Writing and reading dummy subcarriers will not be needed.

7.5.6.2 Frequency Interleaving

The CMTS shall frequency interleave OFDM symbols as described in this clause. The CMTS shall frequency interleave after OFDM symbols have been time interleaved.

The frequency interleaver works on individual OFDM symbols. Each symbol to be interleaved consists of N_I subcarriers. These N_I subcarriers are made up of N_D data subcarriers and N_S scattered pilot placeholders. Although N_D and N_S are not the same for every symbol, the value of N_I is a constant for all OFDM symbols in a given system configuration. See clause 7.5.6.3 for details on interleaving scattered pilots.

There is a 2-D store comprising 128 rows and *K* columns. If the number of data subcarriers and scattered pilots in the OFDM symbol is N_I, then the number of columns, *K*, is given by the following equation:

$$K = ceil(\frac{N_I}{128})$$

If N_I is not an integer multiple of 128, then the last column will only be partially filled during the frequency interleaving process. The number of data subcarriers in the last column, *C*, is given by:

$$C = N_I - 128(K - 1)$$

The frequency interleaver follows the following process; note that rows are numbered 0 to 127, and columns are numbered from 0 to (*K*-1):

1) Write the subcarriers along rows of the 2-D store. Rows are accessed in bit-reversed order. For example, after writing in row 0, the next writing operation will be in row 64. This will be followed by writing in row 32 and so on. If the row number is less than *C*, then *K* subcarriers will be written in the row. Otherwise only (*K*-1) subcarriers will be written. (If the number N_I is an integer multiple of 128 then *C* will be zero. Then *K* subcarriers will be written in every row.)

2) Rotate columns 0 to (*K*-2) by an amount given by a 6-bit shift linear feedback (maximal length) shift register. This shift register is initialized to a value of 17 at the start of each OFDM symbol. The final column, which may be partially full, is not rotated.

3) Read the columns in bit-reverse order, starting at column 0, then column bit-reverse(1), then column bit-reverse(2), … , ending at column bit-reverse(K-1). When K is not a power-of-2, bit-reverse(x), for x = 0, …, K-1, is defined by:

   bit-reverse(x) = reverse_bits(x), if reverse_bits(x) < K; OR
   x, if reverse_bits(x) ≥ K

where reverse_bits(x) is the number obtained by reversing the order of the bits in the m-bit representation of x, with m being the number of bits in K.

The structure of the two-dimensional store is shown in figure 7.51.

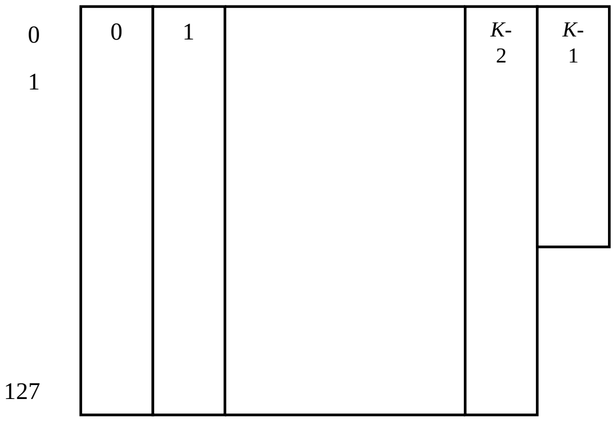

Figure 7.51: Two-Dimensional Block Structure

The linear feedback shift register is defined using the following equation in Galois field GF[2^6].

$$x(i) = \alpha x(i - 1), for\ i = 1, \dots, 127, where\ x(0) = \alpha^5 + \alpha$$

GF[2^6] is defined using the polynomial ($\alpha^6 + \alpha$ +1). As this is primitive, powers of α will generate all 63 non-zero elements of the field. This operation can be represented as the linear feedback shift register, depicted in figure 7.52.

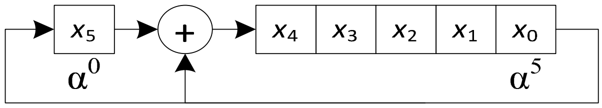

Figure 7.52: Linear Feedback Shift Register

The binary number *x*[5:0] is used to rotate the columns. This number is initialized to 17 at the beginning of each OFDM symbol. The column number 0 is rotated by 17; subsequent columns are rotated by values obtained by clocking the shift register shown in figure 7.52. The rotation applied to the first column is defined in figure 7.53. Subsequent rows are also rotated along the same direction.

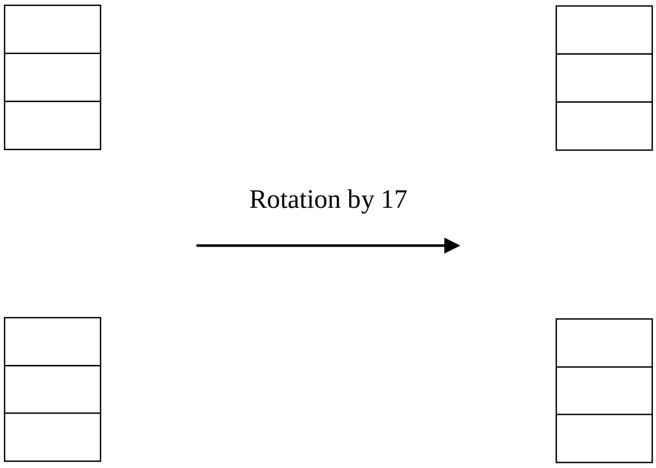

Figure 7.53: Frequency Interleaver Rotation Definition

Note that column (*K-1*) is not rotated, regardless of whether it is full: because all other columns are rotated by a non-zero amount, there is no need to rotate column (*K-1*). The C code for interleaver implementation is given in Annex F.

7.5.6.3 Interleaving Impact on Continuous Pilots, Scattered Pilots, PLC and Excluded Spectral Regions

DOCSIS 3.1 transmissions contain continuous pilots for receiver synchronization and scattered pilots for channel estimation. In addition, there could be excluded regions to accommodate legacy channels. There will also be a physical layer link channel (PLC).

The CMTS interleaves scattered pilots and data subcarriers, but does not interleave continuous pilots, the PLC, and subcarriers belonging to excluded regions. With respect to scattered pilots, it is noted here that CMTS actually interleaves the subcarriers that are tagged to act as placeholders for scattered pilots, since at the time of interleaving the scattered pilots have not yet been inserted. The actual BPSK modulation to these placeholder subcarriers is applied after interleaving as described in clause 7.5.15.

The CMTS inserts scattered pilot placeholders prior to time and frequency interleaving such that when these placeholders get time and frequency interleaved, the resulting placeholders conform to the required scattered pilot pattern described in clause 7.5.15.

To accomplish this, the CMTS has to retain a reference pattern for inserting scattered pilot placeholders prior to interleaving. Since the scattered pilot pattern repeats every 128 symbols, this pattern is a (N_I x 128) two-dimensional bit pattern. A value of one in this bit-pattern indicates the location of a scattered pilot. The CMTS inserts data subcarriers where this reference pattern has a zero and scattered pilot placeholders where this pattern has a one.

This reference pattern may be derived from the following procedure:

1) In the time-frequency plane, create a two-dimensional bit-pattern of zeros and ones from the transmitted "diagonal" scattered pilot patterns described in clause 7.5.15. This pattern has a periodicity of 128 symbols and has a value of one for a scattered pilot location and zero otherwise. Let the time axis be horizontal and the frequency axis vertical.

2) Delete all horizontal lines containing continuous pilots, excluded subcarriers, and PLC from the above mentioned two-dimensional bit pattern; note the some scattered pilots could coincide with continuous pilots. These locations are treated as continuous pilot locations.

3) Send the resulting bit-pattern through the frequency de-interleaver and the time de-interleaver in succession. This will give another two-dimensional bit pattern that has a periodicity of 128 symbols. The appropriate 128-symbol segment of this bit-pattern is chosen as the reference bit pattern referred to above.

Note that the CMTS has to synchronize the scattered pilot pattern to the PLC preamble, as described in clause 7.5.15. This uniquely defines the 128-symbol segment that has to be used as the reference pattern.

Scattered pilots are not in the same subcarrier location in every symbol; hence some scattered pilots can coincide with continuous pilots in some OFDM symbols. The size of the overlap between the set of scattered pilots and the set of continuous pilots will change from symbol to symbol. As a result, the number of data subcarriers in a symbol will not be the same for all OFDM symbols. Note that in the nomenclature used below, when a scattered pilot coincides with a continuous pilot, then that pilot is referred to as a continuous pilot.

Although the number of data subcarriers can change from symbol to symbol, the number of data subcarriers and scattered pilots are the same for every symbol. This is referred to as N_I in this clause. Let N_D denote the number of data subcarriers in a symbol and N_S denote the number of scattered pilots in a symbol. These two parameters, i.e. N_D and N_S, will change from symbol to symbol. However, the sum of these two, i.e. N_I is a constant for a given system configuration.

$$N_I = N_S + N_D$$

Hence the number of OFDM subcarriers that are interleaved does not change from symbol to symbol. This is important, because if not for this, the output of the convolutional time interleaver may have dummy or unused subcarriers in the middle of interleaved OFDM symbols.

The insertion of continuous pilots, PLC and excluded regions happens after both time and frequency interleaving.

Interleaving data and scattered pilots together has another important advantage. This is to do with bit loading. A transmitted profile is said to have non-uniform bit loading if the QAM constellation that is applied to subcarriers is not constant over the entire frequency band. If the data subcarriers are interleaved and scattered pilots are added later, then the data subcarriers will have to be shifted to accommodate the scattered pilots. This shift will be different from symbol to symbol, and this complicates non-uniform bit-loading. Hence, having the scattered pilots in-place during the bit-loading process greatly simplifies the bit loading operation. The insertion of continuous pilots, PLC and excluded regions also results in shift of data subcarriers, but this shift is the same for every symbol, and can easily be accounted for in the bit loading process.

The CMTS only interleaves data subcarriers and scattered pilots, and therefore only needs information about the number of data subcarriers and scattered pilots per symbol. In addition, the interleaver does not need to know what modulation has been applied to an individual data subcarrier. Regardless of modulation scheme, all OFDM symbols will have the same number of data subcarriers and scattered pilots, and the modulation pattern of these data subcarriers may change from symbol to symbol.

7.5.7 IDFT

7.5.7.1 Downstream Transmitter Inverse Discrete Fourier Transform

The CMTS transmitter shall use the IDFT definition and subcarrier referencing method described in this clause.

This clause defines the inverse discrete Fourier transform (IDFT) used in the CMTS transmitter for DOCSIS 3.1. OFDM subcarrier referencing for other definitions such as PLC location, continuous pilots, exclusion bands and bit loading is also described.

The OFDM signal assembled in the frequency domain consists of 4 096 subcarriers for the 4K FFT and 8 192 subcarriers for the 8K FFT. The OFDM signal is composed of:

- Data subcarriers
- Scattered pilots
- Continuous pilots
- PLC subcarriers
- Excluded subcarriers that are zero valued

This signal is described according to the following IDFT equation:

$$x(i) = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} X(k) \exp\left(j\frac{2\pi i\left(k-\frac{N}{2}\right)}{N}\right), \quad for \;\; i = 0, 1, \ldots, (N-1)$$

The resulting time domain discrete signal, $x(i)$, is a baseband complex-valued signal, sampled at 204,8 Msamples per second.

In this definition of the IDFT:

$X(0)$ is the lowest frequency component;

$X(N\text{-}1)$ is the highest frequency component.

The IDFT is illustrated in figure 7.54.

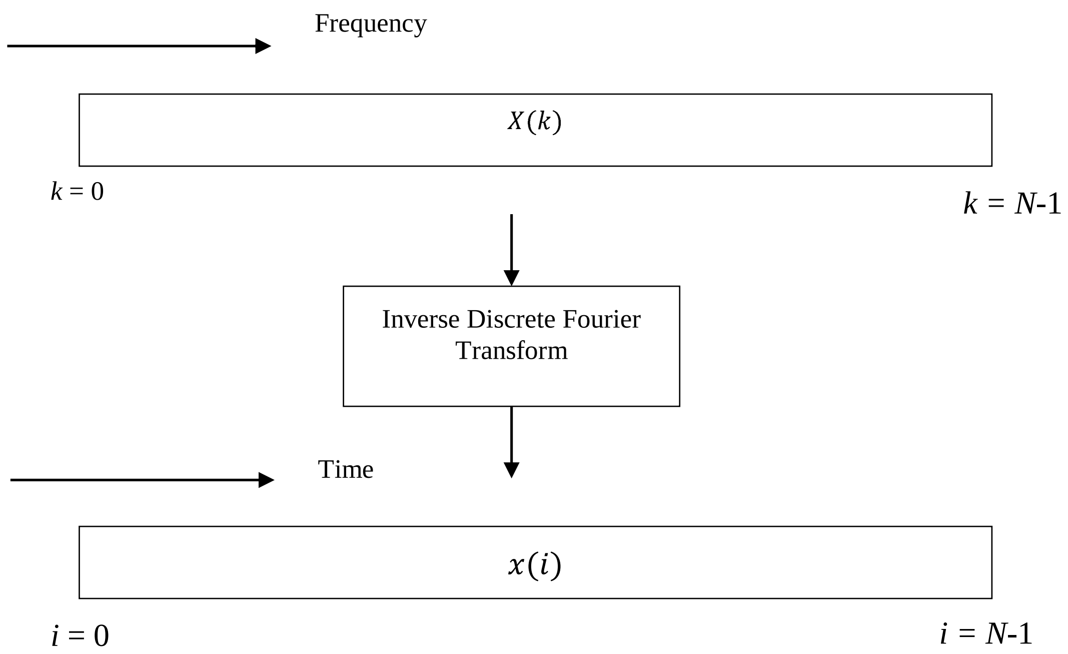

Figure 7.54: Inverse Discrete Fourier Transform

The sample rate in the time domain is 204,8 Msamples/s. Hence, the N samples of the discrete Fourier transform cover a frequency range of 204,8 MHz. This gives the subcarrier spacing shown in table 7.33.

Table 7.33: Subcarrier Spacing

IDFT Size *N*	Carrier Spacing
4 096	50 kHz
8 192	25 kHz

The maximum channel bandwidth is 192 MHz; this corresponds to 3 841 subcarriers in 4K mode and 7 681 subcarriers in 8K mode. The active bandwidth of the channel is expected to be 190 MHz; this corresponds to 3 800 subcarriers in 4K mode and 7 600 subcarriers in 8K mode.

7.5.7.2 Subcarrier Referencing

It is necessary to refer to specific OFDM subcarriers for several definitions:

a) Defining continuous pilot locations

b) Defining exclusion bands and excluded individual subcarriers

c) Defining bit loading profiles

Each of these definitions uses the index *k* of the equation defined in the preceding section to refer to a specific subcarrier.

The subcarrier index goes from 0 to 4 095 for the 4K FFT and from 0 to 8 191 for the 8K FFT; each of these definitions is limited to these subcarrier indices.

The PLC is also defined with reference to $k = 0$. The OFDM template carried by the PLC defines the subcarrier index of the lowest frequency subcarrier of the PLC. Hence, once the CM detects the PLC, the CM knows the location of $k = 0$. Since the FFT size is also known, it is possible to precisely compute the FFT of the data channel containing the PLC.

Note that scattered pilot placement is not referenced to $k = 0$; instead, it is referenced directly to the PLC preamble.

7.5.8 Cyclic Prefix and Windowing

This clause describes how cyclic prefixes are inserted and how a window is applied to the output of the IDFT at the CMTS and how they are handled by the CM.

The addition of a cyclic prefix enables the receiver to overcome the effects of inter-symbol-interference caused by micro-reflections in the channel. Windowing maximizes channel capacity by sharpening the edges of the spectrum of the OFDM signal. Spectral edges occur at the two ends of the spectrum of the OFDM symbol, as well as at the ends of internal exclusion bands.

The number of active OFDM subcarriers can be increased by sharpening these spectral edges. However, sharper spectral edges in the frequency domain imply longer tapered regions in the time domain, resulting in increased symbol duration and reduction in throughput. Therefore, there is an optimum amount of tapering that maximizes channel capacity. This optimum is a function of channel bandwidth as well as the number of exclusion bands.

The CMTS shall follow the procedure described in clause 7.4.11.1 for cyclic prefix insertion and windowing, using CMTS specific cyclic prefix and roll-off period values.

The CMTS shall support cyclic prefix extension and windowing as described in clause 7.4.11.1.

The CMTS shall support the cyclic prefix values defined in table 7.34 for both 4K and 8K FFTs.

The CM shall support the cyclic prefix values listed defined table 7.34 for both 4K and 8K FFTs.

Table 7.34: Cyclic Prefix (CP) Values

Cyclic Prefix (µs)	Cyclic Prefix Samples (N_{cp})
0,9375	192
1,25	256
2,5	512
3,75	768
5,0	1 024

The cyclic prefix (in µs) are converted into samples using the sample rate of 204,8 Msamples/s and is an integer multiple of: 1/64 * 20 µs.

The CMTS shall support the five parameter values specified for this roll-off listed in table 7.35.

Table 7.35: Roll-off Period (RP) Values

Roll-Off Period (µs)	Roll-Off Period Samples (N_{cp})
0	0
0,3125	64
0,625	128
0,9375	192
1,25	256

The CMTS shall not allow a configuration in which the RP value is ≥ the CP value.

7.5.9 Fidelity Requirements

7.5.9.0 Fidelity Requirements Overview

For the purposes of the present document, the number of occupied CEA channels of an OFDM channel is the occupied bandwidth of the OFDM channel divided by 6 MHz.

CMTSs capable of generating N-channels of legacy DOCSIS plus N_{OFDM}-channels of OFDM per RF port, for purposes of the DRFI output electrical requirements, the device is said to be capable of generating N_{eq}-channels per RF port, where $N_{eq} = N + 32*N_{OFDM}$ "equivalent legacy DOCSIS channels."

An N_{eq}-channel per RF port CMTS shall comply with all requirements operating with all N_{eq} channels on the RF port, and shall comply with all requirements for an N_{eq}'-channel per RF port device operating with N_{eq}' active channels on the RF port for all values of N_{eq}' less than N_{eq}.

For an OFDM channel there is a) the occupied bandwidth, b) the encompassed spectrum, c) the modulated spectrum, and d) the number of equivalent legacy DOCSIS channels.

The encompassed spectrum in MHz is 204,8 MHz minus the number of subcarriers in the Band edge Exclusion Sub-band for the upper and lower band edges (combined) times the subcarrier spacing in MHz. For example, with subcarrier spacing of 50 kHz and 150 lower band edge subcarriers and 152 upper band edge subcarriers for a total of 302 subcarriers in the two Band edge Exclusion Sub-bands, the encompassed spectrum = 204,8 - 302*(0,05) = 189,7 MHz. The encompassed spectrum is also equal to the centre frequency of the highest frequency modulated subcarrier minus the centre frequency of the lowest frequency modulated subcarrier in an OFDM channel, plus the subcarrier spacing.

The occupied bandwidth is a multiple of 6 MHz, with a minimum of 24 MHz, and consists of all CEA channels which include the encompassed spectrum plus taper region shaped by the OFDM channels' transmit windowing; the out-of-band spurious emissions requirements (except for gap channel spurious emissions requirements) apply outside the occupied bandwidth. With a 1 MHz taper region on each band edge of the OFDM channel, shaped by the transmit windowing function, encompassed spectrum of 189,7 MHz may provide 192 MHz of occupied bandwidth.

The modulated spectrum of an OFDM channel is the encompassed spectrum minus the total spectrum in the Internal Excluded Sub-bands of the channel, where the total spectrum in the Internal Excluded Sub-bands is equal to the number of subcarriers in all of the Internal Excluded Sub-bands of the OFDM channel multiplied by the subcarrier spacing of the OFDM channel. In the previous example, if there are 188 subcarriers total in three Internal Exclusion Sub-bands, then the total spectrum in the Internal Excluded Sub-bands (in MHz) is 188*0,05 = 9,4 MHz, and the modulated spectrum is 189,7 MHz - 9,4 MHz = 180,3 MHz.

The number of equivalent active legacy DOCSIS channels in the OFDM channel N_{eq}' is the ceiling function applied to the modulated spectrum divided by 6 MHz. For the example, the number of equivalent legacy DOCSIS channels in the OFDM channel is ceiling (180,3 MHz / 6 MHz) = 31.

For an N_{eq}-channel per RF port device, the applicable maximum power per channel and spurious emissions requirements are defined using a value of $N^* = \text{minimum}(4N_{eq}', \text{ceiling}[N_{eq}/4])$ for $N_{eq}' < N_{eq}/4$, and $N^* = N_{eq}'$ otherwise.

These specifications assume that the CMTS will be terminated with a 75 ohm load.

7.5.9.1 CMTS Output Electrical Requirements

For OFDM, all modulated subcarriers in an OFDM channel are set to the same average power (except pilots which are boosted by 6 dB). For purposes of spurious emissions requirements, the "commanded transmit power per channel" for an equivalent legacy DOCSIS channel is computed as follows:

- CMTS power is configured by power per CEA channel and number of occupied CEA channels for each OFDM channel.
- For each OFDM channel, the total power is Power per CEA channel + $10\log_{10}$(Number of occupied CEA channels) for that OFDM channel.
- CMTS calculates power for data subcarrier and pilots (using total number of non-zero valued (non-excluded) subcarriers).
- CMTS calculates power in 6 MHz containing PLC.
- For the spurious emissions requirements, power calculated for the 6 MHz containing the PLC is the commanded average power of an equivalent DOCSIS legacy channel for that OFDM channel.

A CMTS shall output an OFDM RF modulated signal with the characteristics defined in tables 7.37 and 7.38. Legacy DOCSIS RF modulated signal characteristics are provided in clause 6.2.22.

The condition for these requirements is all N_{eq}' combined channels, legacy DOCSIS channels and equivalent legacy DOCSIS channels, commanded to the same average power, except for the Single Channel Active Phase Noise, Diagnostic Carrier Suppression, OFDM Phase Noise, OFDM Diagnostic Suppression, and power difference requirements, and except as described for Out-of-Band Noise and Spurious Requirements.

Table 7.36: RF Output Electrical Requirements

Parameter	Value
Downstream Lower Edge Band of a CMTS	CMTS shall support 258 MHz CMTS should support 108 MHz
Downstream Upper Edge Band of a CMTS	CMTS shall support 1 218 MHz CMTS may support 1 794 MHz
Level	Adjustable. See table 7.37
Modulation Type	See clause 7.5.6.1
OFDM channels' subcarrier spacing	25 kHz and 50 kHz
Inband Spurious, Distortion, and Noise 528 MHz total occupied bandwidth, 6 MHz gap (Internal Excluded subcarriers) 88 equivalent legacy DOCSIS channels. (See Notes 4, 6) For measurements below 600 MHz For measurements from 600 MHz to 1 002 MHz For measurements 1 002 MHz to 1 218 MHz	 ≤ -50 dBr Average over centre 400 kHz subcarriers within gap ≤ -47 dBr Average over centre 400 kHz subcarriers within gap ≤ -45 dBr Average over centre 400 kHz subcarriers within gap

Parameter	Value
MER in 192 MHz OFDM channel occupied bandwidth 528 MHz total occupied bandwidth, 88 equivalent legacy DOCSIS channels. (See Notes 2, 4, 5, 6) For measurements below 600 MHz For measurements from 600 MHz to 1 002 MHz For measurements 1 002 MHz to 1 218 MHz	≥ 48 dB Any single subcarrier. See Note 1 ≥ 50 dB Average over the complete OFDM channel. See Note 1 ≥ 45 dB Any single subcarrier. See Note 1 ≥ 47 dB Average over the complete OFDM channel. See Note 1 ≥ 43 dB Any single subcarrier. See Note 1 ≥ 45 dB Average over the complete OFDM channel. See Note 1 Minimal test receiver equalization: See Note 7 2 dB relief for above requirements (e.g. MER > 48 dB becomes MER > 46 dB)
MER in 24 MHz OFDM channel occupied bandwidth, single OFDM channel only, 24 MHz total occupied bandwidth: See Notes 1, 2, 4, 8 For measurements below 600 MHz For measurements from 600 MHz to 1 002 MHz For measurements 1 002 MHz to 1 218 MHz	≥ 48 dB Average over the complete OFDM channel. ≥ 45 dB Average over the complete OFDM channel. ≥ 43 dB Average over the complete OFDM channel.
Phase noise, double sided maximum, Full power CW signal 1 002 MHz or lower Full power 192 MHz OFDM channel block with 6 MHz in centre as Internal Exclusion subband + 0 dBc CW in centre, with block not extending beyond 1 002 MHz [CW not processed via FFT] Full power 192 MHz OFDM channel block with 24 MHz in centre as Internal Exclusion subband + 0 dBc CW in centre, with block not extending beyond 1 002 MHz [CW not processed via FFT] Full power 192 MHz OFDM channel block with 30 MHz in centre as Internal Exclusion subband + 7 dBc CW in centre, with block not extending beyond 1 002 MHz [CW not processed via FFT]	1 kHz - 10 kHz: -48 dBc 10 kHz - 100 kHz: -56 dBc 100 kHz - 1 MHz: -60 dBc 1 MHz - 10 MHz: -54 dBc 10 MHz - 100 MHz: -60 dBc 1 kHz - 10 kHz: -48 dBc 10 kHz - 100 kHz: -56 dBc 100 kHz - 1 MHz: -60 dBc 1 MHz - 10 MHz: -53 dBc
Output Impedance	75 ohms
Output Return Loss (Note 3)	> 14 dB within an active output channel from 88 MHz to 750 MHz > 13 dB within an active output channel from 750 MHz to 870 MHz > 12 dB within an active output channel from 870 MHz to 1 218 MHz > 12 dB in every inactive channel from 54 MHz to 870 MHz > 10 dB in every inactive channel from 870 MHz to 1 218 MHz

Parameter	Value
Connector	F connector per [17] or [19]
NOTE 1: Receiver channel estimation is applied in the test receiver; test receiver does best estimation possible. Transmit windowing is applied to potentially interfering channel and selected to be sufficient to suppress cross channel interference. NOTE 2: MER (modulation error ratio) is determined by the cluster variance caused by the transmit waveform at the output of the ideal receive matched filter. MER includes all discrete spurious, noise, subcarrier leakage, clock lines, synthesizer products, distortion, and other undesired transmitter products. Phase noise up to ±50 kHz of the subcarrier is excluded from inband specification, to separate the phase noise and inband spurious requirements as much as possible. In measuring MER, record length or carrier tracking loop bandwidth may be adjusted to exclude low frequency phase noise from the measurement. MER requirements assume measuring with a calibrated test instrument with its residual MER contribution removed. NOTE 3: Frequency ranges are edge-to-edge. NOTE 4: Phase noise up to 10 MHz offset is mitigated in test receiver processing or by test equipment (latter using hardline carrier from modulator, which requires special modulator test port and functionality). NOTE 5: Up to 5 subcarriers in one OFDM channel can be excluded from this requirement. NOTE 6: The measured OFDM channel is allocated 204,8 MHz of spectrum which is free from the other OFDM channel and 24 SC-QAM channels which together comprise 528 MHz of occupied bandwidth. NOTE 7: The estimated channel impulse response used by the test receiver is limited to half of length of smallest transmit cyclic prefix. NOTE 8: A single subcarrier in the OFDM channel can be excluded from this requirement, no windowing is applied and minimum CP is selected.	

7.5.9.2 Power per Channel for CMTS

A CMTS shall generate an RF output with power capabilities as defined in table 7.37.

The CMTS shall be capable of adjusting channel RF power on a per channel basis as stated in table 7.37.

If the CMTS has independent modulation capability on a per channel basis for legacy DOCSIS channels, then the CMTS shall be capable of adjusting power on a per channel basis for the legacy DOCSIS channels, with each channel independently meeting the power capabilities defined in table 7.37.

Table 7.37: CMTS Output Power

$\mathbf{for\ N^* \equiv \begin{cases} \boldsymbol{minimum}[4N_{eq}{}', \mathbf{ceiling}\left[\frac{N_{eq}}{4}\right]], & N_{eq}{}' < N_{eq}/4 \\ N_{eq}{}', & N_{eq}{}' \geq N_{eq}/4 \end{cases}}$**, **Adjusted Number of Active Channels Combined per RF Port**	
Parameter	**Value**
Required power per channel for N_{eq}' channels combined onto a single RF port:	Required power in dBmV per channel 60 - ceil [3,6*$\log_2$(N*)] dBmV
Range of commanded transmit power per channel	≥ 8 dB below required power level specified below maintaining full fidelity over the 8 dB range
Range of commanded power per channel; adjusted on a per channel basis	CMTS shall: 0 dBc to -2 dBc relative to the highest commanded transmit power per channel, within an 8 dB absolute window below the highest commanded power. CMTS may: *required power* (in table below) to *required power - 8 dB*, independently on each channel.
Commanded power per channel step size	≤ 0,2 dB Strictly monotonic
Power difference between any two adjacent channels in the 108 - 1 218 MHz downstream spectrum (with commanded power difference removed if channel power is independently adjustable and/or accounting for pilot density variation and subcarrier exclusions)	≤ 0,5 dB
Power difference between any two non-adjacent channels in a 48 MHz contiguous bandwidth block (with commanded power difference removed if channel power is independently adjustable)	≤ 1 dB
Power difference (normalized for bandwidth) between any two channels OFDM channel blocks or legacy DOCSIS channels in the 108 - 1 218 MHz downstream spectrum (with commanded power difference removed if channel power is independently adjustable)	≤ 2 dB

for $\mathbf{N}^* \equiv \begin{cases} \boldsymbol{minimum}[\mathbf{4N_{eq}}', \mathbf{ceiling}\left[\frac{\mathbf{N_{eq}}}{\mathbf{4}}\right]], & \mathbf{N_{eq}}' < \mathbf{N_{eq}/4} \\ \mathbf{N_{eq}}', & \mathbf{N_{eq}}' \geq \mathbf{N_{eq}/4} \end{cases}$**, Adjusted Number of Active Channels Combined per RF Port**

Parameter	Value
Power per channel absolute accuracy	±2 dB
Diagnostic carrier suppression (3 modes) Mode 1: One channel suppressed Mode 2: All channels suppressed except one Mode 3: All channels suppressed	1) ≥ 50 dB carrier suppression within the occupied bandwidth in any one active channel. When suppressing the carrier ≥ 50 dB within the occupied bandwidth in any one active channel the CMTS shall control transmissions such that no service impacting discontinuity or detriment to the unsuppressed channels occurs. 2) 50 dB carrier suppression within the occupied bandwidth in every active channel except one. The suppression is not required to be glitchless, and the remaining unsuppressed active channel is allowed to operate with increased power such as the total power of the N' active channels combined. 3) 50 dB carrier suppression within the occupied bandwidth in every active channel. The CMTS shall control transmissions such that in all three diagnostic carrier suppression modes the output return loss of the suppressed channel(s) complies with the Output Return Loss requirements for active channels given in table 7.39. The total noise and spur requirement is the combination of noise power from the 50 dBc suppressed channel and the normal noise and spur requirement for the CMTS output when operating with all channels unsuppressed.
RF output port muting	≥ 73 dB below the unmuted aggregate power of the RF modulated signal, in every 6 MHz CEA channel from 54 MHz to 1 218 MHz. The specified limit applies with all active channels commanded to the same transmit power level. Commanding a reduction in the transmit level of any, or all but one, of the active channels does not change the specified limit for measured muted power in 6 MHz. When the CMTS is configured to mute an RF output port, the CMTS shall control transmissions such that the output return loss of the output port of the muted device complies with the Output Return Loss requirements for inactive channels given in table 7.39.

7.5.9.3 Out-of-Band Noise and Spurious Requirements for CMTS

One of the goals of the DOCSIS DRFI specification is to provide the minimum intended analog channel CNR protection of 60 dB for systems deploying up to 119 DRFI-compliant QAM channels. A new DOCSIS 3.1 PHY goal is to provide protection for legacy DOCSIS channels and for high density constellations of OFDM channel blocks if they are generated from another DRFI-compliant device.

The specification assumes that the transmitted power level of the digital channels will be 6 dB below the peak envelope power of the visual signal of analog channels, which is the typical condition for 256-QAM transmission. It is further assumed that the channel lineup will place analog channels at lower frequencies than digital channels, and in systems deploying DOCSIS 3.1 modulators, analog channels will be placed at centre frequencies below 600 MHz. An adjustment of $10*\log_{10} (6\ MHz / 4\ MHz)$ is used to account for the difference in a 6 MHz equivalent digital channel, versus an analog channel's noise power bandwidth. With the assumptions above, for a 119-6 MHz equivalent channel system, the specification in item 5 of table 7.38 equates to an analog CNR protection of 60 dB.

With all digital channels at the same equivalent power per 6 MHz channel, the specification provides for 58 dB SNR protection for analog channels below 600 MHz (even with transmissions above 600 MHz) with 192 MHz occupied bandwidth (one full OFDM channel) and 51 dB SNR protection for digital channels below 600 MHz with transmission of 960 MHz modulated spectrum (160 equivalent legacy DOCSIS channels). The SNR protection between 600 MHz and 1 002 MHz is 55 dB for analog channels operating above a 192 MHz occupied bandwidth generated by a DOCSIS 3.1 compliant device, and is 48 dB between 600 MHz and 1 002 MHz for digital channels operating above 960 MHz occupied bandwidth generated by a DOCSIS 3.1 compliant device.

Table 7.38 lists the out-of-band spurious requirements. In cases where the N' combined channels are not commanded to the same power level, "dBc" denotes decibels relative to the strongest channel among the active channels. When commanded to the same power level, "dBc" should be interpreted as the average channel power, averaged over the active channels, to mitigate the variation in channel power across the active channels (see table 7.37), which is allowed with all channels commanded to the same power.

The CMTS modulator shall satisfy the out-of-band spurious emissions requirements of table 7.38 in measurements below 600 MHz and outside the encompassed spectrum when the active channels are contiguous or when the ratio of modulated spectrum to gap spectrum within the encompassed spectrum is 4:1 or greater.

The CMTS modulator shall satisfy the out-of-band spurious emissions requirements of table 7.38, with 1 dB relaxation, in measurements within gaps in modulated spectrum below 600 MHz and within the encompassed spectrum when the ratio of modulated spectrum to gap spectrum within the encompassed spectrum is 4:1 or greater.

The CMTS modulator shall satisfy the out-of-band spurious emissions requirements of table 7.38, with 3 dB relaxation, when the ratio of modulated spectrum to gap spectrum within the encompassed spectrum is 4:1 or greater, in measurements with 603 MHz ≤ centre frequency ≤ 999 MHz, outside the encompassed spectrum or in gap channels within the encompassed spectrum.

The CMTS modulator shall satisfy the out-of-band spurious emissions requirements of table 7.38, with 5 dB relaxation, when the ratio of modulated spectrum to gap spectrum within the encompassed spectrum is 4:1 or greater, in measurements with 999 MHz < centre frequency ≤ 1 215 MHz, outside the encompassed spectrum or in gap channels within the encompassed spectrum.

The CMTS modulator shall satisfy the out-of-band spurious emissions requirements of table 7.38, in addition to contributions from theoretical transmit windowing, with permissible configurations of lower edge and upper edge subband exclusions of at least 1 MHz each, FFT Size, cyclic prefix length (N_{cp}) and windowing roll-off period (N_{rp}) values. Recommendations for configuration parameters are provided in Annex I. The test limit for determining compliance to the spurious emissions requirements is the power sum of the spurious emissions requirements taken in accordance with table 7.38; and the contributions from the theoretical transmit windowing for the configured transmissions.

When the N_{eq}' combined active channels are not contiguous, and the ratio of modulated spectrum to gap spectrum within the encompassed spectrum is 4:1 or greater, the spurious emissions requirements are determined by summing the spurious emissions power allowed in a given measurement bandwidth by each of the contiguous sub-blocks among the occupied bandwidth. In the gap channels within the encompassed spectrum and below 600 MHz there is a 1 dB relaxation in the spurious emissions requirements, so that within the encompassed spectrum the spurious emissions requirements (in absolute power) are 26 % higher power in the measurement band determined by the summing of the contiguous sub-blocks' spurious emissions requirements. In all channels above 600 MHz there is a 3 dB relaxation in the spurious emissions requirements, so that the spurious emissions requirements (in absolute power) are double the power in the measurement band determined by the summing of the contiguous sub-blocks' spurious emissions requirements. The following three paragraphs provide the details of the spurious emissions requirements for non-contiguous channel operation outside the encompassed spectrum; within the encompassed spectrum the same details apply except there is an additional 1 dB allowance below 600 MHz; and 3 dB allowance is applied above 600 MHz for all channels.

The full set of N_{eq}' channels is referred to throughout the present document as the modulated channels or the active channels. However, for purposes of determining the spurious emissions requirements for non-contiguous transmitted channels, each separate contiguous sub-block of channels within the active channels is identified, and the number of channels in each contiguous sub-block is denoted as N_{eqi}, for i = 1 to K, where K is the number of contiguous sub-blocks. Therefore, $N'_{eq} = \sum_{i=1}^{k} N_{eq,i}$. Note that K = 1 when and only when the entire set of active channels is contiguous. Also note that an isolated transmit channel, i.e. a transmit channel with empty adjacent channels, is described by $N_i = 1$ and constitutes a sub-block of one contiguous channel. Any number of the "contiguous sub-blocks" may have such an isolated transmit channel; if each active channel was an isolated channel, then K = N'.

When $N_{eq}' \geq N_{eq}/4$, table 7.38 is used for determining the noise and spurious power requirements for each contiguous sub-block, even if the sub-block contains fewer than $N_{eq}/4$ active channels. When $N_{eq}' < N_{eq}/4$, table 7.38 is used for determining the noise and spurious power requirements for each contiguous sub-block. Thus, the noise and spurious power requirements for all contiguous sub-blocks of transmitted channels are determined from table 7.38, where the applicable table is determined by N_{eq}' being greater than or equal to $N_{eq}/4$, or not. The noise and spurious power requirements for the ith contiguous sub-block of transmitted channels is determined from table 7.38 using the value N_i for the "number of active channels combined per RF port", and using "dBc" relative to the highest commanded power level of a 6 MHz equivalent channel among all the active channels, and not just the highest commanded power level in the ith contiguous sub-block, in cases where the N_{eq}' combined channels are not commanded to the same power. The noise and spurious emissions power in each measurement band, including harmonics, from all K contiguous sub-blocks, is summed (absolute power, NOT in dB) to determine the composite noise floor for the non-contiguous channel transmission condition.

For the measurement channels adjacent to a contiguous sub-block of channels, the spurious emissions requirements from the non-adjacent sub-blocks are divided on an equal "per Hz" basis for the narrow and wide adjacent measurement bands. For a measurement channel wedged between two contiguous sub-blocks, adjacent to each, the measurement channel is divided into three measurement bands, one wider in the middle and two narrower bands each abutting one of the adjacent transmit channels. The wideband spurious and noise requirement is split into two parts, on an equal "per Hz" basis, to generate the allowed contribution of power to the middle band and to the farthest narrowband. The ceiling function is applied to the resulting sum of noise and spurious emissions, per note 1 of table 7.38 to produce a requirement of ½ dB resolution.

Items 1 through 4 list the requirements in channels adjacent to the commanded channels.

Item 5 lists the requirements in all other channels further from the commanded channels. Some of these "other" channels are allowed to be excluded from meeting the Item 5 specification. All the exclusions, such as 2nd and 3rd harmonics of the commanded channel, are fully identified in the table.

Item 6 lists the requirements on the $2N'$ 2nd harmonic channels and the $3N'$ 3rd harmonic channels.

Table 7.38: CMTS Output Out-of-Band Noise and Spurious Emissions Requirements

<table>
<tr><td colspan="2">for $\mathbf{N}^* \equiv \begin{Bmatrix} \boldsymbol{minimum}[4\mathbf{N_{eq}}', \mathbf{ceiling}\left[\frac{}{4}\right]], & \mathbf{N_{eq}}' < \mathbf{N_{eq}}/4 \\ \mathbf{N_{eq}}', & \mathbf{N_{eq}}' \geq \mathbf{N_{eq}}/4 \end{Bmatrix}$</td><td>Adjusted Number of Active Channels Combined per RF Port</td></tr>
<tr><td></td><td>Band</td><td>Requirement (in dBc)</td></tr>
<tr><td>1</td><td>Adjacent channel up to 750 kHz from channel block edge</td><td>For N* = 1, 2, 3, 4: ≤ -58;

For N* ≥ 5: ≤ 10*$\log_{10}[10^{-58/10} + (0,75/6)*(10^{-65/10} + (N^*-2)*10^{-73/10})]$</td></tr>
<tr><td>2</td><td>Adjacent channel (750 kHz from channel block edge to 6 MHz from channel block edge)</td><td>For N* = 1: ≤ -62;

For N* ≥ 2: ≤ 10*$\log_{10}[10^{-62/10} + (5,25/6)*(10^{-65/10} + (N^*-2)*10^{-73/10})]$</td></tr>
<tr><td>3</td><td>Next-adjacent channel (6 MHz from channel block edge to 12 MHz from channel block edge)</td><td>≤ 10*$\log_{10}[10^{-65/10} + (N^*-1)*10^{-73/10}]$</td></tr>
<tr><td>4</td><td>Third-adjacent channel (12 MHz from channel block edge to 18 MHz from channel block edge)</td><td>For N* = 1: ≤ -73;
For N* = 2: ≤ -70;
For N* = 3: ≤ -67;
For N* = 4: ≤ -65;
For N* = 5: ≤ -64,5;
For N* = 6, 7: ≤ -64;

For N* ≥ 8: ≤ -73 + 10*$\log_{10}(N^*)$</td></tr>
<tr><td>5</td><td>Noise in other channels (47 MHz to 1 218 MHz) Measured in each 6 MHz channel excluding the following:
a) Desired channel(s)
b) 1st, 2nd, and 3rd adjacent channels (see Items 1, 2, 3, 4 in this table)
c) Channels coinciding with 2nd and 3rd harmonics (see Item 6 in this table)</td><td>For N* = 1: ≤ -73;
For N* = 2: ≤ -70;
For N* = 3: ≤ -68;
For N* = 4: ≤ -67;

For N* ≥ 5: ≤ -73 + 10*$\log_{10}(N^*)$</td></tr>
<tr><td>6</td><td>In each of 2N' contiguous 6 MHz channels or in each of 3N' contiguous 6 MHz channels coinciding with 2nd harmonic and with 3rd harmonic components respectively (up to 1 218 MHz)</td><td>≤ -73 + 10*$\log_{10}(N^*)$ dBc, or -63, whichever is greater</td></tr>
<tr><td>7</td><td>Lower out of band noise in the band of 5 MHz to 47 MHz
Measured in 6 MHz channel bandwidth</td><td>≤ -50 + 10*$\log_{10}(N^*)$</td></tr>
<tr><td>8</td><td>Higher out of band noise in the band of 1 218 MHz to 3 000 MHz
Measured in 6 MHz channel bandwidth</td><td>For N* ≤ 8: ≤ -55 + 10*$\log_{10}(N^*)$

For N* > 8: ≤ -60 + 10*$\log_{10}(N^*)$</td></tr>
<tr><td colspan="3">NOTE 1: All equations are Ceiling(Power, 0,5) dBc. Use "Ceiling(2*Power) / 2" to get 0,5 steps from ceiling functions that return only integer values. For example Ceiling(-63,9 and 0,5) = -63,5 dBc.
NOTE 2: Add 3 dB relaxation to the values specified above for noise and spurious emissions requirements in all channels with 603 MHz ≤ centre frequency of the noise measurement ≤ 999 MHz. For example -73 dBc becomes -70 dBc. Add 5 dB relaxation to the values specified above for noise and spurious emissions requirements in all channels with 999 MHz < centre frequency of the noise measurement ≤ 1 215 MHz. For example -73 dBc becomes -68 dBc.
NOTE 3: Add 1 dB relaxation to the values specified above for noise and spurious emissions requirements in gap channels with centre frequency below 600 MHz. For example -73 dBc becomes -72 dBc.</td></tr>
</table>

7.5.10 CMTS Transmitter Output Requirements

7.5.10.0 CMTS Transmitter Output Requirements

CMTSs capable of generating *N*-channels of legacy DOCSIS plus N_{OFDM}-channels of OFDM per RF port, for purposes of the output electrical requirements, the device is said to be capable of generating N_{eq}-channels per RF port, where $N_{eq} = N + 32*N_{OFDM}$ "equivalent legacy DOCSIS channels.

When operating with all N_{eq} channels on the RF port the CMTS shall comply with all requirements for a device having all N_{eq} channels on the RF port as defined in the present document.

When operating with N_{eq}' active channels on the RF port for all values of N_{eq}' less than N_{eq}, THE CMTS shall comply with all requirements for a device having N_{eq}' channels per RF port as defined in the present document.

The number of equivalent legacy DOCSIS channels in the OFDM channel is the ceiling function applied to the modulated spectrum divided by 6 MHz. For the example, the number of equivalent legacy DOCSIS channels in the OFDM channel is ceiling(180,3 MHz / 6 MHz) = 31.

For an N_{eq}-channel per RF port device, the applicable maximum power per channel and spurious emissions requirements are defined using a value of $N^* = \text{minimum}(4N_{eq}', \text{ceiling}[N_{eq}/4])$ for $N_{eq}' < N_{eq}/4$, and $N^* = N_{eq}'$ otherwise.

These specifications assume that the CMTS will be terminated with a 75 ohm load.

7.5.10.1 CMTS Output Electrical Requirements

A CMTS shall output an RF modulated signal with the characteristics defined in table 7.39. The condition for these requirements is the entire OFDM block commanded to a constant power spectral density, except for Phase Noise, Diagnostic Carrier Suppression, and power difference requirements (table 7.37), and except as described for Out-of-Band Noise and Spurious Requirements listed in clause 7.5.9.

When the CMTS generates ≥ 54 MHz (nine or more equivalent channels) on a single RF output port, it shall provide for independent selection of centre frequency when the ratio of active bandwidth to excluded bandwidth in the encompassed spectrum is at least 2:1, and with each channel independently meeting the requirements of table 7.39 except for spurious emissions requirements defined in clause 7.5.9.

When the CMTS generates ≥ 54 MHz (nine or more equivalent channels) on a single RF output port it shall meet the requirements of table 7.39 when the ratio of active bandwidth to excluded bandwidth in the encompassed spectrum is at least 4:1.

Table 7.39: CMTS Output Requirements

Parameter	Value
Lower Band Edge	Refer to clause 7.2
Upper Band Edge	Refer to clause 7.2
Signal Type	OFDM
Single FFT Block Bandwidth	192 MHz
Minimum Active Signal Bandwidth	24 MHz
Subcarrier Spacing / OFDM Symbol Rate FFT duration	25 kHz / 40 µs 50 kHz / 20 µs
FFT Size	50 kHz: 4 096 (4K FFT) 25 kHz: 8 192 (8K FFT)
Maximum Number of Subcarriers per FFT	4K: 3 800 8K: 7 600
Number of Data Subcarriers per FFT	4K: 3 800 - number of pilot tones 8K: 7 600 - number of pilot tones
Continuous Pilot Tones	Continuous pilot placement is defined in clause 7.5.15.1.2. Minimum number of continuous pilots is 16 and the maximum number is 128. Locations of 8 continuous pilots are pre-defined with reference to the PLC location. Locations of remaining continuous pilots are defined using PLC messages.
Scattered Pilot Tones	4K FFT: One out of every 128 subcarriers, staggered by 1 8K FFT: One out of every 128 subcarriers, staggered by 2
Cyclic Prefix Options	Samples µs 192 0,9375 256 1,25 512 2,5 768 3,75 1 024 5,0 ** sampling rate of 204,8 MHz*
OFDM Shaping Windowing Options	Raised cosine (Tukey) absorbed by CP Samples µs

Parameter	Value
	0 0 64 0,3125 128 0,625 192 0,9375 256 1,25 * *sampling rate of 204,8 MHz*
Level	Adjustable, see table 7.37
Inband Spurious, Distortion, and Noise Inband Spurious and Noise Out of Band Spurious and Noise	See table 7.36 See table 7.36 See table 7.38
In band Phase Noise Allowable Degradation	See table 7.36 1,5 dB
Output Impedance	See table 7.36
Output Return Loss	See table 7.36
Connector	F connector per [19]

7.5.10.2 Independence of Individual Channel Within the Multiple Channels on a Single RF Port

A potential use of a CMTS is to provide a universal platform that can be used for high-speed data services or for video services. For this reason, it is essential that interleaver depth be set on a per channel basis to provide a suitable transmission format for either video or data as needed in normal operation. Any *N*-channel block of a CMTS shall be configurable with at least two different interleaver depths, using any of the interleaver depths defined in clause 7.5.6.1. Although not as critical as per-channel interleaver depth control, there are strong benefits for the operator if the CMTS is provided with the ability to set RF power, centre frequency, and modulation type on a per-channel basis:

1) A multiple-channel CMTS shall be configurable with at least two different legacy interleaver depths among the legacy channels on an RF output port in addition to each OFDM channel which is configurable independently.

2) A multiple-channel CMTS shall provide for 3 modes of carrier suppression of RF power for diagnostic and test purposes, see table 7.37 for mode descriptions and carrier RF power suppression level.

3) A multiple-channel CMTS may provide for independent adjustment of RF power in a per channel basis for legacy channels with each RF carrier independently meeting the requirements defined in table 7.37.

4) A multiple-channel CMTS may provide for independent selection of centre frequency on a per channel basis, thus providing for non-contiguous channel frequency assignment with each channel independently meeting the requirements in table 7.39. A multiple-channel CMTS capable of generating nine or more channels on a single RF output port shall provide for independent selection of centre frequency with the ratio of number of active channels to gap channels in the encompassed spectrum being at least 2:1, and with each channel independently meeting the requirements in table 7.39 except for spurious emissions. A multiple-channel CMTS capable of generating nine or more channels on a single RF output port shall meet the requirements of table 7.39 when the ratio of number of active channels to gap channels in the encompassed spectrum is at least 4:1. (A ratio of number of active channels to gap channels of at least 4:1 provides that at least 80 % of the encompassed spectrum contains active channels, and the number of gap channels is at most 20 % of the encompassed spectrum.)

5) A multiple-channel CMTS may provide for independent selection of modulation order, either 64-QAM or 256-QAM, on a per channel basis for legacy channels, with each channel independently meeting the requirements in table 7.39.

6) A CMTS shall provide a test mode of operation, for out-of-service testing, configured for N channels but generating one-CW-per-channel, one channel at a time at the centre frequency of the selected channel; all other combined channels are suppressed. One purpose for this test mode is to support one method for testing the phase noise requirements of table 7.39. As such, the CMTS generation of the CW test tone should exercise the signal generation chain to the fullest extent practicable, in such manner as to exhibit phase noise characteristics typical of actual operational performance; for example, repeated selection of a constellation symbol with power close to the constellation RMS level would seemingly exercise much of the modulation and up-conversion chain in a realistic manner. The CMTS test mode shall be capable of generating the CW tone over the full range of Center Frequency in table 7.39.

7) A CMTS shall provide a test mode of operation, for out-of-service testing, generating one-CW-per-channel, at the centre frequency of the selected channel, with all other N - 1 of the combined channels active and containing valid data modulation at operational power levels. One purpose for this test mode is to support one method for testing the phase noise requirements of table 7.39. As such, the generation of the CW test tone should exercise the signal generation chain to the fullest extent practicable, in such manner as to exhibit phase noise characteristics typical of actual operational performance. For example, a repeated selection of a constellation symbol, with power close to the constellation RMS level, would seemingly exercise much of the modulation and upconversion chain in a realistic manner. For this test mode, it is acceptable that all channels operate at the same average power, including each of the N - 1 channels in valid operation, and the single channel with a CW tone at its centre frequency. When operating in one-CW-per-channel test mode the CMTS shall be capable of generating the CW tone over the full range of Center Frequency in table 7.39.

8) A CMTS shall be capable of glitchless reconfiguration over a range of active channels from ceiling[7*N_{eq}/8] to N_{eq}. Channels which are undergoing configuration changes are referred to as the "changed channels." The channels which are active and are not being reconfigured are referred to as the "continuous channels".

   Glitchless reconfiguration consists of any of the following actions while introducing no discontinuity or detriment to the continuous channels, where the modulator is operating in a valid DOCSIS 3.1-required mode both before and after the reconfiguration with an active number of channels staying in the range {ceiling[7*N_{eq}/8], N_{eq}}: adding and/or deleting one or more channels, and/or moving some channels to new RF carrier frequencies, and/or changing the interleaver depth, modulation, power level, or frequency on one or more channels. Any change in the modulation characteristics (power level, modulation density, interleaver parameters, centre frequency) of a channel excuses that channel from being required to operate in a glitchless manner. For example, changing the power per channel of a given channel means that channel is not considered a continuous channel for the purposes of the glitchless modulation requirements. Glitchless operation is not required when N_{eq} is changed, even if no reconfigurations accompany the change in N_{eq}.

7.5.11 Cable Modem Receiver Input Requirements

The CM shall be able to accept any range of OFDM subcarriers defined between Lower Frequency Boundary and Upper Frequency Boundary simultaneously. Active subcarrier frequencies, loading, and other OFDM characteristics are described by OFDM configuration settings and CM exclusion bands and profile definition. The OFDM signals and CM interfaces will have the characteristics and limitations defined in table 7.40.

The CM shall support the requirements in table 7.40 unless otherwise noted.

Table 7.40: Electrical Input to CM

Parameter	Value
Lower Frequency Boundary	258 MHz should support 108 MHz Note: applies if f_{umax} is 85 MHz or less
Upper Frequency Boundary	1 218 MHz should support 1 794 MHz
Frequency Boundary Assignment Granularity	25 kHz 8K FFT 50 kHz 4K FFT
Signal Type	OFDM
Single FFT Block Bandwidth	192 MHz
Minimum Contiguous-Modulated OFDM Bandwidth	24 MHz
Number of FFT Blocks	Support minimum of 2 FFT Blocks AND 24 SC-QAM Channels
Subcarrier Spacing/FFT Duration	25 kHz / 40 µs 50 kHz / 20 µs
Modulation Type	QPSK, 16-QAM, 64-QAM, 128-QAM, 256-QAM, 512-QAM, 1024-QAM, 2048-QAM, 4096-QAM may support 8192-QAM, 16384-QAM
Variable Bit Loading	Support with subcarrier granularity Support zero bit loaded subcarriers
Total Input Power	< 40 dBmV, 54 MHz - 1,794 GHz ** Assuming negligible power outside this range*
Level Range (24 MHz min occupied BW) Equivalent PSD to SC-QAM of -15 dBmV to + 15 dBmV per 6 MHz	-9 dBmV/24 MHz to 21 dBmV/24 MHz
Maximum average power of any 24 MHz input to the CM from 54 MHz to 1 218 MHz OR From 258 MHz to 1,794 GHz	Let **X** = Average power of lowest power 24 MHz BW for demodulation Additional Demodulated Bandwidth, B_{remond} : ≤ Min [**X** + 10 + 10*log**(**B_{remond}/24) ; 21**+** 10*log**(**B_{remond}/24)] Additional Non-Demodulated Bandwidth, $B_{no\text{-}demod}$: ≤ Min [**X** + 10 + 10*log**(**$B_{no\text{-}demod}$/24) ; 26**+** 10*log**(**$B_{no\text{-}demod}$/24)] For up to 12 MHz of occupied bandwidth (analog, OOB, QAM, OFDM) ≤ Min [**X** + 10 + 10*log**(**$B_{no\text{-}demod}$/24) ; 21**+** 10*log**(**$B_{no\text{-}demod}$/24)] for all remaining bandwidth Level range does not imply anything about BER performance or capability vs. QAM. CM BER performance is separately described.
Input Impedance	75 ohms
Input Return Loss	> 6 dB (258 MHz - 1 218 MHz) > 6 dB (108 MHz - 1 218 MHz) (see note 1) > 6 dB (258 MHz - 1,794 GHz) (see note 2)
Connector	F connector per [17] or [19]
NOTE 1: Applies when lower frequency boundary is 108 MHz. NOTE 2: Applies when upper frequency boundary is 1,794 GHz.	

7.5.12 Cable Modem Receiver

7.5.12.0 Cable Modem Receiver Capabilities

The required level for CM downstream post-FEC error ratio is defined as less than or equal to 10^{-6} PER (packet error ratio) with 1 500 byte Ethernet packets. This clause describes the conditions at which the CM is required to meet this error ratio.

7.5.12.1 CM Error Ratio Performance in AWGN Channel

Implementation loss of the CM shall be such that the CM achieves the required error ratio when operating at a CNR as shown in table 7.41, under input load and channel conditions as follows:

- Any valid transmit combination (frequency, subcarrier clock frequency, transmit window, cyclic prefix, pilot, PLC, subcarrier exclusions, interleaving depth, multiple modulation profile configuration, etc.) as defined in the present specification.
- P6AVG (the measured channel power divided by number of occupied CEA channels) ≤ 15 dBmV.
- Up to fully loaded spectrum of 54 - 1 218 MHz, including up to 48 analog channels placed lower in the spectrum than the digital channels.
- Power in (both above and below) 4 adjacent 6 MHz channels ≤ P6AVG+3 dB.
- Power in any 6 MHz channel over the spectrum ≤ P6AVG+6 dB.
- Peak envelope power in any analog channel over the spectrum ≤ P6AVG+6 dB.
- Average power per channel across spectrum ≤ P6AVG+3 dB.
- OFDM channel phase noise as in the CMTS specification.
- No other artifacts (reflections, burst noise, tilt, etc.).

Table 7.41: CM Minimum CNR Performance in AWGN Channel

Constellation	CNR[NOTE 1, NOTE 2] (dB) Up to 1 GHz	CNR[1,2] (dB) 1 GHz to 1,2 GHz	Min P_{6AVG} dBmV
4 096	41,0	41,5	-6
2 048	37,0	37,5	-9
1 024	34,0	34,0	-12
512	30,5	30,5	-12
256	27,0	27,0	-15
128	24,0	24,0	-15
64	21,0	21,0	-15
16	15,0	15,0	-15
NOTE 1: CNR is defined here as total signal power in occupied bandwidth divided by total noise in occupied bandwidth. NOTE 2: Channel CNR is adjusted to the required level by measuring the source inband noise including phase noise component and adding the required delta noise from an external AWGN generator. NOTE 3: Applicable to an OFDM channel with 192 MHz of occupied bandwidth.			

7.5.13 Physical Layer Link Channel (PLC)

7.5.13.0 PLC Overview

This clause contains the description of the Physical layer Link Channel (PLC) that the CMTS follows during the construction of the PLC.

The aim of the PLC is for the CMTS to convey to the CM the physical properties of the OFDM channel. In a blind acquisition, that is, in an acquisition without prior knowledge of the physical parameters of the channel, the CM first acquires the PLC, and from this extract the parameters needed to acquire the complete OFDM channel.

7.5.13.1 PLC Placement

The CMTS shall transmit a PLC for every downstream OFDM channel.

The CMTS shall place the PLC at the centre of a 6 MHz encompassed spectrum with no excluded subcarriers. For 4K FFT OFDM, this 6 MHz will contain 56 subcarriers followed by the 8 PLC subcarriers followed by another 56 subcarriers. For 8K FFT OFDM, this 6 MHz will contain 112 subcarriers followed by the 16 PLC subcarriers followed by another 112 subcarriers.

The CMTS shall place the 6 MHz encompassed spectrum containing the PLC on a 1 MHz grid, that is, the centre frequency of the lowest frequency subcarrier of the 6 MHz encompassed spectrum containing the PLC has to be an integer when the frequency is measured in units of MHz.

7.5.13.2 PLC Structure

The CMTS shall place the PLC so that it occupies the same set of contiguous subcarriers in every OFDM symbol.

The CMTS shall place 8 OFDM subcarriers in the PLC of every OFDM symbol when using 4K FFT OFDM (i.e. a subcarrier spacing of 50 kHz).

The CMTS shall place 16 OFDM subcarriers in the PLC of every OFDM symbol when using 8K FFT OFDM (i.e. a subcarrier spacing of 25 kHz).

Table 7.42: PLC Components

DFT size	Subcarrier spacing	Number of PLC subcarriers (N_p)
4 096	50 kHz	8
8 192	25 kHz	16

The CMTS shall use a 16-QAM constellation for the PLC subcarriers.

The CMTS shall construct the PLC as 8 symbols of preamble followed by 120 symbols of data subcarriers, as shown in figure 7.55.

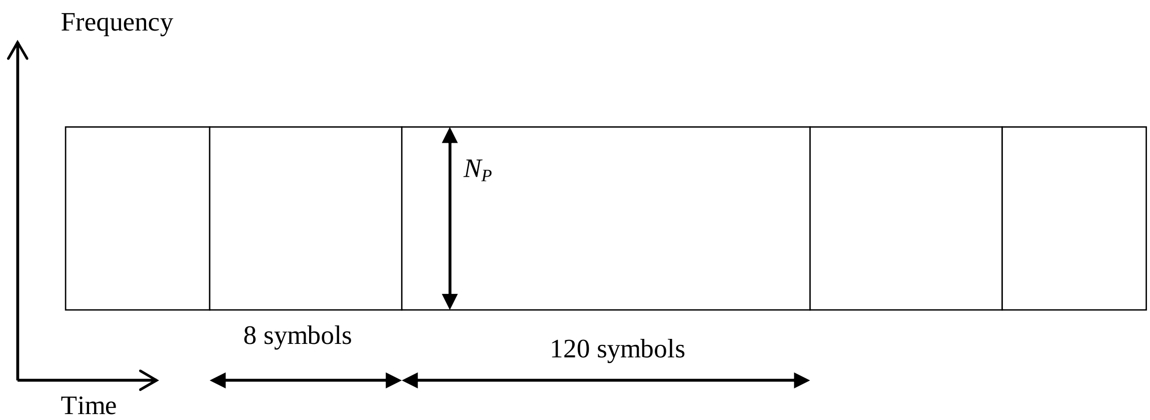

Figure 7.55: Structure of the PLC

The CMTS shall place the PLC at the centre of a 6 MHz of active frequency range. For 4K FFT OFDM, this 6 MHz channel, in the increasing order of frequency, will consist of 56 subcarriers followed by the 8 PLC subcarriers followed by another 56 subcarriers. For 8K FFT OFDM, this 6 MHz channel, in the increasing order of frequency, will consist of 112 subcarriers followed by the 16 PLC subcarriers followed by another 112 subcarriers.

The CMTS shall not insert any exclusion zones or excluded subcarriers within this 6 MHz band that contains the PLC.

The CMTS shall insert 8 continuous pilots in this 6 MHz bandwidth, 4 on each side of the PLC, as defined in the section on downstream pilots.

The CMTS shall interleave the PLC subcarriers on their own, as described in the section on "PLC Interleaving".

The CMTS shall not interleave the PLC preamble.

The CMTS shall synchronize the scattered pilot pattern to the PLC preamble as defined in clause 7.5.15 such that in the OFDM symbol that follows the last symbol of the preamble sequence, the subcarrier next to the highest-frequency subcarrier in the PLC is a scattered pilot.

The CMTS shall not insert any scattered pilots or continuous pilots within the PLC frequency band.

The CMTS shall synchronize the downstream data randomizer to the PLC preamble as described in clause 7.5.5.3. That is, the CMTS will initialize the downstream randomizer just before the lowest frequency data subcarrier of the first OFDM symbol following the preamble.

The CMTS shall synchronize the downstream PLC randomizer to the PLC preamble as described in clause 7.5.13.8. That is, the CMTS shall initialize the downstream PLC randomizer just before the lowest frequency PLC subcarrier of the first OFDM symbol following the preamble.

The CMTS shall place the 6 MHz bandwidth containing the PLC within the active bandwidth of the OFDM channel.

Two possible locations for the PLC channel are illustrated in the example of figure 7.56. In this example there are three contiguous OFDM spectral bands in the 192 MHz channel, of width 22, 12 and 5 MHz. There are two exclusion zones between these. The spectrum outside these bands is also excluded.

It is not necessary to place the PLC in the largest contiguous spectral segment of the OFDM channel. The 6 MHz channel containing the PLC at its centre may be anywhere provided it contains 6 MHz of spectrum without any excluded subcarriers. In the example given the one possible location for the PLC channel is in the 12 MHz wide segment.

Since the downstream channel will contain a minimum of 22 MHz of contiguous OFDM spectrum, there will always be a spectral band to place the PLC. It may be noted that it not necessary to place the PLC at the centre of the 22 MHz bandwidth.

The CMTS is expected to place the PLC in part of the spectrum that is less susceptible to noise and interference.

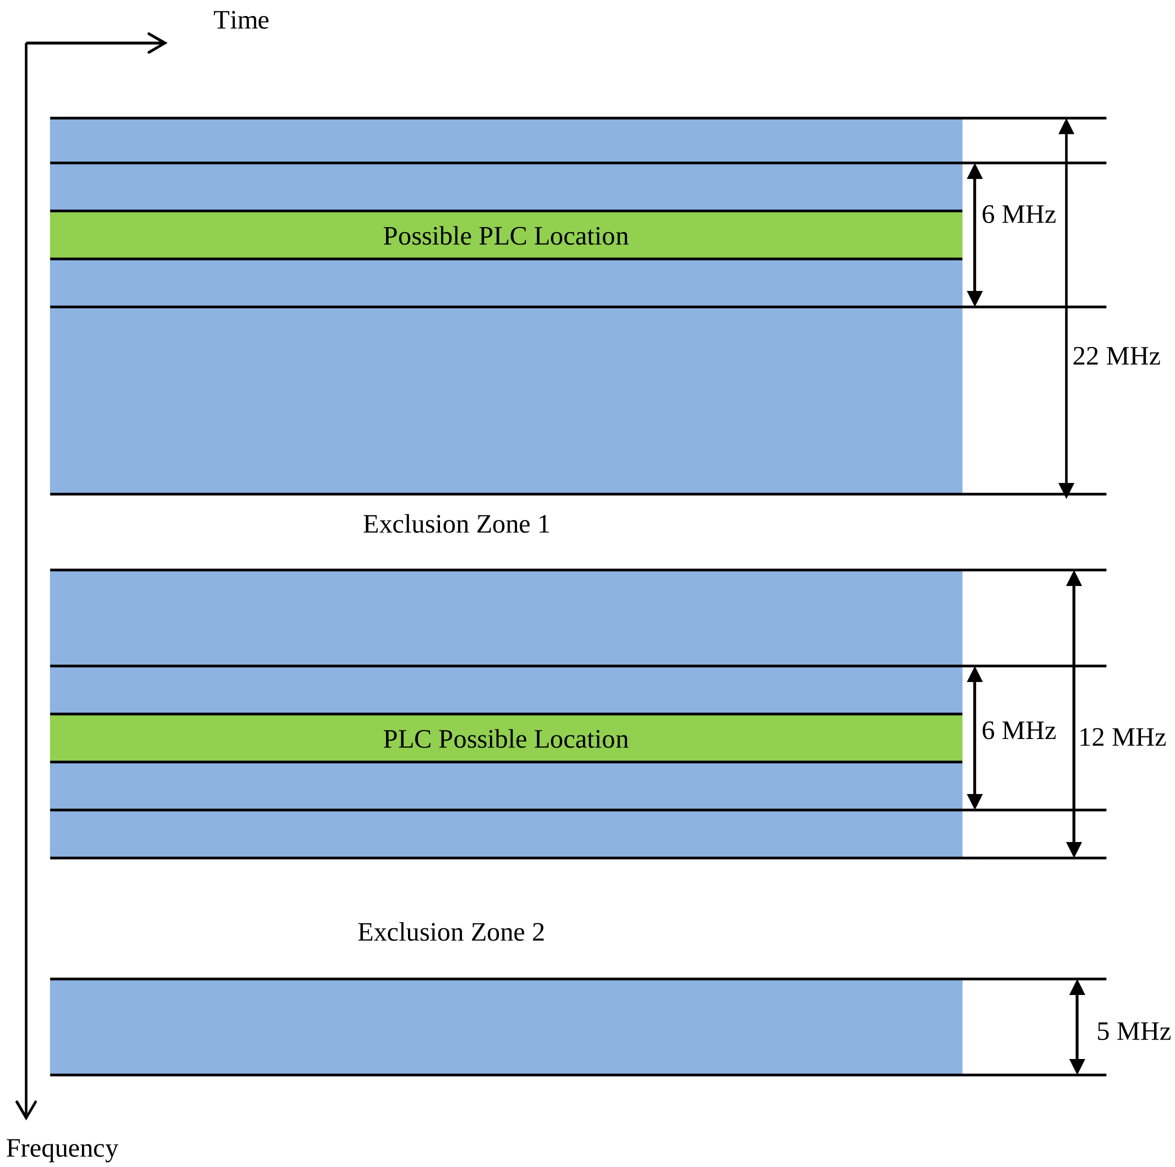

Figure 7.56: Examples of PLC Placement

The CMTS shall generate the PLC as shown in figure 7.57.

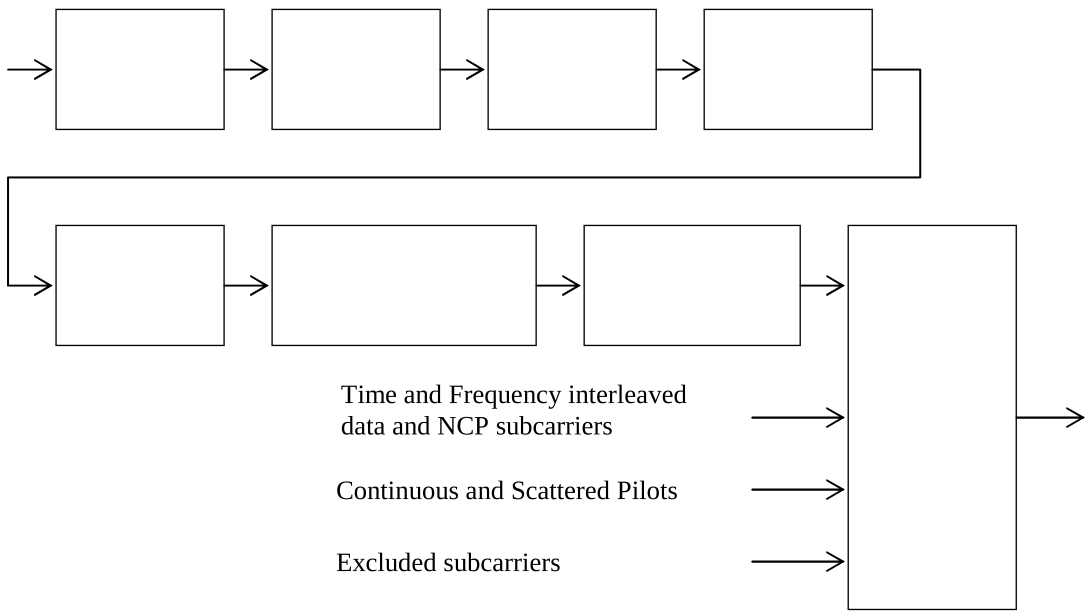

Figure 7.57: Physical Layer Operations for Forming the PLC Subcarriers

7.5.13.3 PLC Preamble Modulation

The CMTS shall modulate the subcarriers in the PLC preamble using binary phase-shift keying (BPSK), as described in this clause.

For 4K FFT, the preamble size is 8 subcarriers. Thus, an array of size (8 x 8) is defined as follows:

Table 7.43: PLC Preamble for 4K FFT

	Symbol 1	Symbol 2	Symbol 3	Symbol 4	Symbol 5	Symbol 6	Symbol 7	Symbol 8
Subcarrier 1	0	0	1	0	1	1	0	1
Subcarrier 2	1	0	0	0	1	1	1	0
Subcarrier 3	0	1	1	1	1	0	0	1
Subcarrier 4	0	1	0	0	0	1	1	0
Subcarrier 5	1	1	1	0	1	1	1	1
Subcarrier 6	1	0	0	0	0	0	0	1
Subcarrier 7	0	1	0	1	0	0	1	1
Subcarrier 8	0	0	1	0	0	0	1	1

Table 7.44: PLC Preamble for 8K FFT

	Symbol 1	Symbol 2	Symbol 3	Symbol 4	Symbol 5	Symbol 6	Symbol 7	Symbol 8
Subcarrier 1	1	0	0	1	0	1	0	0
Subcarrier 2	0	1	1	0	0	1	0	0
Subcarrier 3	0	1	1	1	0	0	0	1
Subcarrier 4	0	0	0	1	0	1	1	1
Subcarrier 5	1	1	0	0	1	0	1	0
Subcarrier 6	0	0	0	1	1	0	0	1
Subcarrier 7	0	1	1	1	0	1	1	0
Subcarrier 8	1	1	1	0	0	0	1	0
Subcarrier 9	0	1	1	1	1	0	0	1
Subcarrier 10	1	1	1	1	0	1	1	1
Subcarrier 11	1	1	1	0	0	0	0	0
Subcarrier 12	1	1	0	1	0	1	0	1
Subcarrier 13	1	1	0	0	1	1	0	0
Subcarrier 14	1	0	1	1	1	0	1	0
Subcarrier 15	0	1	0	1	1	0	0	0
Subcarrier 16	0	0	1	0	0	0	0	1

The CMTS shall map each of the above binary bits to a BPSK constellation point in the complex plane using the following transformation:

0 → (1 + j0)

1 → (-1 + j0)

7.5.13.4 PHY Parameters Carried by the PLC

The PLC carries two sets of PHY parameters from the CMTS to cable modems: the Downstream Profile Descriptor and the OFDM Channel Descriptor. Contents of each of these descriptors are described in [4]. This clause contains only a brief description of the physical layer parameters carried by the PLC. For formatting and other details, reference is made to the MULPI specification.

The inverse discrete Fourier transform that defines the OFDM signal at the CMTS is given by the following equation:

$$x(i) = \frac{1}{\sqrt{N}}\sum_{k=0}^{N-1} X(k)\exp\left(\frac{j2\pi i\left(k-\frac{N}{2}\right)}{N}\right);\ \text{for } i = 0, 1, \ldots, N-1 \qquad (1)$$

The sampling rate in the previous equation is 204,8 Msamples/s and the value of N is either 4 096 or 8 192. The CMTS shall specify this value of N via the PLC.

The CMTS shall define, via the PLC, the frequency of the subcarrier $X(0)$ in equation (1) as a 32-bit positive integer in units of Hz.

The PLC subcarriers constitute a set of contiguous subcarriers given by:

$$\{X(k),\ k = L, L+1, \ldots, L+N_P-1\} \qquad (2)$$

The CMTS shall define the value of L to define the location of the PLC within an OFDM channel.

The CMTS shall define the locations of the continuous pilots, excluding the eight predefined ones, via the PLC.

The CMTS shall define the locations of excluded subcarriers via the PLC.

The CMTS shall define the bit loading profile for all 4 096 or 8 192 subcarriers of equation (1), excluding continuous pilots and excluded subcarriers, via the PLC.

The CMTS shall use the indices k of equation (1) to specify the locations of subcarriers in all of the above definitions.

In addition to above, the CMTS shall define the following physical parameters of the OFDM channel:

- Cyclic prefix length (five possible settings)

- Roll-off (five possible settings)
- Time interleaver depth (any integer from 0 to 32)
- Modulation of the NCP (QPSK, QAM-16 or QAM-64)

7.5.13.5 Mapping of Bytes to a Bit Stream

The CMTS shall convert the stream of bytes received by the PLC into a stream of bits, MSB first, as illustrated in figure 7.58.

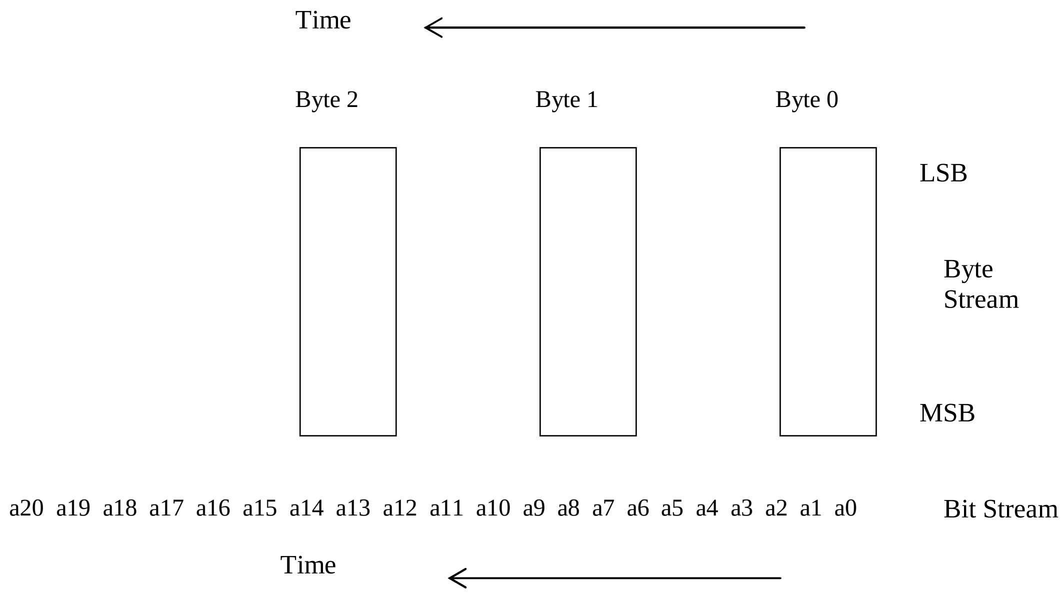

Figure 7.58: Mapping Bytes into a Bit Stream for FEC Encoding

7.5.13.6 Forward Error Correction code for the PLC

The CMTS shall encode the PLC data using (384,288) puncturing LDPC encoder, see clause 7.4.4.3 for the definition of puncturing encoder.

The puncturing encoder uses the same mother encoder for fine ranging FEC (see clause 7.4.16.2.3), that is the rate 3/5 (480,288) LDPC encoder listed by table 7.15.

Denote the information bits sent to the mother code encoder by $(a_0, \cdots, a_{287})$ and let the encoder output being $(a_0, \cdots, a_{287}, b_{288}, \cdots, b_{479})$, where $b_{288}, \cdots, b_{479}$ are parity-check bits. The coordinates to be deleted by the puncturing step are:

- Period 1: 48 consecutive coordinates $a_{48}, \cdots, a_{95}$
- Period 2: 48 consecutive coordinates $b_{384}, \cdots, b_{431}$

NOTE: Also see figure 7.75.

The puncturing is described in figure 7.59.

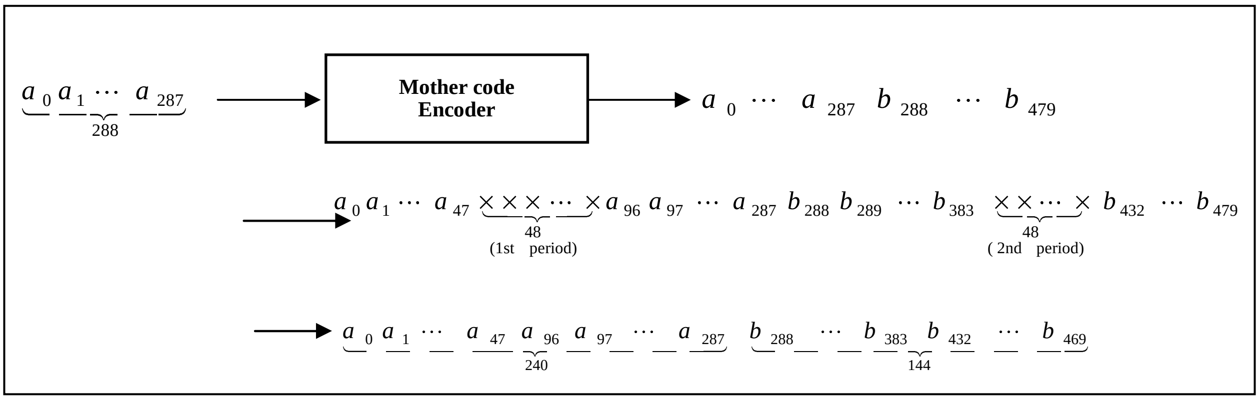

Figure 7.59: Puncturing Encoder for the PLC FEC

7.5.13.7 Block Interleaving of the PLC Subcarriers

The preceding section shows 240 data bits entering the LDPC encoder and 384 encoded bits exiting the LDPC encoder. This sequence is in effect is time-reversed order. The time-ordered sequence takes the form shown in figure 7.60.

The CMTS shall map these 384 data bits into 96 4-bit nibbles $\{u_i,\ i = 0,1, ...\ , 95\}$ as described below using figure 7.60 before interleaving.

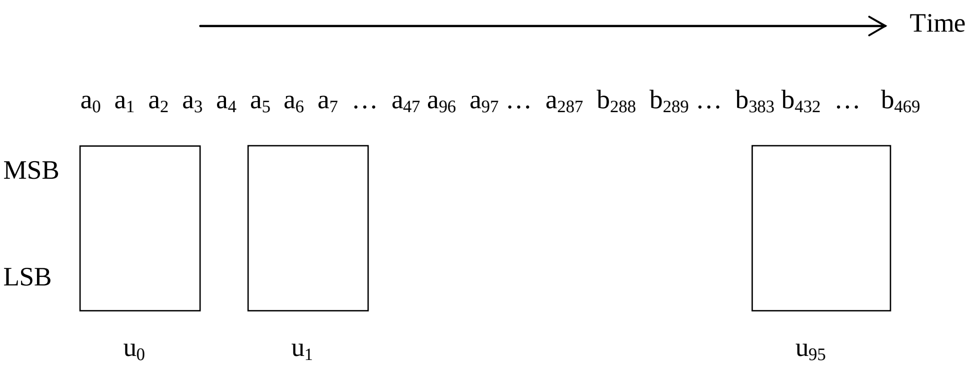

Figure 7.60: Mapping Encoded Bit Stream into a Stream of Nibbles

The CMTS shall interleave this 96-nibble sequence $\{u_0\ u_1\ u_2\ ...\ \ u_{95}\}$ as described below.

For 4K FFT, the CMTS shall use an (8x12) array. The CMTS shall write the values u_i along the rows of this two-dimensional array, as shown in figure 7.61.

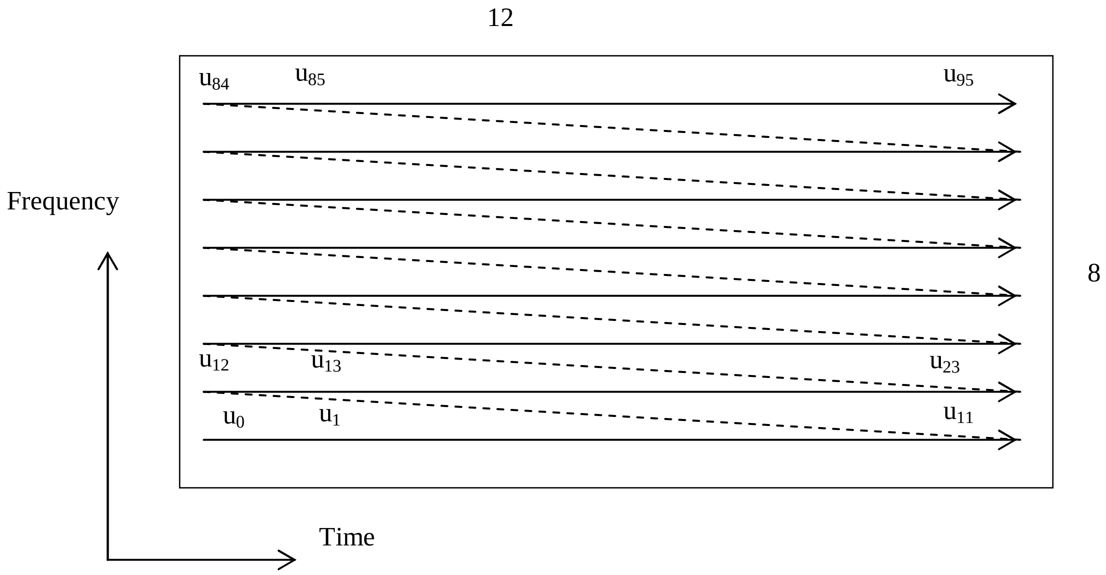

Figure 7.61: Block Interleaving of PLC Subcarriers for 4K FFT

The CMTS shall then read this two-dimensional array along vertical columns to form the two-dimensional sequence $\{v_{t,f}, t = 0, 1, \ldots, 11 \text{ and } f = 0, 1, \ldots, 7\}$. This operation is mathematically represented as:

$$v_{t,f} = u_{t+12f}$$

The CMTS shall map each of the 8-point sequence given below to the 8 successive PLC subcarriers of an OFDM symbol after randomization described in the next section.

$V_t = \{v_{t.f}, \ f = 0, 1, \ldots, 7\}$ for 12 successive OFDM symbols t =0, 1, … , 11

Therefore, each FEC codeword will occupy the PLC segment of twelve successive 4K FFT OFDM symbols. There will be ten such codewords in an 128-symbol PLC frame, including the 8-symbol preamble.

The CMTS shall map ten complete FEC codewords into one 4K FFT PLC frame.

For 8K FFT, the CMTS shall use a (16 x 6) array. The CMTS shall write the values u_i along the rows of this two-dimensional array, as shown figure 7.62.

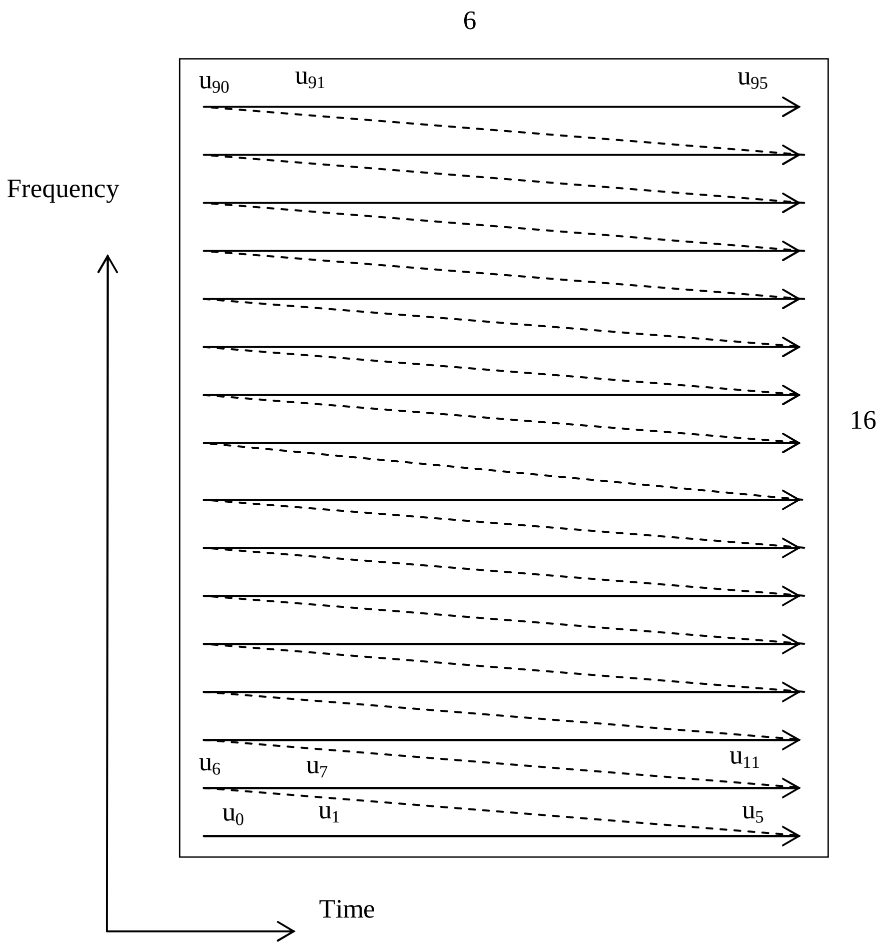

Figure 7.62: Block Interleaving of PLC Subcarriers for 8K FFT

The CMTS shall then read this two-dimensional array along vertical columns to form the two-dimensional sequence $\{v_{t,f}, t = 0, 1, \dots, 5 \text{ and } f = 0, 1, \dots, 15\}$. This operation is mathematically represented as:

$$v_{t,f} = u_{t+6f}$$

The CMTS shall map each of the 16-point sequence given below to the 16 successive PLC subcarriers of an OFDM symbol after randomization described in the next section.

$V_t = \{v_{t,f},\ f = 0, 1, \dots, 15\}$ for 6 successive OFDM symbols t =0, 1, … , 5

Therefore, each FEC codeword will occupy the PLC segment of six successive 8K FFT OFDM symbols. There will be twenty such codewords in a 128-symbol PLC frame, including the 8-symbol preamble.

The CMTS shall map twenty complete FEC codewords into one 8K FFT PLC frame.

7.5.13.8 Randomizing the PLC Subcarriers

The CMTS shall randomize QAM symbols forming the data section of the PLC frame using a copy of the linear feedback shift register in GF[2^{12}] used for randomizing the data subcarriers. This is shown in figure 7.63.

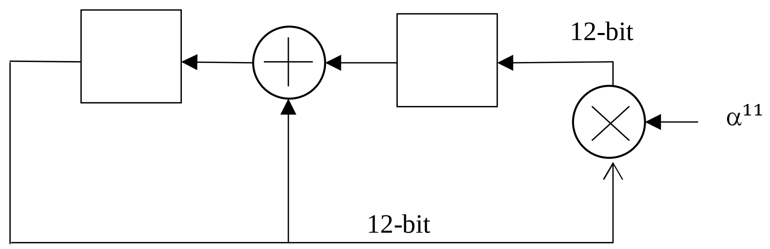

Figure 7.63: Linear Feedback Shift Register for PLC Randomization

The LFSR is defined by the following polynomial in GF[2^{12}].

$$x^2 + x + \alpha^{11}$$

The GF[2^{12}] is defined through polynomial algebra modulo the polynomial:

$$\alpha^{12} + \alpha^6 + \alpha^4 + \alpha + 1$$

This LFSR is initialized to the hexadecimal numbers given below:

D0 = "4A7"

D1 = "B4C"

The CMTS shall initialize the LFSR with the above two 12-bit numbers at the beginning of the first OFDM symbol following the PLC preamble. The CMTS shall clock the LFSR once after randomizing one PLC subcarrier. The CMTS shall randomize each subcarrier through an exclusive-OR operation of the 4 bits representing the subcarrier ($v_{t,f}$) with the four LSBs of register D0.

The first subcarrier to be randomized is the lowest frequency subcarrier of the PLC in the OFDM symbol immediately after the preamble. This will be randomized using the four LSBs of the initialized D0, namely 0x7. The LFSR will be clocked once after randomizing each PLC subcarrier of the OFDM symbol. After randomizing the highest frequency PLC subcarrier of an OFDM symbol the CMTS shall clock the LFSR before randomizing the lowest frequency PLC subcarrier in the next OFDM symbol.

The CMTS shall use the bit ordering given below to perform randomization. The four LSBs of D0 are defined as the coefficients of $\{\alpha^3\alpha^2\alpha^1\alpha^0\}$ of the Galois field polynomial representing D0. The LSB is defined as the coefficient of α^0 of the polynomial representing D0. The ordering of the four bits representing the subcarrier is defined with reference to figure 7.60. Assume that the FEC block shown in figure 7.60 is the first FEC block in the PLC frame. Then, since the location of the first nibble does not change as a result of interleaving:

$$v_{0,0} = \{a_0\ a_1\ a_2\ a_3\}$$

Then the randomization operation (i.e. exclusive-OR with 0x7) is given by:

$$\{y_0, y_1, y_2, y_3\} = \{a_0 + 1,\ a_1 + 1,\ a_2 + 1,\ a_3 + 0\}$$

The addition operations in the above equation are defined in GF[2], that is, these are bit-wise exclusive-OR operations. The LFSR is clocked once before randomizing the next nibble $v_{0,1}$.

The CMTS shall not randomize the PLC preamble.

7.5.13.9 Mapping to 16-QAM Subcarriers

The CMTS shall map each randomized nibble $\{y_0y_1y_2y_3\}$ into a complex number using the 16-QAM constellation mapping shown in figure 9.6.

The CMTS shall multiply the real and imaginary parts by $1/\sqrt{10}$ to ensure that mean-square value of the QAM constellation is unity.

7.5.13.10 PLC Timestamp Reference Point

The PLC subcarriers following the preamble may contain a timestamp.

The CMTS shall define this timestamp with reference to the first OFDM symbol following the preamble if such a timestamp exists. This OFDM symbol is indicated by an arrow in figure 7.64.

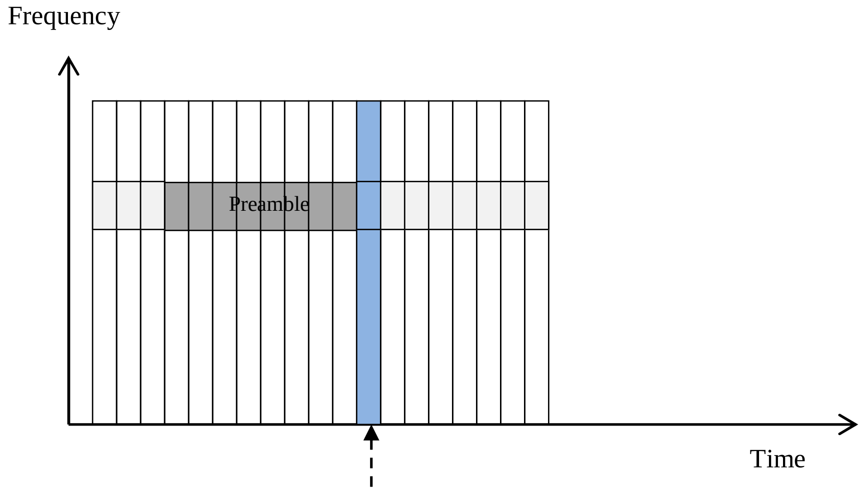

Figure 7.64: Time - Frequency Plane Representation of PLC Timestamp Synchronization

Time domain version of the OFDM symbol is shown in figure 7.65. The inverse discrete Fourier transform of the symbol of figure 7.64 results in the set of 4 096 or 8 192 samples occupying the FFT duration shown. After this the CMTS will introduce a configurable cyclic prefix (CP), window the symbol and overlap successive symbols in the time domain.

The CMTS shall use the time of the first sample of the FFT duration as the timestamp.

To clarify this further, individual time domain samples are also shown in figure 7.65. (This is for illustration only; actual samples are complex-valued.) The sample rate is 204,8 Msamples/s. The dotted arrow points to the first sample of the FFT symbol duration.

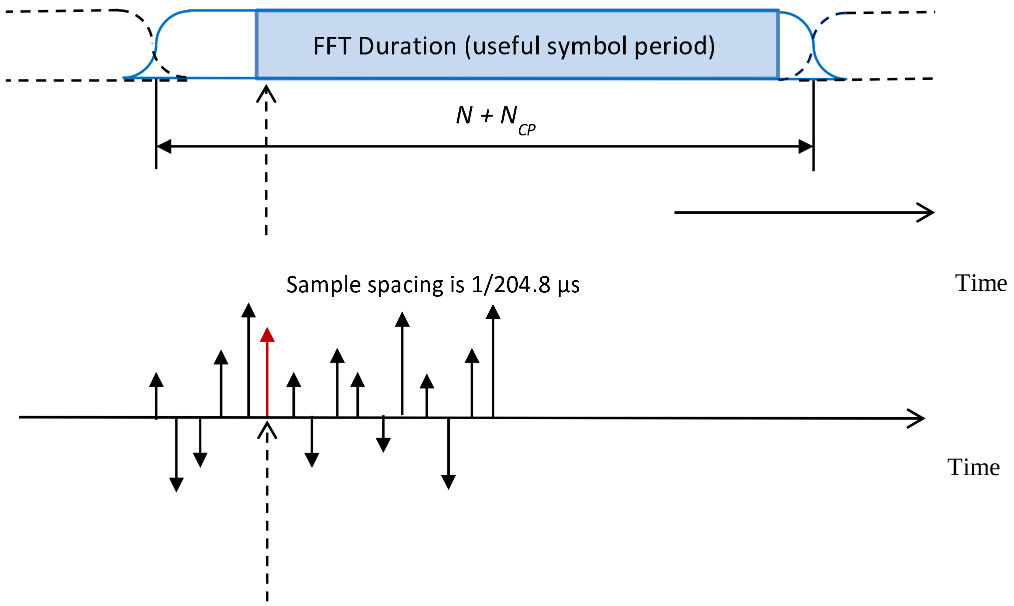

Figure 7.65: Time Domain Representation of PLC Timestamp Synchronization

7.5.14 Next Codeword Pointer

7.5.14.1 Mapping of Bytes to Bits

Each NCP consists of three bytes as defined in clause 8.3.4. The first byte (Byte 0) contains the profile identifier as the four MSBs and four control bits as the four LSBs. The other two bytes (Byte 1 and Byte 2) contain the start pointer.

The CMTS shall map the three NCP bytes into 24-bit serial bit stream $\{a_{23}\ a_{22}\ \ldots\ \ a_0\}$ for the purpose of LDPC encoding, as shown in figure 7.66. Note that the LDPC encoder is also defined using the same bit pattern $\{a_{23}\ a_{22}\ \ldots\ \ a_0\}$.

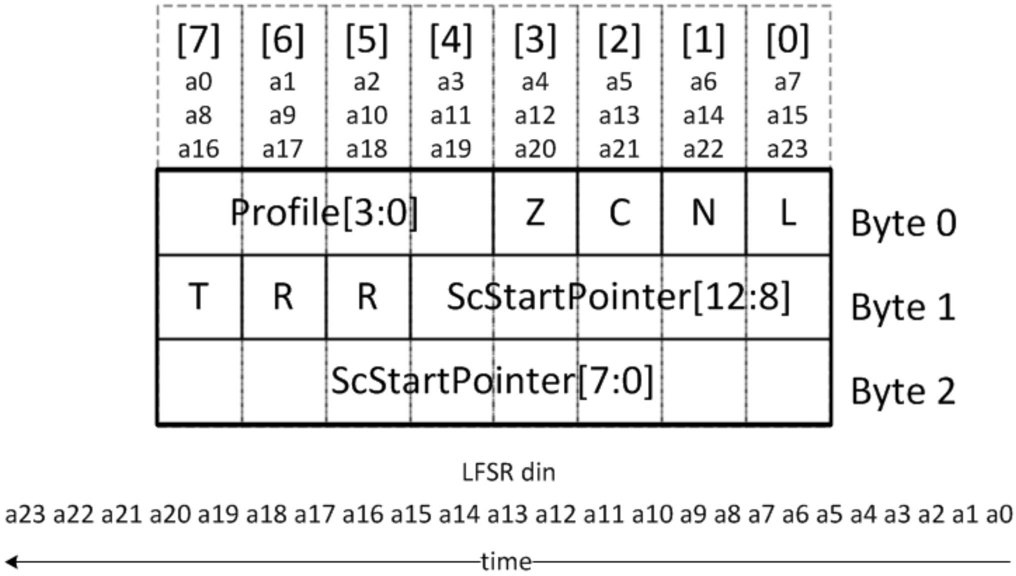

Figure 7.66: Mapping NCP Bytes into a Bit Stream for FEC Encoding

Figure 7.67 depicts the NCP bytes to input stream bits mapping after FEC encoding, including the FEC parity bits. FEC parity bits are specified in clause 7.5.14.2.

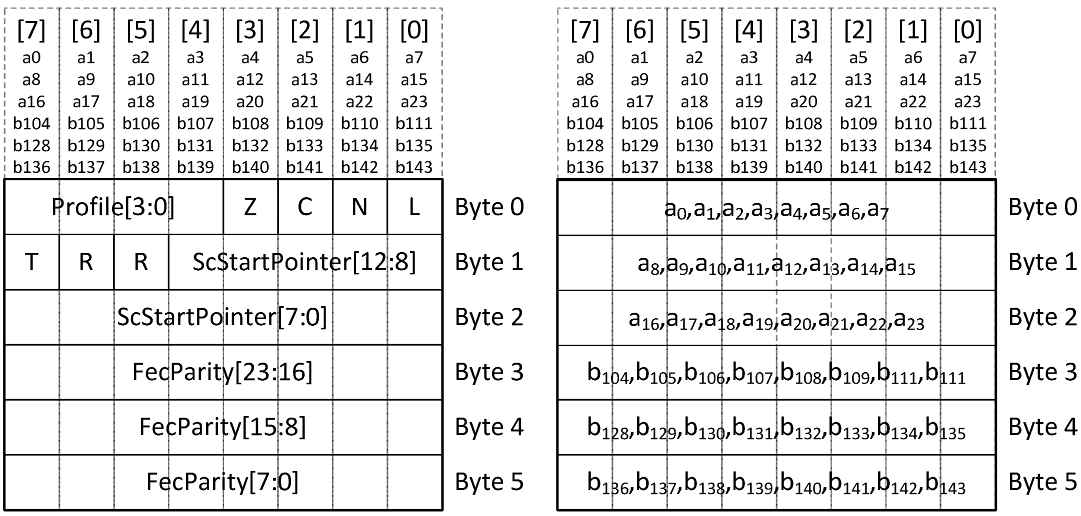

Figure 7.67: Mapping FEC Encoded NCP Bytes into a Bit Stream

7.5.14.2 CRC-24-D

The last NCP in the NCP field of each symbol contains a CRC-24-D message which is calculated across all NCPs in the NCP field of the symbol as specified in Annex E.

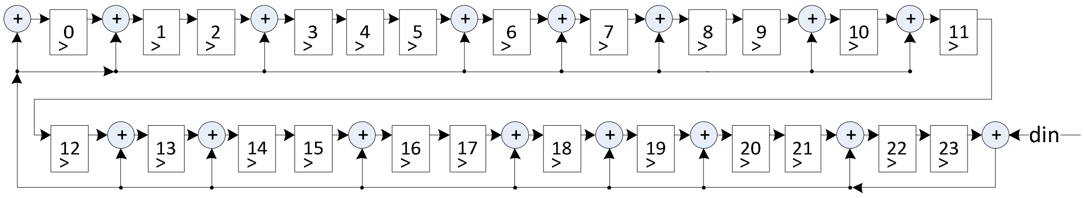

x^24 + x^22 + x^20 + x^19 + x^18 + x^16 + x^14 + x^13 + x^11 + x^10 + x^8 + x^7 + x^6 + x^3 + x + 1

NCP data is fed into the CRC-24-D encoder in the same order as the FEC encoder, as depicted in figure 7.68.

[7]	[6]	[5]	[4]	[3]	[2]	[1]	[0]	
a0	a1	a2	a3	a4	a5	a6	a7	
a8	a9	a10	a11	a12	a13	a14	a15	
a16	a17	a18	a19	a20	a21	a22	a23	
Profile[3:0]				Z	C	N	L	Byte 0
T	R	R	ScStartPointer[12:8]					Byte 1
ScStartPointer[7:0]								Byte 2

LFSR din

a23 a22 a21 a20 a19 a18 a17 a16 a15 a14 a13 a12 a11 a10 a9 a8 a7 a6 a5 a4 a3 a2 a1 a0

← time

Figure 7.68: Mapping NCP Data into the CRC-24-D Encoder

The 24-bit CRC output is represented as:

```
                   MSB          LSB
CRC-24-D-LFSR[23:0]= p0, p1,.. p22, p23
```

Figure 7.69 describes the mapping of the CRC-NCP bytes to input bit stream including the FEC parity bits.

[7]	[6]	[5]	[4]	[3]	[2]	[1]	[0]	
c0	c1	c2	c3	c4	c5	c6	c7	
c8	c9	c10	c11	c12	c13	c14	c15	
c16	c17	c18	c19	c20	c21	c22	c23	
b104	b105	b106	b107	b108	b109	b110	b111	
b128	b129	b130	b131	b132	b133	b134	b135	
b136	b137	b138	b139	b140	b141	b142	b143	
CRC-24-D-LFSR[23:16]								Byte 0
CRC-24-D-LFSR[15:0]								Byte 1
CRC-24-D LFSR[7:0]								Byte 2
FecParity[23:16]								Byte 3
FecParity[15:8]								Byte 4
FecParity[7:0]								Byte 5

[7]	[6]	[5]	[4]	[3]	[2]	[1]	[0]	
c0	c1	c2	c3	c4	c5	c6	c7	
c8	c9	c10	c11	c12	c13	c14	c15	
c16	c17	c18	c19	c20	c21	c22	c23	
b104	b105	b106	b107	b108	b109	b110	b111	
b128	b129	b130	b131	b132	b133	b134	b135	
b136	b137	b138	b139	b140	b141	b142	b143	
$c_0,c_1,c_2,c_3,c_4,c_5,c_6,c_7$								Byte 0
$c_8,c_9,c_{10},c_{11},c_{12},c_{13},c_{14},c_{15}$								Byte 1
$c_{16},c_{17},c_{18},c_{19},c_{20},c_{21},c_{22},c_{23}$								Byte 2
$b_{104},b_{105},b_{106},b_{107},b_{108},b_{109},b_{110},b_{111}$								Byte 3
$b_{128},b_{129},b_{130},b_{131},b_{132},b_{133},b_{134},b_{135}$								Byte 4
$b_{136},b_{137},b_{138},b_{139},b_{140},b_{141},b_{142},b_{143}$								Byte 5

Figure 7.69: Mapping FEC Encoded CRC-NCP Bytes into a Bit Stream

7.5.14.2 Forward Error Correction code for the NCP

The CMTS shall encode the 24 information bits of a Next Codeword Pointer using (48, 24) shortening and puncturing LDPC encoder, see clause 7.4.16.1.4 for the definition of shortening and puncturing encoder.

The shortening puncturing encoder uses the same mother encoder for initial ranging FEC (see clause 7.4.4.3), that is the rate 1/2 (160, 80) LDPC encoder listed by table 7.14.

Denote the information bits sent to the mother code encoder by $(a_0, \cdots, a_{79})$ and let the encoder output being $(a_0, \cdots, a_{79}, b_{80}, \cdots, b_{159})$, where $b_{80}, \cdots, b_{159}$ are parity-check bits. Then the shortening and puncturing steps can be described as follows, also see figure 7.70: Shortening and Puncturing Encoder for the NCP FEC:

The shortening step fills 0 to 56 consecutive coordinate starting at position 24, i.e. let $a_{24} = a_{25} = \cdots = a_{79} = 0$. The rest 24 bits i.e. $a_0, \cdots, a_{23}$, are NCP information data.

The coordinates to be deleted by the puncturing step are:

- period 1: 24 consecutive coordinates $b_{80}, \cdots, b_{103}$
- period 2: 16 consecutive coordinates $b_{112}, \cdots, b_{127}$
- period 3: 16 consecutive coordinates $b_{144}, \cdots, b_{159}$

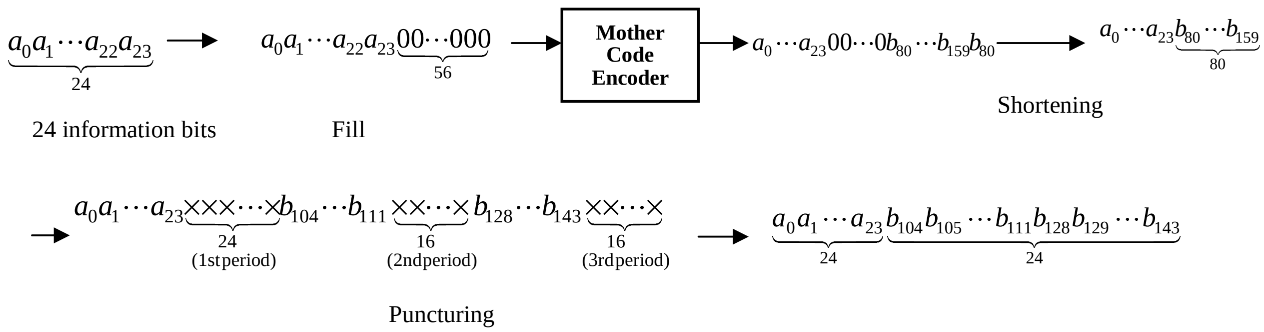

Figure 7.70: Shortening and Puncturing Encoder for the NCP FEC

7.5.14.3 Mapping LDPC Encoded Bits into OFDM Subcarriers

The LDPC encoder outputs a stream of 48 bits:

$$\{b_{143}\, b_{142}\, \ldots\, b_{128}\, b_{111}\, b_{110} \ldots b_{104}\, a_{23}\, a_{22}\, \ldots\, a_0\}$$

The NCP QAM constellation can be a member of the set {QPSK, QAM-16, QAM-64}.

For QAM-64 the CMTS shall map the LDPC encoded bits into eight 6-bit QAM constellation points as defined below:

$$\{y_{0,0}\, y_{0,1}\, y_{0,2}\, y_{0,3}\, y_{0,4}\, y_{0,5}\} = \{a_5\, a_4\, a_3\, a_2\, a_1\, a_0\}$$

$$\{y_{1,0}\, y_{1,1}\, y_{1,2}\, y_{1,3}\, y_{1,4}\, y_{1,5}\} = \{a_{11}\, a_{10}\, a_9\, a_8\, a_7\, a_6\}$$

$$\cdots$$

$$\{y_{7,0}\, y_{7,1}\, y_{7,2}\, y_{7,3}\, y_{7,4}\, y_{7,5}\} = \{b_{143}\, b_{142}\, b_{141}\, b_{140}\, b_{139}\, b_{138}\}$$

The mapping of these 6-bit integers to points in the complex plane is given by the figure below. Hexadecimal notation has been used to represent the 6-bit numbers $\{y_{i,0} y_{i,1} y_{i,2} y_{i,3} y_{i,4} y_{i,5}\}$.

The CMTS shall multiply the real and imaginary parts by 1/√42 to ensure that mean-square value of the QAM constellation is unity.

	-7	-5	-3	-1	1	3	5	7
7	20	22	2A	28	08	0A	02	00
5	21	23	2B	29	09	0B	03	01
3	25	27	2F	2D	0D	0F	07	05
1	24	26	2E	2C	0C	0E	06	04
-1	34	36	3E	3C	1C	1E	16	14
-3	35	37	3F	3D	1D	1F	17	15
-5	31	33	3B	39	19	1B	13	11
-7	30	32	3A	38	18	1A	12	10

Figure 7.71: 64-QAM Constellation Mapping of $\{y_{i,0}y_{i,1}y_{i,2}y_{i,3}y_{i,4}y_{i,5}\}$

For QAM-16 the CMTS shall map the LDPC encoded bits into twelve 4-bit QAM constellation points as defined below:

$$\{y_{0,0}\ y_{0,1}\ y_{0,2}\ y_{0,3}\} = \{a_3\ a_2\ a_1\ a_0\}$$

$$\{y_{1,0}\ y_{1,1}\ y_{1,2}\ y_{1,3}\} = \{a_7\ a_6\ a_5\ a_4\}$$

$$\ldots$$

$$\{y_{11,0}\ y_{11,1}\ y_{11,2}\ y_{11,3}\} = \{b_{143}\ b_{142}\ b_{141}\ b_{140}\}$$

The mapping of these 4-bit integers to points in the complex plane is given by the figure below. Hexadecimal notation has been used to represent the 4-bit numbers$\{y_{i,0}y_{i,1}y_{i,2}y_{i,3}\}$.

The CMTS shall multiply the real and imaginary parts by 1/√10 to ensure that mean-square value of the QAM constellation is unity.

	-3	-1	1	3
3	8	A	2	0
1	9	B	3	1
-1	D	F	7	5
-3	C	E	6	4

Figure 7.72: 16QAM Constellation Mapping of $\{y_{i,0}\ y_{i,1}\ y_{i,2}\ y_{i,3}\}$

For QPSK the CMTS shall map the LDPC encoded bits into twenty four 2-bit QAM constellation points as defined below:

$$\{y_{0,0}\ y_{0,1}\} = \{a_1\ a_0\}$$

$$\{y_{1,0}\ y_{1,1}\} = \{a_3\ a_2\}$$

$$\dots$$

$$\{y_{23,0}\ y_{23,1}\} = \{b_{143}\ b_{142}\}$$

The mapping of these 2-bit integers to points in the complex plane is given by the figure below. Hexadecimal notation has been used to represent the 2-bit numbers $\{y_{i,0} y_{i,1}\}$.

The CMTS shall multiply the real and imaginary parts by $1/\sqrt{2}$ to ensure that mean-square value of the QAM constellation is unity.

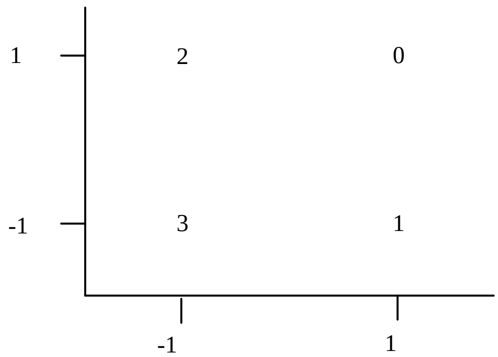

Figure 7.73: QPSK Constellation Mapping of $\{y_{i,0} y_{i,1}\}$

7.5.14.4 Placement of NCP Subcarriers

The CMTS shall place the NCP subcarriers beginning from the frequency location of the highest frequency active data subcarrier of the OFDM symbol, and going downwards along active data subcarriers of the OFDM symbols before they are time and frequency interleaved.

Therefore the first subcarrier of the first NCP occupies the frequency location of the highest frequency active data subcarrier of the OFDM symbol. The term active data subcarrier is used to indicate a subcarrier that is neither excluded and that is neither a continuous pilot nor a scattered pilot. This highest frequency active subcarrier may not occur at the same frequency in every symbol owing to the presence of scattered pilots.

The OFDM symbol, prior to time and frequency interleaving at the CMTS, will have subcarriers assigned to be scattered pilot placeholders. Furthermore, the NCP profile may indicate subcarriers that are to be zero-bit-loaded. The CMTS shall skip both of these types of subcarriers during the placement of NCP subcarriers.

7.5.14.5 Randomization and Interleaving of NCP Subcarriers

The CMTS shall randomize the NCP constellation points $\{y_{i,j}\}$ described in clause 7.5.14.3 using the algorithm applied to the data subcarriers, described in clause 7.5.5.3.

The CMTS shall time and frequency interleave the NCP subcarriers using the algorithm applied to data subcarriers and this is described in the interleaving section.

7.5.15 Downstream Pilot Patterns

7.5.15.0 Overview

Downstream pilots are subcarriers modulated by the CMTS with a defined modulation pattern that is known to all the CMs in the system to allow interoperability.

There are two types of pilots: continuous and scattered. Continuous pilots occur at fixed frequencies in every symbol. Scattered pilots occur at different frequency locations in different symbols. Each of these pilot types for DOCSIS 3.1 is defined in the following clauses.

7.5.15.1 Scattered Pilots

7.5.15.1.0 Scattered Pilots Overview

The main purpose of scattered pilots is the estimation of the channel frequency response for the purpose of equalization. There are two scattered pilot patterns, one for 4K FFT and one for 8K FFT. Although these pilots occur at different frequency locations in different OFDM symbols, the patterns repeat after every 128 OFDM symbols; in other words, the scattered pilot pattern has a periodicity of 128 OFDM symbols along the time dimension.

7.5.15.1.1 Scattered Pilot Pattern for 4K FFT

The CMTS shall create scattered pilots for 4K FFTs in the manner described in this clause.

Figure 7.74 shows the 4K FFT scattered pilot pattern for OFDM transmissions.

The scattered pilot pattern is synchronized to the PLC as shown in figure 7.74. The first OFDM symbol after the PLC preamble has a scattered pilot in the subcarrier just after the highest frequency subcarrier of the PLC. Two such scattered pilots that are synchronized to the PLC preamble are marked as red circles in figure 7.77.

The remainder of the scattered pilot pattern is linked to the scattered pilot synchronized to the PLC preamble, using the following rules:

1) In each symbol scattered pilots are placed every 128 subcarriers.

2) From symbol to symbol, scattered pilots are shifted by one subcarrier position in the increasing direction of the frequency axis. This will result in scattered pilots placed in the exclusion band and in the PLC band.

3) Scattered pilots are zero-valued in the exclusion bands.

4) Scattered pilots are zero-valued when these coincide with excluded subcarriers.

5) In the PLC, normal PLC signals (i.e. PLC data or the PLC preamble) are transmitted instead of scattered pilots. The CMTS shall not transmit scattered pilots in the PLC band.

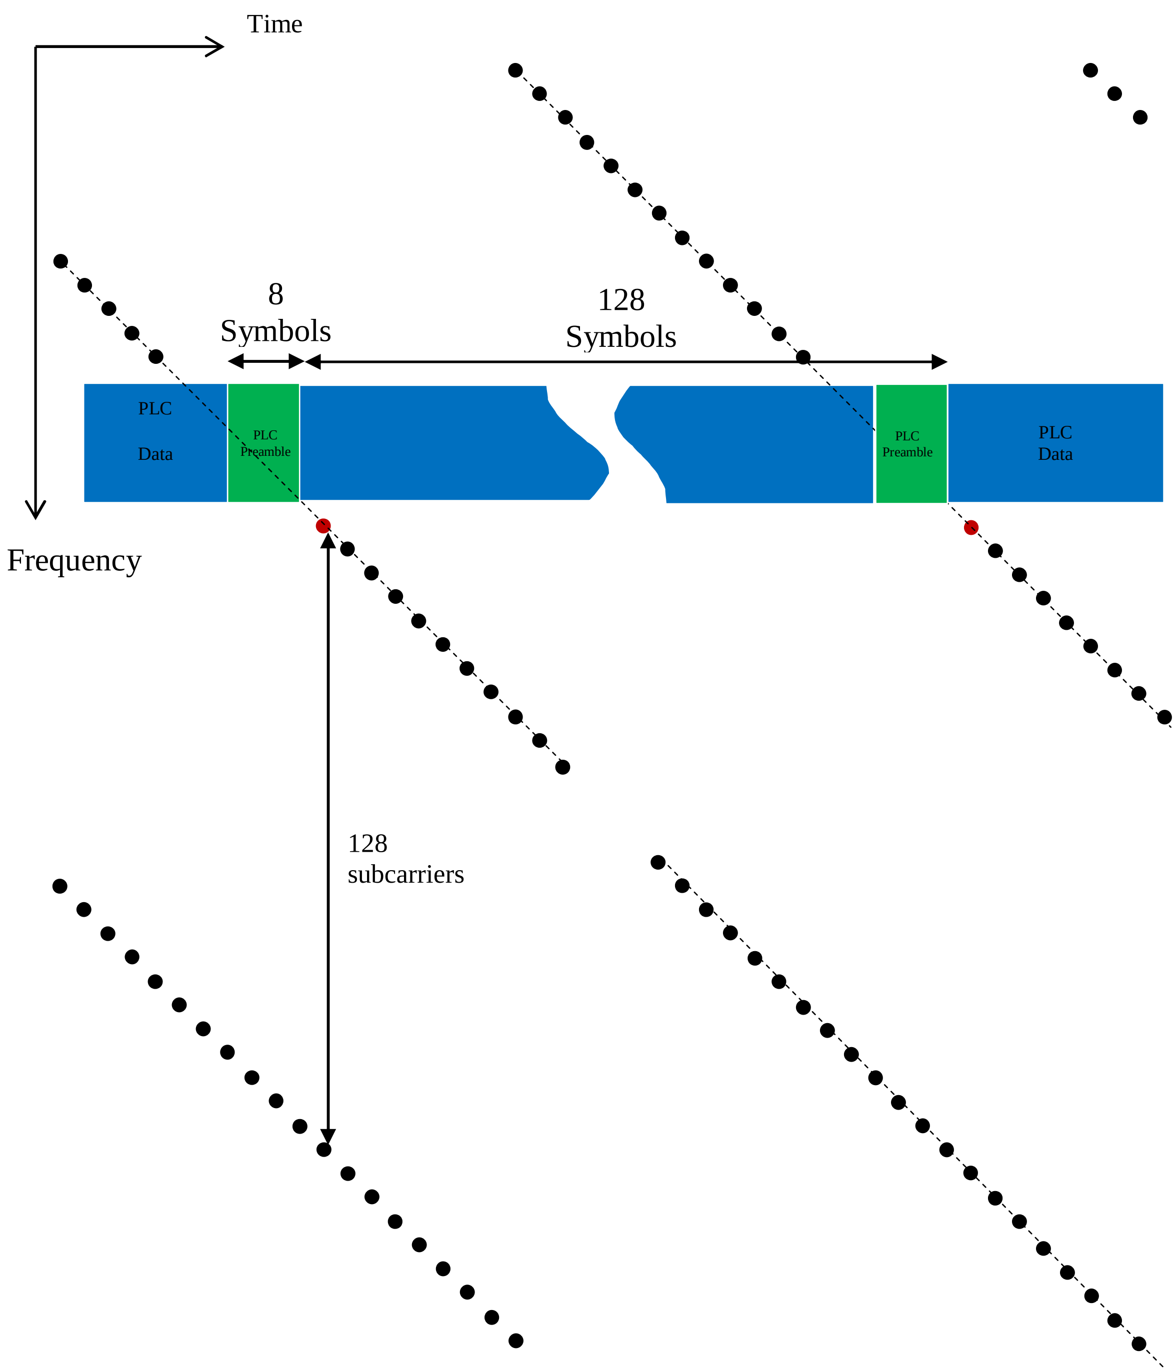

Figure 7.74: 4K FFT Downstream Pilot Pattern

There are 8 preamble symbols in the PLC; for 4K FFT, there are 8 PLC subcarriers in each symbol.

Mathematically, the scattered pilot pattern for a 4K FFT is defined as follows. Let a subcarrier (depicted in red in the above figure just after the PLC preamble) be referred to as $x(m,n)$, where:

m is the frequency index

n is the time index (i.e. the OFDM symbol number)

The scattered pilots in the 128 symbols following (and including symbol n) are given by:

Symbol n: x(n, m±128i), for all non-negative integers i

Symbol (n+1): x(n+1, m±128i + 1) , for all non-negative integers i

Symbol (n+2): x(n+2, m±128i + 2) , for all non-negative integers i

⋮

Symbol (n+127): x(n+127, m±128i + 127) , for all non-negative integers i

Each of the above locations is a scattered pilot, provided that it does not fall on a continuous pilot, on the PLC, on an exclusion zone or on a excluded subcarrier. If the scattered pilot coincides with a continuous pilot it is treated as a continuous pilot and not as a scattered pilot.

This pattern repeats every 128 symbols. That is, symbol (128+n) has the same scattered pilot pattern as symbol n.

7.5.15.1.2 Scattered Pilot Pattern for 8K FFT

The CMTS shall create scattered pilots for 8K FFTs in the manner described in this clause.

Figure 7.75 shows a scattered pilot pattern that may be used for OFDM transmissions employing 8K FFT. This is used here for explanation purposes only and to help with the derivation of the scattered pilot pattern actually used in 8K FFT OFDM transmissions depicted in figure 7.76.

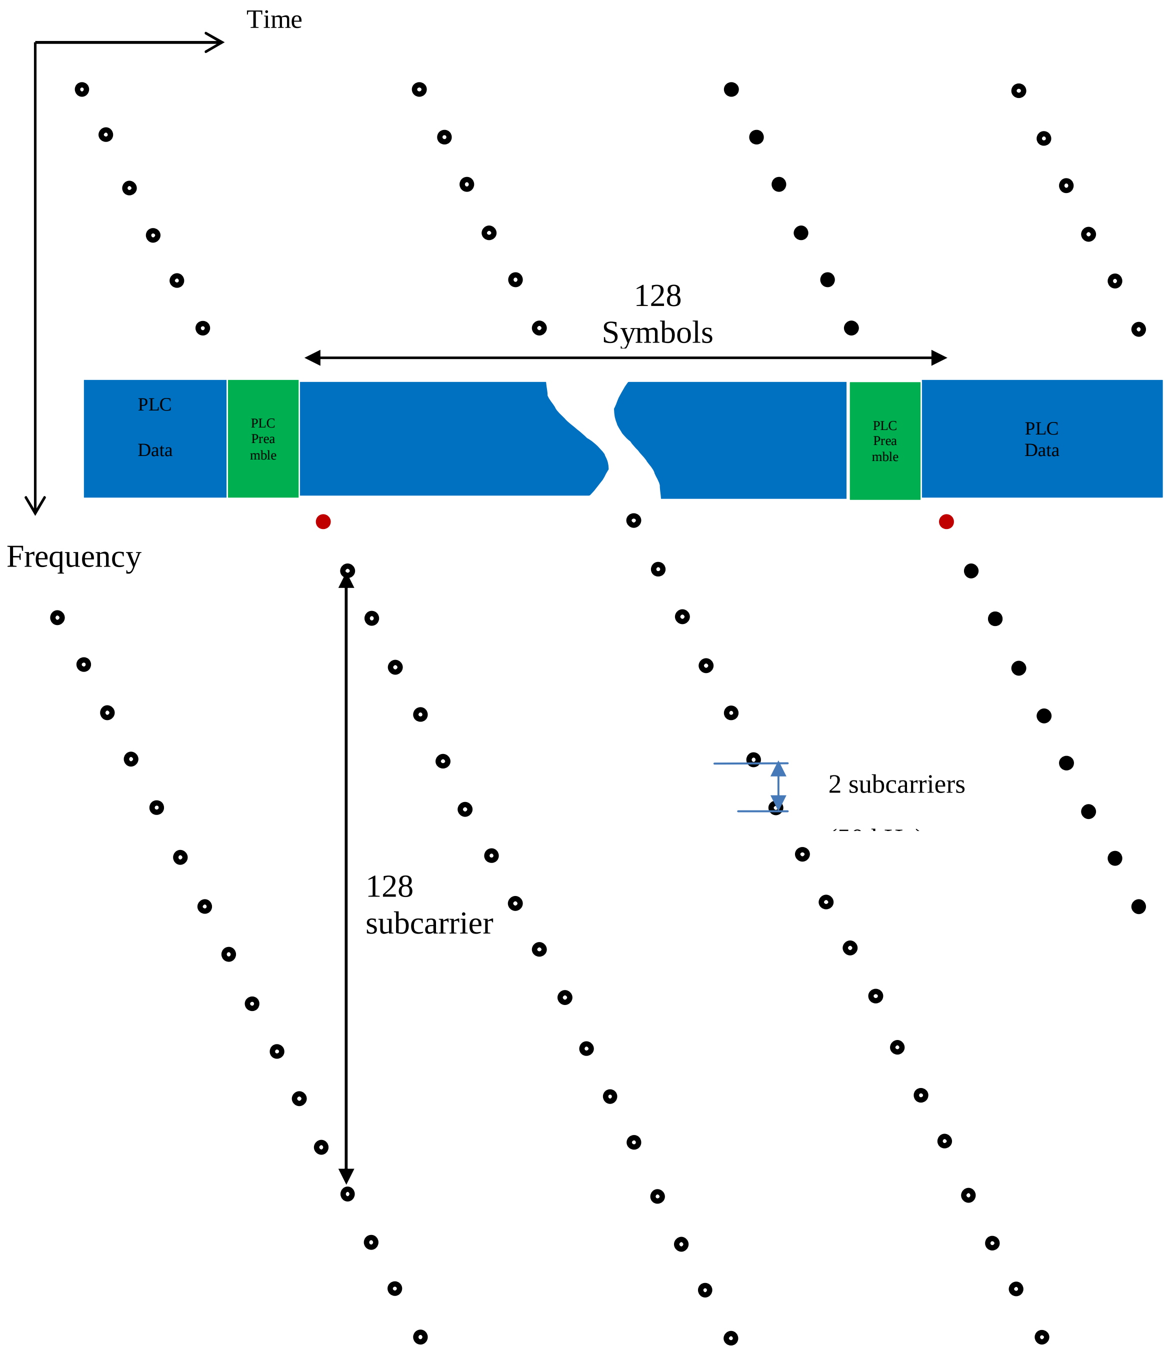

Figure 7.75: A Downstream Scattered Pilot Pattern for 8K FFT (for Explanation Purposes Only)

The scattered pilot pattern is synchronized to the PLC as shown in figure 7.74. The first OFDM symbol after the PLC preamble has a scattered pilot in the subcarrier just after the highest frequency subcarrier of the PLC. Two such scattered pilots that are synchronized to the PLC preamble are marked as red circles in figure 7.75.

In the case of an 8K FFT, pilots are stepped by two subcarriers from one OFDM symbol to the next. Since the pilot spacing along the frequency axis is 128, this results in a pilot periodicity of 64 in the time dimension. When figure 7.74 and figure 7.75 are compared, it is clear that the periodicity is half for the 8K scattered pilot pattern. However, because an 8K symbol is twice as long as a 4K symbol, the scattered pilot periodicity in terms of actual time is approximately the same for both the 4K and 8K FFTs. This allows channel estimates for 8K FFTs to be obtained in approximately the same amount of time as for the 4K FFT. However, scattered pilots for 8K FFTs do not cover all subcarrier locations and hence intermediate channel estimates have to be obtained through interpolation.

Noise can also be estimated using scattered pilots, and again, the noise at subcarrier locations not covered by scattered pilots in the 8K FFT can be obtained through interpolation. Note that this interpolation operation could fail in the presence of narrowband ingress; interpolation could also be problematic when there are excluded subcarriers.

To overcome these interpolation problems, the entire 8K scattered pilot location can be shifted by one subcarrier location after 64 subcarriers, as illustrated in figure 7.76. This may be treated as the interlacing of two identical scattered pilot patterns. The set of purple scattered pilots are shifted one subcarrier space in relation to the set of green scattered pilots. As a result the scattered pilots cover all subcarrier locations; noise at every subcarrier location can be estimated without interpolation. Note that periodicity of the 8K FFT scattered pilot pattern is now 128, not 64.

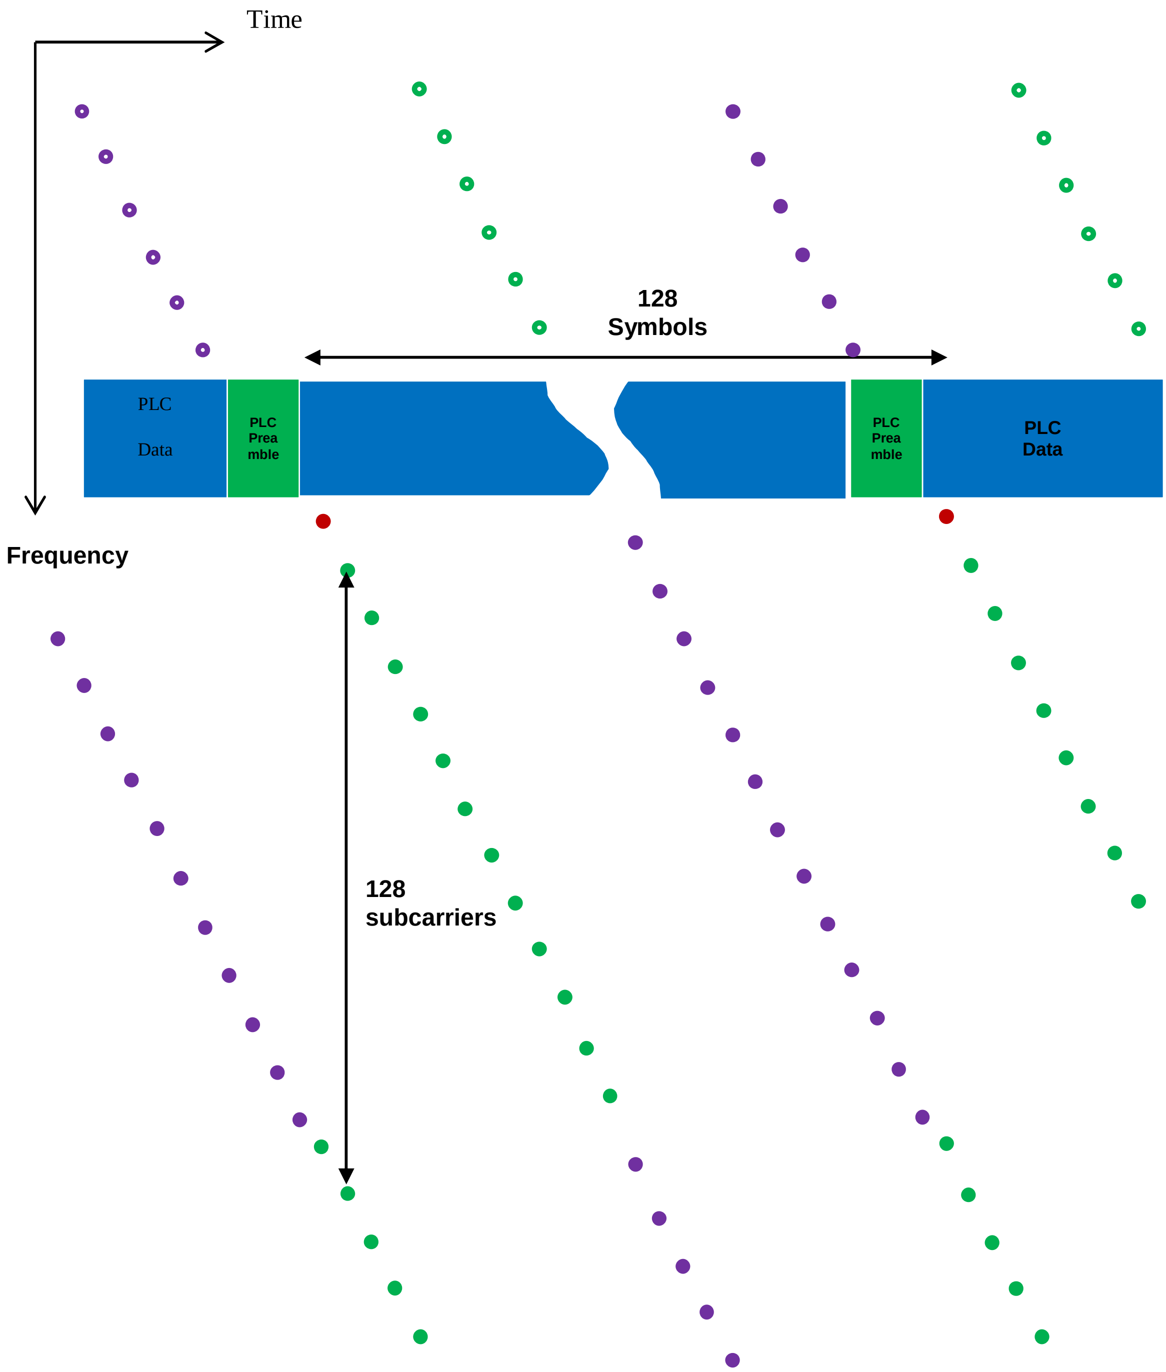

Figure 7.76: 8K FFT Downstream Scattered Pilot Pattern

Mathematically, the scattered pilot pattern for an 8K FFT is defined as follows. Let the subcarrier (depicted in red in figure 7.76 just after the PLC preamble) be referred to as $x(m, n)$ where:

m is the frequency index

n is the time index (i.e. the OFDM symbol number)

The scattered pilots in the first 64 symbols following and including symbol *n* are given by:

Symbol n: $\mathrm{x(n, m \pm 128i)}$, for all non-negative integers i

Symbol (n+1): $\mathrm{x(n+1, m \pm 128i + 2)}$, for all non-negative integers i

Symbol (n+2): $\mathrm{x(n+2, m \pm 128i + 4)}$, for all non-negative integers i

⋮

Symbol (n+63): $\mathrm{x(n+63, m \pm 128i + 126)}$, for all non-negative integers i

The scattered pilot sequence of the next 64 symbols is the same as above, but with a single subcarrier shift in the frequency dimension.

Symbol (n+64): $\mathrm{x(n+64, m \pm 128i + 1)}$, for all non-negative integers i

Symbol (n+65): $\mathrm{x(n+65, m \pm 128i + 3)}$, for all non-negative integers i

Symbol (n+66): $\mathrm{x(n+66, m \pm 128i + 5)}$, for all non-negative integers i

⋮

Symbol (n+127): $\mathrm{x(n+127, m \pm 128i + 127)}$, for all non-negative integers i

⋮

Each of the above locations is a scattered pilot, provided that it does not fall on a continuous pilot, on the PLC, on an exclusion band or on an excluded subcarrier. If the scattered pilot coincides with a continuous pilot it is treated as a continuous pilot and not as a scattered pilot.

This pattern repeats every 128 symbols. That is, symbol (128+n) has the same scattered pilot pattern as symbol n.

7.5.15.2 Continuous Pilots

7.5.15.2.0 Continuous Pilots Overview

Continuous pilots occur at the same frequency location in all symbols and are used for receiver synchronization. Placement of continuous pilots is determined in two ways:

a) Predefined continuous pilot placement around the PLC

b) Continuous pilot placement defined via PLC messages

Note that continuous and scattered pilots can overlap; the amount of overlap, in terms of number of carriers, changes from symbol to symbol. Overlapping pilots are treated as continuous pilots.

7.5.15.2.1 Predefined Continuous Pilots Around the PLC

As discussed in clause 7.5.13.1, the PLC is placed at the centre of a 6 MHz spectral region. Four pairs of predefined continuous pilots are placed symmetrically around the PLC as shown in figure 7.77. The spacing between each pilot pair and the PLC are different to prevent all pilots from being impacted at the same time by echo or interference.

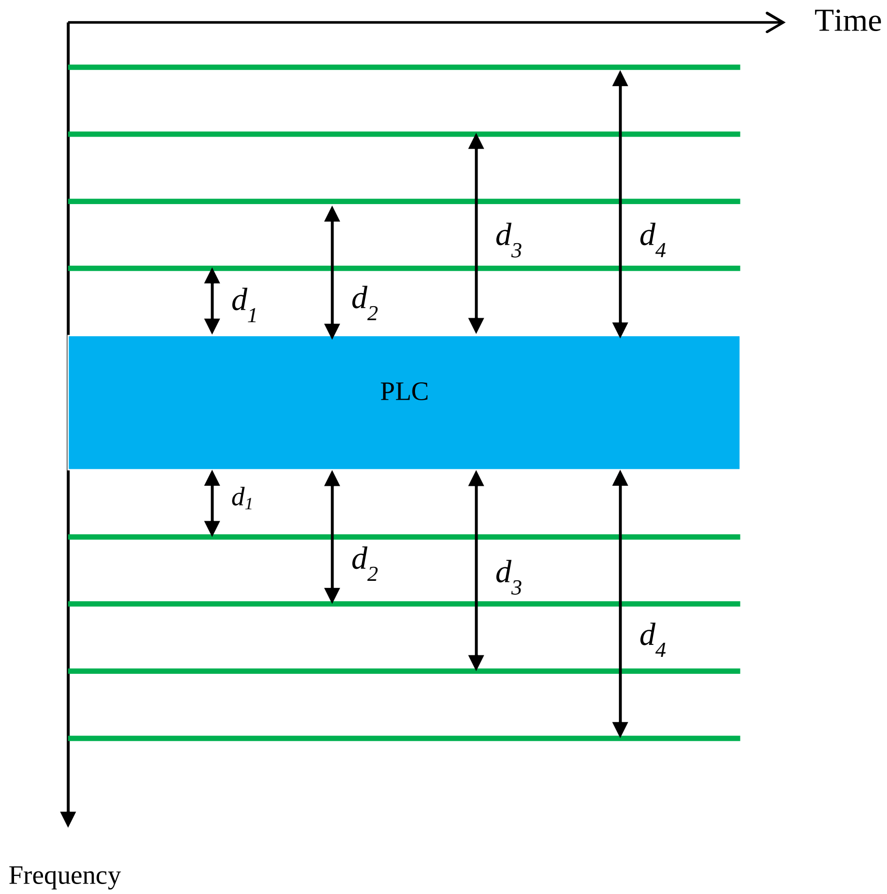

Figure 7.77: Placement of Predefined Continuous Pilots around the PLC

The locations of the continuous pilots are defined with reference to the edges of the PLC band. Hence, once the PLC has been detected, these continuous pilots also become known to the receiver.

Table 7.45 provides the values of d_1, d_2, d_3, and d_4, measured in number of subcarriers from the PLC edge. That is, d_x is absolute value of the difference between the index of the continuous pilot and the index of the PLC subcarrier at the PLC edge nearest to the continuous pilot. The index of a subcarrier is the integer k of the IDFT definition given in clause 7.5.2.8. For example, let the lowest frequency subcarrier of the PLC have the IDFT index k equal to 972. Then according to table 7.45 for the 4K FFT mode the continuous pilot nearest to this lowest frequency PLC subcarrier will have the IDFT index k of (972-15)=957. The index k of the highest frequency PLC subcarrier of this OFDM channel is 979. Hence continuous pilot that is nearest upper frequency edge of the PL has an index k of 994.

The table provides the number of subcarriers from the edge of the PLC to the placement of the pilot for the two FFT sizes. For each distance (d_x) defined in table 7.45, the CMTS shall place two pilots: one d_x subcarriers above and one d_x subcarriers below the edge of the PLC band.

Table 7.45: Subcarrier Distances for Placement of Predefined Pilots

		d_1	d_2	d_3	d_4
4K FFT	**PLC 8 subcarriers**	15	24	35	47
8K FFT	**PLC 16 subcarriers**	30	48	70	94

7.5.15.2.2 Continuous Pilot Placement Defined by PLC Message

The CMTS shall define a set of continuous pilots distributed as uniformly as possible over the entire OFDM spectrum in addition to the predefined continuous pilots described in the preceding section.

The CMTS shall ensure that there are no isolated active OFDM spectral regions that are not covered by continuous pilots.

It is not practical to predefine the locations of this set of continuous pilots because of exclusion bands and excluded subcarriers.

The CMTS shall provide the continuous pilot placement definition via the PLC in accordance with messaging formats contained in the MULPI specification.

The CMTS shall adhere to the rules given below for the definition of this set of continuous pilot locations conveyed to the CM via PLC messaging. It is noted that these rules do not apply to the eight predefined pilots.

The CMTS shall place the continuous pilots generated using these rules in every OFDM symbol, in addition to the eight predefined continuous pilots.

The CMTS shall obtain the value of N_{CP} using the following formula:

$$N_{CP} = \min\left(\max\left(8, ceil\left(M * \left(\frac{F_{max}-F_{min}}{190e6}\right)\right)\right), 120\right) \quad (1)$$

In this equation F_{max} refers to frequency in Hz of the highest frequency active subcarrier and F_{min} refers to frequency in Hz of the lowest frequency active subcarrier of the OFDM channel. It is observed that the number of continuous pilots is linearly proportional to the frequency range of the OFDM channel. It may also be observed that the minimum number of continuous pilots defined using the PLC cannot be less than 8, and the maximum number of continuous pilots defined using the PLC cannot exceed 120. Therefore, the total number of continuous pilots, including the predefined ones, will be in the range 16 to 128, both inclusive.

The value of M in equation (1) is kept as a parameter that can be adjusted by the CMTS. Nevertheless, the CMTS shall ensure that M is in the range given by the following equation:

$$120 \geq M \geq 48 \quad (2)$$

The typical value proposed for M is 48.

The CMTS shall use the algorithm given below for defining the frequencies for the location of these continuous pilots.

Step 1:

Merge all the subcarriers between F_{min} and F_{max} eliminating the following:

- Exclusion bands
- 6 MHz band containing the PLC
- Known regions of interference, e.g. LTE
- Known poor subcarrier locations, e.g. CTB/CSO

Let the merged frequency band be defined as the frequency range $[0, F_{merged\_max}]$.

Step 2:

Define a set of N_{CP} frequencies using the following equation:

$$F_i = \frac{F_{merged\_max}}{2N_{CP}} + \frac{i*F_{merged\_max}}{N_{CP}}; \text{ for } i = 0, 1, \ldots, N_{CP} - 1 \quad (3)$$

This yields a set of uniformly spaced NCP frequencies:

$$\left\{\frac{F_{merged\_max}}{2N_{CP}}, \frac{3F_{merged\_max}}{2N_{CP}}, \ldots, F_{merged\_max} - \frac{F_{merged\_max}}{2N_{CP}}\right\} \quad (4)$$

Step 3:

Map the set of frequencies given above to the nearest subcarrier locations in the merged spectrum. This will give a set of N_{CP} approximately uniformly spaced subcarriers in the merged domain.

Step 4:

De-merge the merged spectrum through the inverse of the operations through which the merged spectrum was obtained in step 1.

Step 5:

If any continuous pilot is within 1 MHz of a spectral edge, move this inwards (but avoiding subcarrier locations impacted by interferences like CSO/CTB) so that every continuous pilot is at least 1 MHz away from a spectral edge. This is to prevent continuous pilots from being impacted by external interferences. If the width of the spectral region does not allow the continuous pilot to be moved 1 MHz from the edge then the continuous pilot has to be placed at the centre of the spectral band.

Step 6:

Identify any spectral regions containing active subcarriers (separated from other parts of the spectrum by exclusion bands on each side) that do not have any continuous pilots. Introduce an additional continuous pilot at the centre of every such isolated active spectral region.

In the unlikely event that the inclusion of these extra pilots results in the total number of continuous pilots defined by PLC exceeding 120, return to step 1 and re-do the calculations after decrementing the value of N_{CP} by one.

Step 7:

Test for periodicity in the continuous pilot pattern and disturb periodicity, if any, through the perturbation of continuous pilot locations using a suitable algorithm. A simple procedure would be to introduce a random perturbation of up to ±5 subcarrier locations around each continuous pilot location, but avoiding subcarrier locations impacted by interferences like CSO/CTB.

The CMTS shall transmit this continuous pilot pattern to the CMs in the system using the PLC.

7.5.15.3 Pilot Modulation

For both continuous and scattered pilots, the CMTS shall modulate these subcarriers as described in the following section.

Continuous and scattered pilots are BPSK modulated using a pseudo-random sequence. This pseudo-random sequence is generated using a 13-bit linear feedback shift register, shown in figure 7.78 with polynomial (x^13+x^12+x^11+x^8+1).

This linear feedback shift register is initialized to all ones at the *k*=0 index of the 4K or 8K discrete Fourier transform defining the OFDM signal (refer to clause 7.5.7). It is then clocked after every subcarrier of the FFT. If the subcarrier is a pilot (scattered or continuous), then the BPSK modulation for that subcarrier is taken from the linear feedback shift register output.

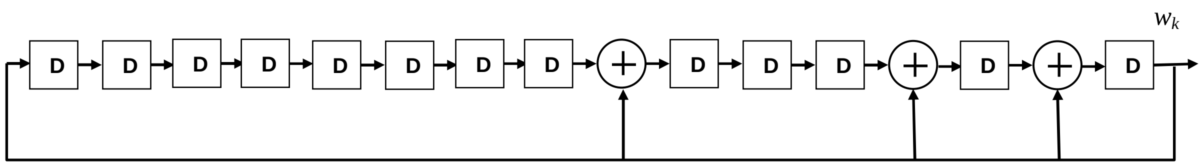

Figure 7.78: 13-Bit Linear Feedback Shift Register for the Pilot Modulation Pseudo-Random Sequence

For example, let the output of the linear feedback shift register be w_k. The BPSK modulation used for the pilot would be:

$$w_k = 0{:}\ BPSK\ Constellation\ Point = 1 + j0$$

$$w_k = 1{:}\ BPSK\ Constellation\ Point = -1 + j0$$

7.5.15.4 Pilot Boosting

The CMTS shall multiply the real and imaginary components of continuous and scattered pilots by a real-valued number such that the amplitude of the continuous and scattered pilots is twice the root-mean-square value of the amplitude of other subcarriers of the OFDM symbol; that is, continuous and scattered pilots are boosted by approximately 6 dB with reference to other subcarriers.

8 PHY-MAC Convergence

8.1 Scope

The present document defines the electrical characteristics and signal processing operations for a cable modem (CM) and Cable Modem Termination System (CMTS). It is the intent of the present document to define an interoperable CM and CMTS such that any implementation of a CM can work with any CMTS. It is not the intent of the present document to imply any specific implementation.

This clause describes CM and CMTS requirements for the convergence logical layer between the MAC and PHY layers for OFDM downstream channels and OFDMA upstream channels. The primary roles of the convergence layer are to map DOCSIS MAC frames into codewords and to map codewords into minislots for transmission from the cable modem to the CMTS. Contents of the Next Codeword Pointer (NCP) message and PHY Link Channel (PLC) are also defined in clause 8.

8.2 Upstream

8.2.1 Profiles

Upstream profiles are comprised of multiple minislots, and are characterized by bit loading and pilot pattern. Bit loading and pilot patterns can vary between minislots within the profile. The bit loading and pilot pattern assignment of minislots can also vary between profiles. An upstream profile maps to an Interval Usage Code defined in an Upstream Channel Descriptor Message.

Different FEC codeword sizes may use portions of a single minislot. The use case for this is as follows: With a 17 kb grant, there needs to be a long codeword to cover the first 16 200 bits and a 1 kb codeword to cover the rest of the bits. The first long codeword can land in the middle of a minislot. In this situation, it does not make sense to require a constant codeword size per profile, as the profile needs to cover a group of minislots.

FEC codewords can cross minislot and frame boundaries.

8.2.2 Upstream Subcarrier Numbering Conventions

Subcarriers are numbered from lower frequency to higher frequency within a FFT block. All subcarriers within the 102,4 MHz bandwidth are numbered, including the outside excluded subcarriers. Numbering starts at 0 and goes to 2 047 for 2K FFT, and 0 to 4 095 for 4K FFT.

Data codewords are mapped into minislots - prior to time and frequency interleaving - as described in clause 8.2.3.10.

8.2.3 Minislots

8.2.3.0 Minislots Overview

Minislots are defined by a size in terms of the number of symbols and number of subcarriers. They include data carried on data subcarriers, pilots carried on pilot subcarriers and complementary pilots that can carry data but at a lower modulation order.

In this clause, BW is defined as the encompassed spectrum on a single OFDMA channel.

8.2.3.1 Minislot Parameters

The CMTS shall define minislot parameters according to table 8.1. The CMTS communicates minislot definition to the CM in UCD messages as defined in [4].

The CM shall use the minislot structure defined by the UCD messages received from the CMTS.

The CMTS shall be capable of receiving minislots structured according to table 8.1.

The CMTS shall apply any subcarrier exclusions to the entire channel independent of upstream profile assignment.

Table 8.1: Minislot Parameters

Parameter	Minimum Value	Maximum Value	Recommended or Typical Value
K Number of symbol periods per frame	6	For BW > 72 MHz 18 for 20 µs FFT duration 9 for 40 µs FFT duration For 48 < BW < 72 MHz 24 for 20 µs FFT duration 12 for 40 µs FFT duration For BW < 48 MHz 36 for 20 µs FFT duration 18 for 40 µs FFT duration	N/A
Q Number of subcarriers per minislot	8 for 20 µs FFT duration 16 for 40 µs FFT duration		
Data bitloading	QPSK	CMTS Mandatory: 1024-QAM CMTS Optional: 2048-QAM 4096-QAM CM Mandatory: 4096-QAM	Plant-dependent
Complementary-pilot bitloading	BPSK	256-QAM	
Modulation order for primary pilots	BPSK	BPSK	
NOTE 1: Data bitloading is constant within a minislot, excepting pilots and complementary pilots. NOTE 2: The bitloading of complementary pilots within a minislot is constant.			

8.2.3.2 Minislot Structure

8.2.3.2.1 Number of OFDM Symbols per Minislot

The CMTS shall follow the rules listed below for the range of the frame size in number of symbols (K):

- K is configurable between 6 (minimal value) and one of the following values:

  With 20 µs FFT duration (2K FFT)

  - K_{max} = 18 for BW > 72 MHz
  - K_{max} = 24 for 48 MHz < BW < 72 MHz
  - K_{max} = 36 for BW < 48 MHz

  With 40 µs FFT duration (4K FFT)

  - K_{max} = 9 for BW > 72 MHz

- K_{max} = 12 for 48 MHz < BW < 72 MHz
- K_{max} = 18 for BW < 48 MHz

8.2.3.2.2 Number of subcarriers per minislot

The CMTS signals the number of subcarriers per minislot to the CM in the UCD.

The CMTS shall use 16-subcarrier minislots when the subcarrier spacing is 25 kHz.

The CMTS shall use 8-subcarrier minislots when the subcarrier spacing is 50 kHz.

8.2.3.3 Modulation of Data Subcarriers

With the exception of pilots and complementary pilots, bit loading is constant within a minislot.

The CMTS shall use the same modulation order for all data subcarriers in a minislot.

8.2.3.4 Location of Pilots

A set of pilot patterns is defined from which the CMTS or operator can select to match the frequency response of the network. Pilot patterns are described in clause 7.4.17.

8.2.3.5 Modulation Order of Pilots

The CMTS shall use BPSK modulation for pilots.

8.2.3.6 Location of Complementary Pilots

The CMTS shall place complementary pilots as defined by the chosen minislot pattern.

8.2.3.7 Modulation Order of Complementary Pilots

The CMTS shall use a modulation order equal to (data modulation order - 4) for complementary pilots.

The CMTS shall use a minimum modulation order of BPSK for complementary pilots.

8.2.3.8 Pilot Overhead

Pilot overhead is dependent on the chosen minislot pattern. Capacity and pilot overhead vary with the length of the minislot (number of symbols) and with the number of subcarriers (8 or 16). Minimum capacity and largest pilot overhead occur with the shortest minislot length (8 symbols).

8.2.3.9 Mapping Minislots to Subcarriers

The CMTS shall construct minislots using only contiguous subcarriers. There is no subcarrier exclusion within a minislot.

8.2.3.10 Ordering of Data Bits within a Minislot

With the exception of initial ranging transmissions, the CM shall fill minislots as follows: prior to interleaving, data would be filled across all symbol periods, subcarrier by subcarrier, transmitted symbol period by symbol period, with complementary pilots filled inline. The fill order is illustrated in figure 8.1.

NOTE: The position of the pilots shown in figure 8.1 is for illustrative purposes only and is not intended to be prescriptive.

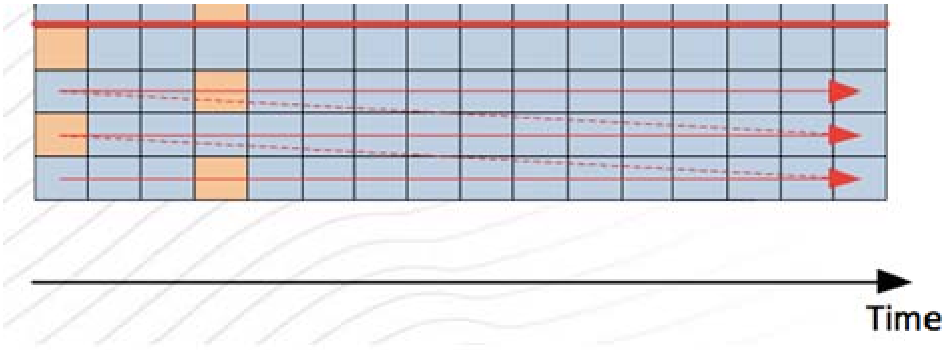

Figure 8.1: Minislot Data Bit Ordering

8.2.3.11 Modulation Order Variability

Different minislots are allowed to have different pilot patterns: pilot patterns are assigned at minislot granularity.

Different minislots are allowed to have different bit loading: bit loading is assigned at minislot granularity. This allows bit loading to vary across the spectrum.

8.2.4 Subslots

8.2.4.1 Subslot Structure

The minislot can be subdivided along time into multiple subslots to provide multiple transmission opportunities for BW requests, as shown in Figure 8.2. The subslots fit within the minislot boundary and can have leftover symbols in the end along the time axis. Leftover symbols are not a part of any subslot and are unused (zero valued subcarriers) when the minislot is granted to IUCs 1 or 2. Gaps between subslots within a single minislot are not permitted.

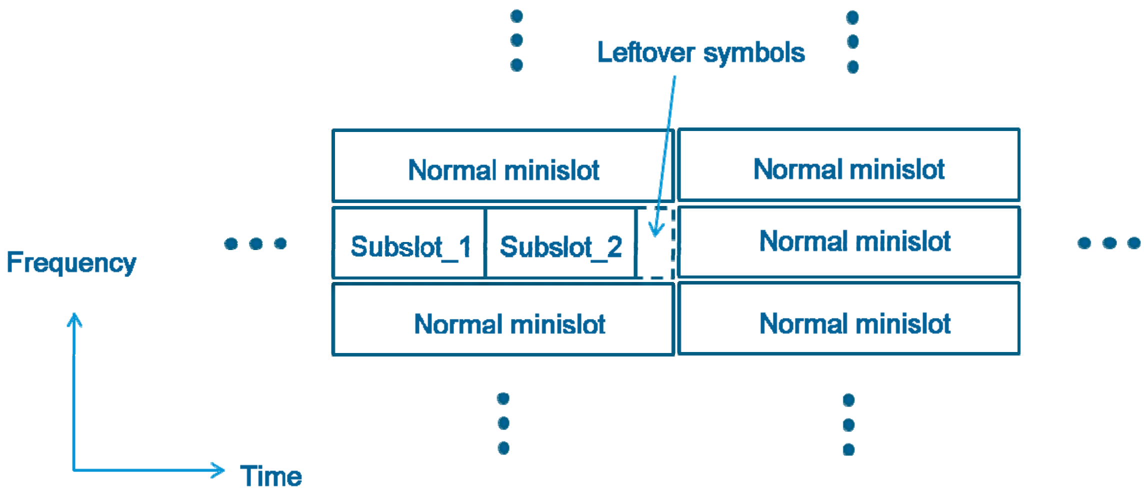

Figure 8.2: Subslot Structure

The subslot length is fixed at 2 symbols with minislots that employ 16 subcarriers, and 4 symbols with minislots that employ 8 subcarriers.

8.2.4.2 Data Mapping Among Subslots

The data mapping within a subslot is as follows: the mapping starts from the lowest subcarriers index, and lowest symbol index, and first along symbol time index, and then goes up on subcarriers index. The pilots are skipped during data mapping.

Data mapping to subcarriers is implemented without time or frequency interleaving.

Figure 8.3 and figure 8.4 illustrate the mapping process. Note that the positions of the pilots shown in figure 8.3 and figure 8.4 are for illustrative purposes only and are not intended to be prescriptive.

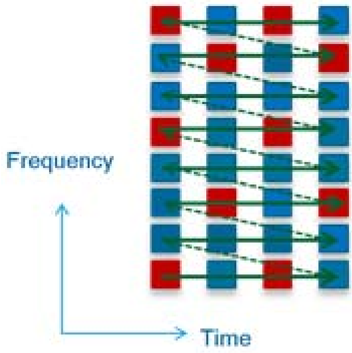

Figure 8.3: Data Mapping for a 4x8 Subslot

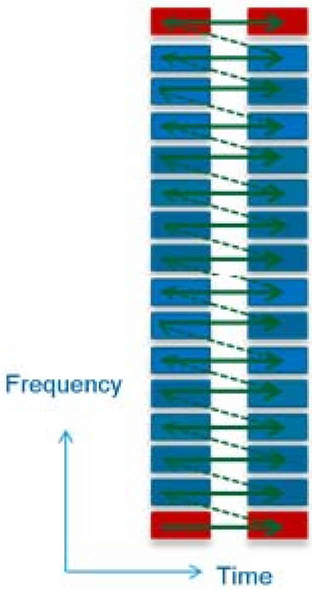

Figure 8.4: Data Mapping for a 2x16 Subslot

8.3 Downstream

8.3.1 Operation

An example implementation of the downstream convergence layer for DOCSIS 3.1 and its association with the stages before and after it is shown in figure 8.5. This block diagram is intended to demonstrate functionality; while it represents one style of implementation, there are no requirements that an implementation needs to adhere directly to this example.

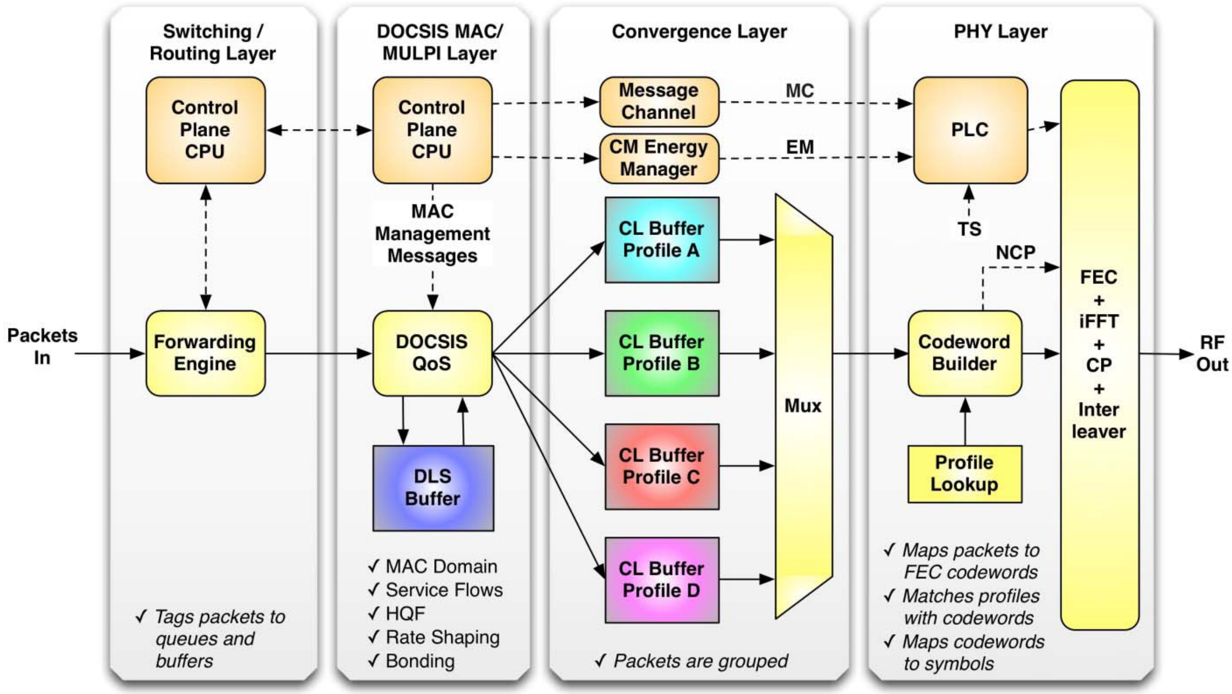

Figure 8.5: Downstream Convergence Layer Block Diagram

The operation of the downstream can be split between the forwarding plane and the control plane. The *forwarding plane* contains the data packets that are destined to the user. The *control plane* carries MAC management messages and other types of control messages.

The forwarding engine in the CMTS forwards packets to a DOCSIS MAC Domain. The MAC Domain performs QoS functions such as hierarchical QoS, per-user rate shaping, aggregate per-channel rate shaping, and aggregate per-bonding-group rate shaping. The MAC engine can also hold packets in a buffer as part of the DOCSIS Light Sleep mode.

A profile is a list of modulations that are used for the subcarriers within an OFDM channel. The downstream can use different profiles for different groups of CMs. Generally, a group of CMs that have similar SNR performance will be grouped into the same profile. When packets are encoded into FEC codewords and transmitted into the OFDM spectrum, a path in the downstream for that packet is created so the PHY layer uses a profile to create a path at the MAC layer.

There can be multiple paths from the CMTS to the same CM. Each path has a different profile. Profiles are typically given a letter designation such as Profile A. Profile A is the boot profile that CMs first begin receiving when they initialize and register. Either the forwarding engine or the DOCSIS QoS engine keeps a lookup table of exactly to what path and what profile each packet needs to be assigned. This profile assignment is used to pick a convergence layer buffer.

There is one convergence layer buffer per profile. These are shallow buffers that hold only a few packets so as to not build up any significant latency. The output of these buffers is fed to the codeword builder. The codeword builder is responsible for mapping DOCSIS frames into codewords. It is also responsible for balancing out the traffic flow between all the profiles so that the latency budgets are observed.

The codeword builder uses the same profile for an entire codeword. It can change profiles at each codeword boundary. The convergence layer buffers do not have to be serviced in any particular order. The DOCSIS MAC layer has already rate-shaped the packet flow to the size of the OFDM channel, so all packets will fit. It is up to the codeword builder to schedule the packets into the FEC codewords as it deems appropriate. Although rate shaping or packet drops are not intended to be performed at the convergence layer, some queues could be treated as low latency while other queues could be treated as high latency.

Since the codeword builder is multiplexing at the codeword level, the packets in the convergence layer buffers are naturally split across codeword boundaries and multiplexed together. The convergence layer buffers are packets in - bytes out. The codeword builder combines bytes from one buffer, adds FEC, and then using the profile modulation vector, it maps the codeword onto one or more OFDM symbols (or partial symbols).

8.3.2 MAC Frame to Codewords

The downstream LDPC codeword shown in figure 8.6 is referred to as (16 200, 14 400). This means that a full codeword is 2 025 bytes (16 200 bits) that are divided into 225 bytes (1 800 bits) of parity and 1 800 bytes (14 400 bits) of LDPC payload. That payload is further divided into 21 bytes (168 bits) of BCH parity, a 2 byte fixed header, and a variable 1 777 byte maximum payload for DOCSIS frames. When the FEC codeword is shortened, only the DOCSIS payload shrinks. All other fields remain the same size.

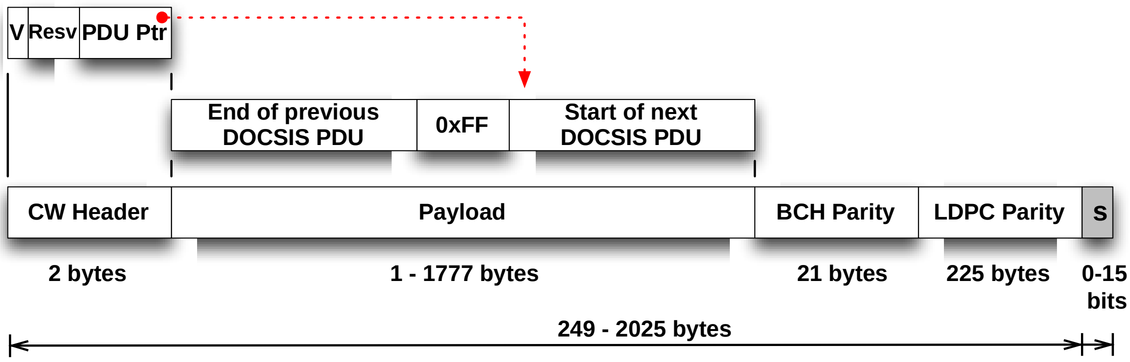

Figure 8.6: DOCSIS Frame to Codeword Mapping

Although the codeword is shown in figure 8.6 as a collection of bytes, this has to be treated as a bit stream for FEC encoding, and for subsequent mapping to OFDM subcarriers. Therefore, the CMTS shall map the stream of bytes into a stream of bits in the MSB-first order. For example, the MSB of the first byte of the codeword is to be mapped to the leftmost bit of the codeword in figure 8.6.

DOCSIS frames are sequentially mapped to codewords that are associated with a common profile. The codewords do not need to be adjacent, although their order is guaranteed. If there is no DOCSIS frame available, the CMTS shall adopt one of the following two options until the next DOCSIS frame becomes available:

1) Insert zero-bit-loaded filler sub-carriers into OFDM symbols as described in clause 7.5.5.5.2.

2) Insert stuffing pattern of 0xFF bytes into codewords.

The codeword header is defined in table 8.2.

Table 8.2: Data Codeword Definition

Name	Length	Value
Valid	1 bit	0 = PDU Pointer is not valid (ignore) 1 = PDU pointer is valid
Reserved	4 bits	Set to 0. Ignore on receive.
PDU Pointer	11 bits	This points to the first byte of the first DOCSIS frame that starts in the payload. A value of zero points to the byte immediately following the codeword header.

When the codeword gets mapped across subcarriers within a symbol, there may be residual bits left over on the last subcarrier within that symbol. Since the number of residual bits may be more or less than 8, the receiver cannot simply round down to a byte boundary. To permit the downstream receiver to discard these residual bits properly, the CMTS shall make the codeword payload an odd number of bytes.

One potential set of algorithms for the CMTS codeword builder and the CM codeword receiver is as follows.

On the CMTS, the algorithm is:

- Number of total bytes is (header + payload + parity) and never exceeds 2 025 bytes.
- IF (total bytes = odd), send to FEC engine.
- IF (total bytes = even), add a0xFF stuff byte to the payload if legal or change the number of bytes.

- CMTS Symbol mapper adds trailing bits (all 1s) to map codeword to a symbol boundary.

On the CM, the corresponding algorithm is:

- CM extracts total bits between two NCP pointers.
- IF total bits > 16 200, use initial 16 200 bits, and a full codeword is declared.
- IF total bits = 16 200, and a full codeword is declared.
- IF total bits < 16 200, round down to the nearest odd number of bytes.
- Discard [(total bits + 8) Modulo 16] bits.

8.3.3 Subcarrier Numbering Conventions

Subcarriers are numbered from lower frequency to higher frequency within a FFT block. All subcarriers within the 204,8 MHz bandwidth are numbered, including the outside excluded subcarriers. Numbering starts at 0 and goes to 4 095 for 4K FFT, and 0 to 8 191 for 8K FFT.

Data codewords are mapped to subcarriers - prior to time and frequency interleaving - from a lower number to a higher number.

8.3.4 Next Codeword Pointer

8.3.4.0 NCP Overview

When the data codewords are mapped to subcarriers within a symbol, a pointer is needed to identify where a data codeword starts. This is known as the Next Codeword Pointer (NCP). The collection of NCP message blocks within a symbol is known as the NCP field. There are a variable number of NCP message blocks (MBs) on each OFDM symbol. To make sure that all subcarriers are used without reserving empty NCP MBs, the mapping of the NCP occurs in the opposite direction of the mapping for data. The relationship of NCP message blocks to the data channel is shown in figure 8.7 (scattered pilots are not shown; last NCP MB is a CRC).

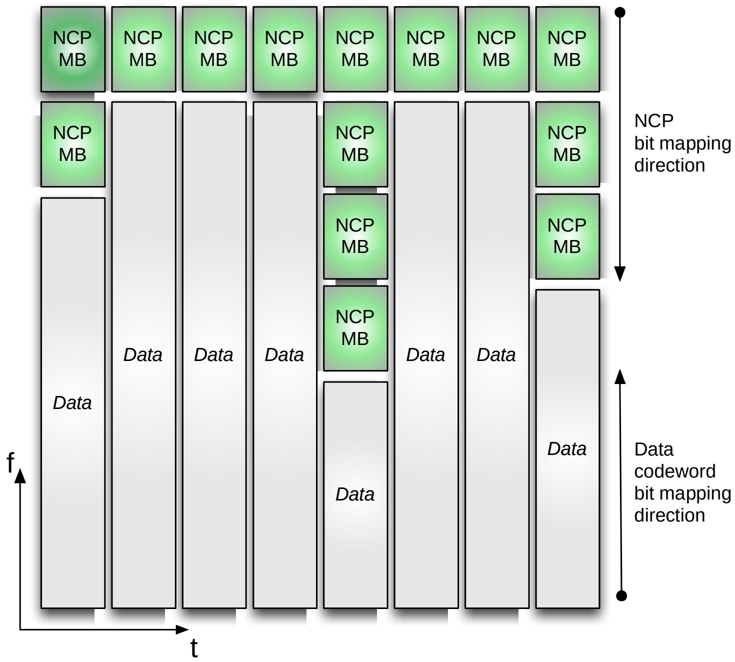

Figure 8.7: Data and NCP Prior to Interleaving

The CMTS shall map data subcarriers within a symbol, starting from a lower subcarrier number and proceeding to a higher subcarrier number.

The CMTS shall map NCP message blocks starting at a higher subcarrier number and moving to a lower subcarrier number.

8.3.4.1 NCP Data Message Block

The format of the NCP data message block is illustrated in figure 8.8 and defined in table 8.3.

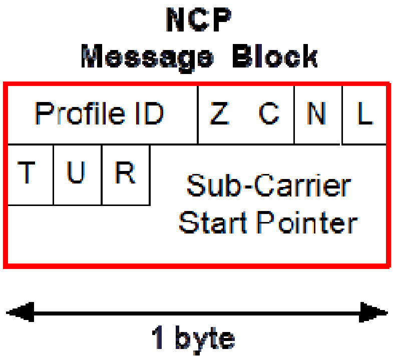

Figure 8.8: NCP Data Message Block

Note that each three byte NCP MB is mapped into a unique FEC codeword that has a 3 byte payload with 3 bytes of FEC. The last FEC codeword is then followed by a 3 byte CRC-24-D (refer to Annex E) that is also placed in its own FEC block.

Table 8.3: NCP Parameters

Field	Size	Description
Profile ID	4 bits	Profile ID for the data channel 0 = Profile A 1 = Profile B … 15 = Profile P
Z	1 bit	Zero Bit Loading 0 = subcarriers follow profile 1 = subcarriers are all zero-bit-loaded
C	1 bit	Data Profile Update 0 = use even profile 1 = use odd profile
N	1 bit	NCP Profile Select 0 = use even profile 1 = use odd profile This bit is equal to the LSB of the NCP profile change count.
L	1 bit	Last NCP Block 0 = This NCP is followed by another NCP. 1 = This is the last NCP in the chain and is followed by an NCP CRC message block.
T	1 bit	Codeword Tagging 0 = This codeword is not included in the codeword counts reported by the CM in the OPT-RSP message. 1 = This codeword is included in the codeword counts reported by the CM in the OPT-RSP message. This bit is applicable only when Codeword Tagging is enabled in the CM for the CM's transition profile. When Codeword Tagging is not enabled for the CM's transition profile, all codewords are counted. If the CM is not conducting OFDM Downstream Profile Usability Testing on the Profile ID of this NCP, or if the CM is conducting testing but Codeword Tagging is disabled for the test, then the Codeword Tagging bit is ignored and all codewords are counted. Codeword Tagging is enabled or disabled in the CM for the CM's transition profile by the CMTS through a parameter passed in the OPT-REQ message. Codeword Tagging enable/disable also applies per profile. Codeword Tagging can only be enabled for a CM's transition profile. See [4] for a more details about OFDM Downstream Profile Testing (OPT) and Codeword Tagging.
U	1 bit	NCP Profile Update indicator Indicates a change in the NCP bit-loading profile. The CMTS sets this bit in each of the 128 symbols immediately preceding an NCP bit-loading profile change. The 128 sequential "U" bits form a specific bit pattern as defined below to indicate the NCP profile change. The CMTS sets NCP Profile Update indicator to value 0 in all other symbols.
R	1 bit	Reserved
Subcarrier pointer	13 bits	This is the number assigned to the first subcarrier used by the codeword. The maximum value is 0x1FFE = 8 190. The value 0x1FFF is reserved as a null pointer.

The NCP structure is predicated upon the following facts:

- FEC codewords are mapped continuously across successive symbols.
- The PHY can determine the first subcarrier of the first NCP message block.
- The PHY can determine the first subcarrier of the data field in the current symbol.

Based upon these facts and combined with the information in the NCP fields, then:

- The PHY can determine the last subcarrier of the last NCP message block.
- The next subcarrier after the last NCP message block CRC is last subcarrier of the data field.

The main task of the NCP message block is to provide a reference to the appropriate profile and a start pointer for codewords. The length of a codeword is determined by the difference between the subcarrier pointer in two successive NCP message blocks.

Data subcarriers may contain FEC codewords or unused sub-carriers. These functions are referred to as fields in the NCP header.

The CMTS shall include one NCP within the same symbol for each start of codeword or a group of unused subcarriers that exists in that symbol. The CMTS shall include a valid Profile ID when the field is a FEC codeword field and no zero-bit-loading (Z bit not asserted).

The CMTS shall assert the Zero load bit ("Z") to mark a set of subcarriers which are not used as described in clause 8.3.4.2.

The CM shall ignore the Profile ID when the Z bit is asserted.

The CM shall use the Data Profile Update bit ("C") to select the odd or even data profile.

The CMTS shall set the value of Data Profile Update bit ("C") to indicate whether the odd or even data profile is in use in the current OFDM symbol.

The CM shall use the NCP Profile Select bit ("N") to select the odd or even NCP profile in the current OFDM symbol.

The CMTS shall set the NCP Profile Select bit to the same value in all the NCP Message Blocks in any one symbol.

The CMTS shall use the NCP Profile Update bit ("U") to indicate a change in the NCP bit-loading profile.

The CMTS shall set the NCP Profile Update bit to the same value in all NCP Message Blocks in any one symbol.

The CMTS shall set the U-bits of the 128 symbols immediately preceding a NCP bit-loading profile change to the pattern Hex "BCB240898BAD833539ED0ABE946E3F85", as illustrated in figure 8.9.

The CMTS shall set the NCP Profile Update Bit to zero in all other symbols, i.e. all symbols except the 128 symbols immediately preceding a NCP bit-loading profile change.

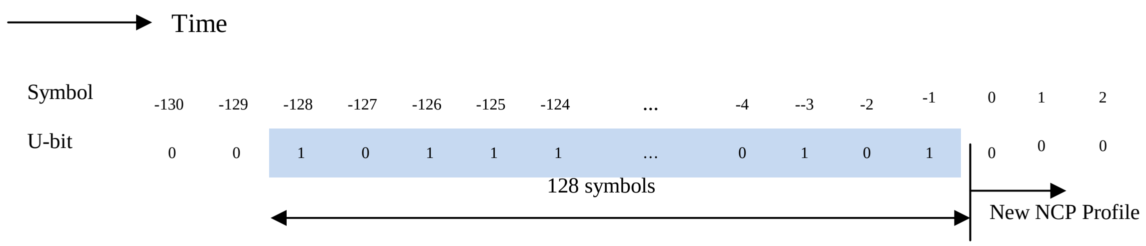

Figure 8.9: NCP Profile Update Bit Setting Immediately Preceding an NCP Bit Loading Profile Change

The CMTS shall assert the Last bit ("L") if the NCP is the last NCP message block.

The CMTS shall follow the NCP block that has its Last bit asserted with a CRC-24-D. The CRC-24-D is calculated across all message blocks in a symbol exclusive of the FEC parity bits.

When a Downstream Profile Usability test is in progress with Codeword Tagging enabled, the CMTS assigns the value of the Codeword Tagging bit ("T") to indicate whether the codeword is to be included in the codeword counts reported in the OPT-REQ message [4].

When a Downstream Profile Usability test is in progress and Codeword Tagging is enabled by the OPT-REQ message for the CM's test profile, the CM is required to respect the Codeword Tagging ("T") bit [4].

A NULL NCP is defined as an NCP with the start pointer set to 0x1FFF. The usage of Null NCPs is defined in the next section. An Active NCP is an NCP that points to valid FEC codeword. Therefore, an Active NCP is an NCP in which Z-bit is Zero and the Start Pointer is not equal to 0x1FFF.

8.3.4.2 NCP Field with CRC and FEC

The last NCP message block is a dedicated CRC block. This block is shown in figure 8.10.

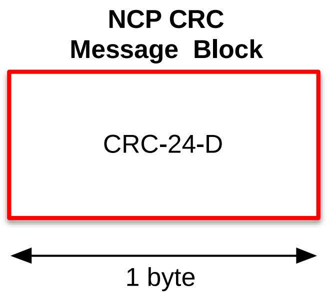

Figure 8.10: NCP CRC Message Block

The CMTS shall include a NCP CRC message block based upon the CRC-24-D calculation after the last NCP data message.

The CM shall check the NCP CRC field.

If the NCP CRC field indicates an error in the NCP field, then the CM shall reject all NCP data message blocks in the NCP field of the current symbol.

The CRC-24-D, defined in Annex A, is applicable to a bit stream. Hence to work out this CRC the CMTS has to first map the NCP Message Blocks into a bit-stream. The CMTS shall implement this conversion to bit-serial format in MSB-first order. That is, the CMTS shall place the first bit of the bit stream as the leftmost bit of the Profile ID shown in figure 8.8, of the topmost NCP Message Block shown in figure 8.7.

NCP has a specific method of mapping NCP message blocks with FEC that is different than the FEC used on the main OFDM data channels. The complete NCP field with FEC parity is shown in figure 8.11.

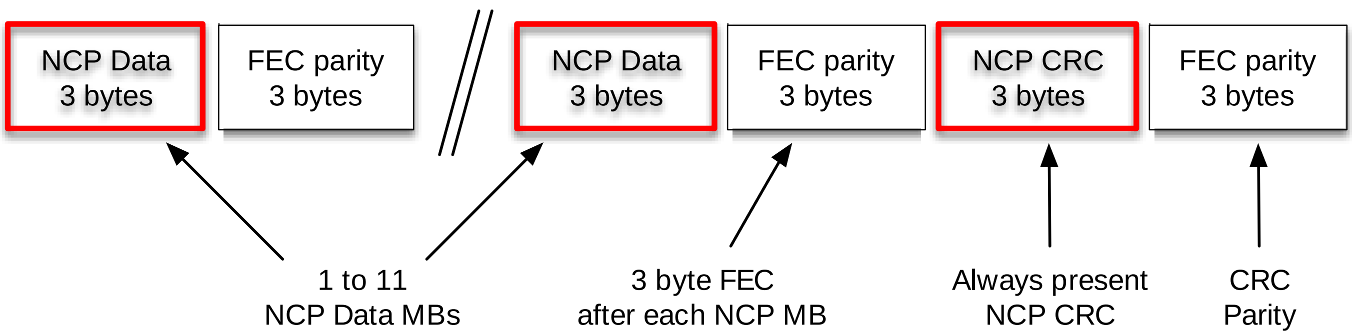

Figure 8.11: NCP Message Blocks Field with FEC

NOTE: Each three byte NCP MB is mapped into a unique FEC codeword that has a 3 byte payload with 3 bytes of FEC. The last FEC codeword is then followed by a 3 byte CRC-24-D (refer to Annex E) that is also placed in its own FEC block.

8.3.4.3 NCP Usage

The CMTS shall not place more than 11 NCP data message blocks plus a CRC for a total of 12 NCP MBs in an 8K OFDM symbol. The CMTS shall not place more than 12 NCP data message blocks plus two CRCs for a total of 14 NCP MBs in any two successive 4K OFDM symbols.

In the case of an 8K FFT OFDM symbol, the 12 NCP MBs will be formed by a maximum of 10 active NCP MBs, the NULL or zero-bit-loaded NCP MB (i.e. NCP with Z-bit set to ONE) and the NCP CRC MB.

In the case of a 4K FFT, in addition to the 10 maximum active NCP MBs over two successive symbols, each of these symbols may have one additional NULL or zero-bit-loaded NCP MB, and each symbol will have a NCP CRC MB. This brings the maximum number of NCP MBs over two successive symbols to 14.

If the data FEC blocks are small in one 8K FFT OFDM symbol, there could be data subcarriers left in the symbol after the placement of 10 active NCPs and the corresponding data. In such a case the CMTS shall include an NCP describing the remaining subcarriers as zero-bit-loaded ("Z" bit asserted).

The CMTS shall not place more than 10 active NCPs in any two consecutive 4K OFDM symbols, i.e. the number of active NCPs in 4K FFT OFDM symbols *n* and *n+1* shall not exceed 10, for any value of *n*.

If the data FEC blocks are small, there could be data subcarriers remaining and unused after the placement of 10 active NCPs and the corresponding data in two consecutive 4K OFDM symbols. In such a case, the CMTS shall include an NCP describing the remaining subcarriers as zero-bit-loaded ("Z" bit asserted). Furthermore, in the case of 4K FFT, if all of the 10 active NCPs are consumed by the symbol *n*, then the CMTS shall place an NCP in symbol *n+1* indicating that the unused subcarriers are zero-bit-loaded. The symbol *n+1* may contain a continuation of a codeword from symbol *n*, but no new codeword can start in symbol *n+1*.

For small bandwidths it is possible that there may not be a beginning or an end of a FEC codeword in a symbol. That is, a codeword may begin in the previous symbol and end in the following symbol. In such a case the CMTS shall insert a NULL NCP in the current symbol.

There may also be scenarios in which a FEC codeword may end within a symbol without leaving sufficient space to include an NCP. In this case, the CMTS shall insert a NULL NCP and move some of the data subcarriers of the FEC codeword to next OFDM symbol.

8.3.4.4 NCP Example

Figure 8.12 shows some examples of how the NCP field is used. This view is prior to interleaving.

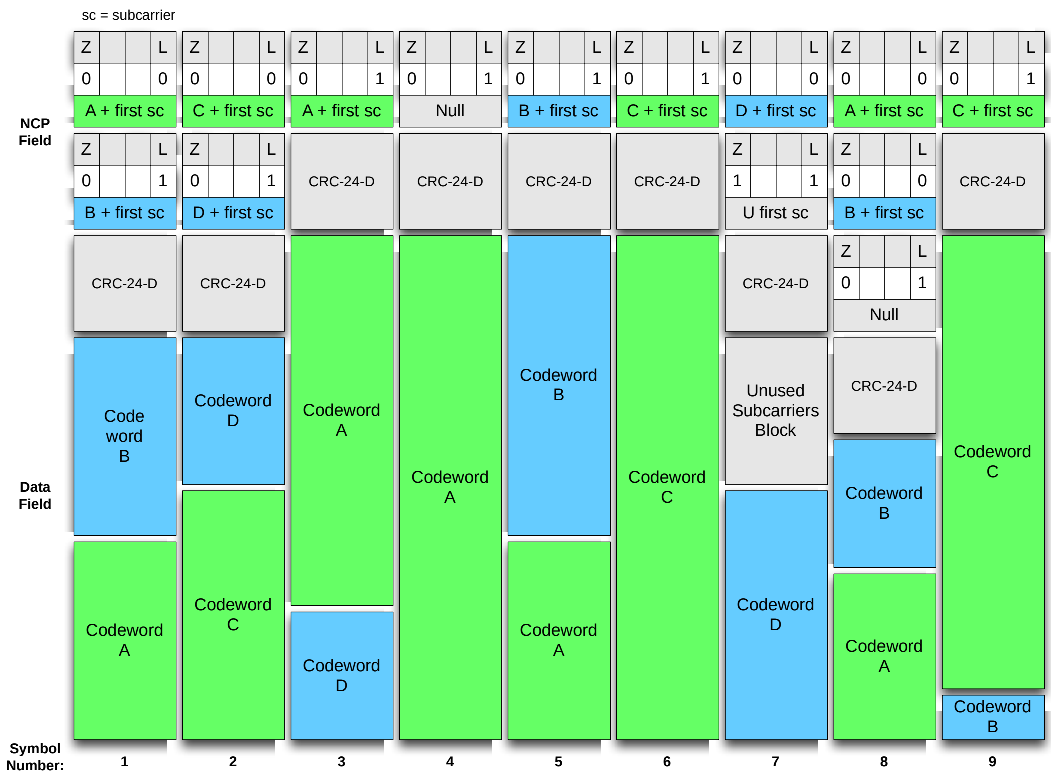

Figure 8.12: NCP Examples

NCP blocks are mapped to sub-carriers starting with the first non-excluded subcarrier at the top of the spectrum and then down in frequency. After the last NCP MB is a CRC-24-D. Data is mapped to the first non-excluded subcarrier at the bottom of the frequency range and then continuing upwards in frequency.

In symbol 1, Codeword A starts at the beginning of the symbol and has a start pointer. Codeword B starts after codeword A and has a start pointer. The length of codeword A is the difference between the codeword A start pointer and the codeword B start pointer.

In symbol 2, Codeword C starts at the beginning of the symbol and has a start pointer. The length of the previous codeword B is derived from the difference between the codeword B start pointer and the codeword C start pointer, taking into account where the last data subcarrier was in symbol 1. Codeword D gets a start pointer.

In symbol 3, Codeword D continues from symbol 2 and finishes. Codeword A follows and is given a start pointer. The length of codeword D is derived from the difference between the codeword C start pointer and the codeword D start pointer, taking into account where the last data subcarrier was in symbol 2.

In symbol 4, Codeword A continues. Since there is no start pointer required, but at least one NCP block is required, an NCP block with a null pointer is included.

In symbol 5, Codeword A ends. Codeword B begins and ends. A single NCP block is created with a start pointer to codeword B.

In symbol 6, Codeword C both starts and ends. A single NCP block is created with a start pointer to codeword C.

In symbol 7, Codeword D starts and ends. There are no more data packets to send, so the remaining subcarriers are unused. A NCP block is assigned for the codeword D start pointer. A second NCP block is assigned to the start pointer of the unused subcarriers. This start pointer is used to determine the length of codeword D.

In symbol 8, Codeword A begins and ends. Codeword B begins and tried to end with a few subcarriers unused between the end of the data codeword and the end of the NCP field. Since no subcarriers can be left unused, and since an NCP would not fit, an NCP with a null pointer was inserted and some of the last few bytes of codeword B were forced into the next symbol. There is an NCP message block for codeword A, codeword B, and the null NCP.

In symbol 9, Codeword C starts a few subcarriers into the symbol. There is one NCP block for codeword C.

9 Proactive Network Maintenance

9.1 Scope

This clause defines the requirements supporting Proactive Network Maintenance (PNM). CMTS and cable modem features and capabilities can be leveraged to enable measurement and reporting of network conditions such that undesired impacts such as plant equipment and cable faults, interference from other systems and ingress can be detected and measured. With this information cable network operations personnel can make modifications necessary to improve conditions and monitor network trends to detect when network improvements are needed.

9.2 System Description

As shown in figure 9.1, the CMTS and CM contain test points which include essential functions of a spectrum analyzer, vector signal analyzer (VSA), and network analyzer, while the cable plant is considered the Device Under Test (DUT). The goal is to rapidly and accurately characterize, maintain and troubleshoot the upstream and downstream cable plant, in order to guarantee the highest throughput and reliability of service. The CMTS and CM make the specified measurements and report the results to the PNM management entity as defined in [i.2] and [i.3]. Unless otherwise specified, the CM shall make all its PNM measurements while in service, without suspending normal operational modes or data transmission and reception. Unless otherwise specified, the CMTS shall make all its PNM measurements while in service, without suspending normal operational modes or data transmission and reception.

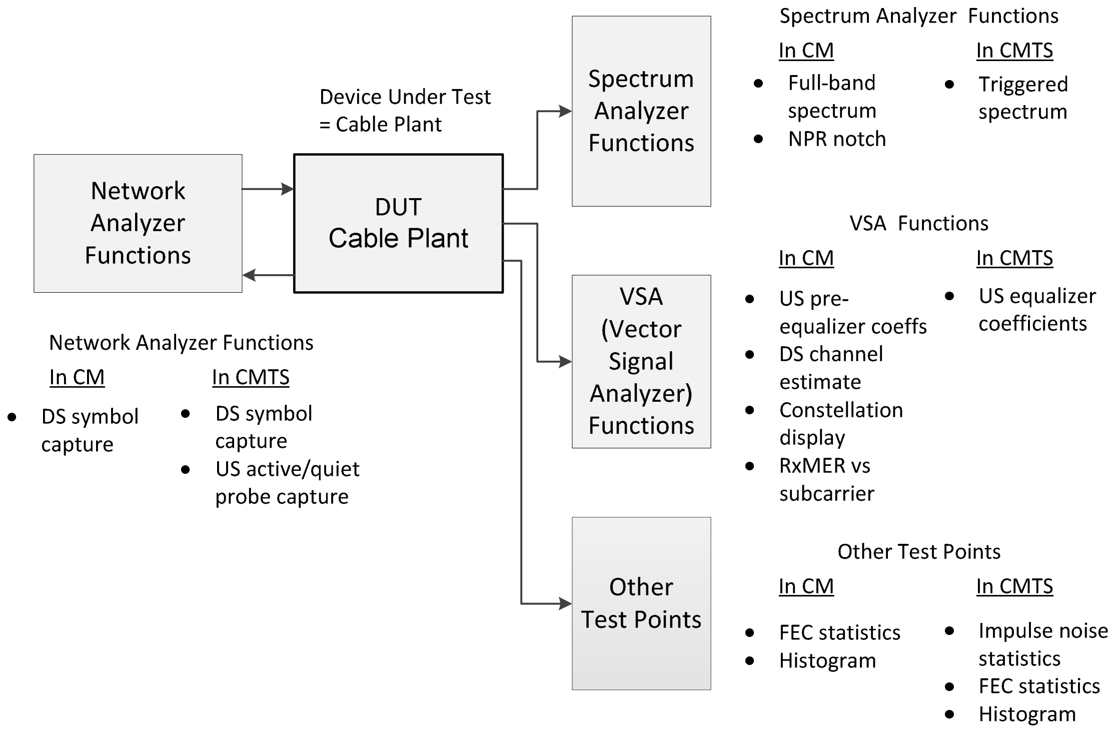

Figure 9.1: Test Points in CM and CMTS Supporting Proactive Network Maintenance

9.3 Downstream PNM Requirements

9.3.1 Downstream Symbol Capture

The purpose of downstream symbol capture is to provide partial functionality of a network analyzer to analyze the response of the cable plant.

At the CMTS, the transmitted frequency-domain modulation values of one full OFDM symbol before the IFFT are captured and made available for analysis. This includes the I and Q modulation values of all subcarriers in the active bandwidth of the OFDM channel, including data subcarriers, pilots, PLC preamble symbols and excluded subcarriers. This capture will result in a number of samples that depends on the OFDM channel width, per clause 7.5.7.1. As examples, for 50 kHz subcarrier spacing in a 192 MHz channel with an active bandwidth of 190 MHz, 3 800 samples will be captured; for 25 kHz subcarrier spacing in a 192 MHz channel with an active bandwidth of 190 MHz, 7 600 samples will be captured; for 25 kHz subcarrier spacing in a 24 MHz channel with an active bandwidth of 22 MHz, 880 samples will be captured.

At the CM, the received I and Q time-domain samples of one full OFDM symbol before the FFT, not including the guard interval, are captured and made available for analysis. This capture will result in a number of data points equal to the FFT length in use, time aligned for receiver FFT processing. The number of captured samples can be reduced for narrower channels if the sampling rate, which is implementation dependent, is reduced. The capture includes a bit indicating if receiver windowing effects are present in the data.

As examples, for 50 kHz subcarrier spacing in a 192 MHz channel with 204,8 MHz sampling rate, 4 096 samples will be captured; for 25 kHz subcarrier spacing in a 192 MHz channel with 204,8 MHz sampling rate, 8 192 samples will be captured; for 50 kHz subcarrier spacing in a 24 MHz channel with a reduced sampling rate of 25,6 MHz, 512 samples will be captured.

Capturing the input and output of the cable plant is equivalent to a wideband sweep of the channel, which permits full characterization of the linear and nonlinear response of the downstream plant. The MAC provides signalling via the PLC Trigger Message to ensure that the same symbol is captured at the CMTS and CM.

The CMTS shall be capable of capturing the modulation values of the full downstream symbol marked by the trigger for analysis.

The CM shall be capable of locating and capturing the time-domain samples of the full downstream symbol marked by the trigger for analysis.

9.3.2 Downstream Wideband Spectrum Analysis

The purpose of downstream wideband spectrum capture is to provide a downstream wideband spectrum analyzer function in the DOCSIS 3.1 CM similar to the capability provided in [i.2] and [i.3].

The CM shall provide a downstream wideband spectrum capture and analysis capability.

The CM should provide the capability to capture and analyze the full downstream band of the cable plant.

The CM shall provide a calibration constant permitting the downstream wideband spectrum capture measurement to be related to the downstream received power measurement of clause 9.3.9.

9.3.3 Downstream Noise Power Ratio (NPR) Measurement

The purpose of downstream NPR measurement is to view the noise, interference and intermodulation products underlying a portion of the OFDM signal. As part of its normal operation or in an out-of-service test, the CMTS can define an exclusion band of zero-valued subcarriers which forms a spectral notch in the downstream OFDM signal for all profiles of a given downstream channel. The CM provides its normal spectral capture measurements per 9.3.2, or symbol capture per 9.3.1, which permit analysis of the notch depth. A possible use case is to observe LTE interference occurring within an OFDM band; another is to observe intermodulation products resulting from signal-level alignment issues. Since the introduction and removal of a notch affects all profiles, causing possible link downtime, this measurement is intended for infrequent maintenance.

9.3.4 Downstream Channel Estimate Coefficients

The purpose of this item is for the CM to report its estimate of the downstream channel response. The reciprocals of the channel response coefficients are typically used by the CM as its frequency-domain downstream equalizer coefficients. The channel estimate consists of a single complex value per subcarrier. [i.2] defines summary metrics to avoid having to send all coefficients on every query.

The CM shall report its downstream channel estimate (full set or summary) for any single OFDM downstream channel upon request.

9.3.5 Downstream Constellation Display

The downstream constellation display provides received QAM constellation points for display. Equalized soft decisions (I and Q) at the slicer input are collected over time, possibly subsampling to reduce complexity, and made available for analysis. Only data-bearing subcarriers with the specified QAM constellation are sampled. Pilots and excluded subcarriers within the range are ignored. Up to 8 192 samples are provided for each query; additional queries may be made to further fill in the plot.

The CM shall be capable of capturing and reporting received soft-decision samples, for a single selected constellation from the set of profiles it is receiving within a single OFDM downstream channel.

9.3.6 Downstream Receive Modulation Error Ratio (RxMER) per Subcarrier

9.3.6.0 RxMER per Subcarrier Overview

This item provides measurements of the receive modulation error ratio (RxMER) for each subcarrier. The CM measures the RxMER using pilots and PLC preamble symbols, which are not subject to symbol errors as data subcarriers would be. Since scattered pilots visit all data subcarriers and the PLC preamble symbols are known, the RxMER of all active subcarriers in the OFDM band can be measured over time. For the purposes of this measurement, RxMER is defined as the ratio of the average power of the ideal QAM constellation to the average error-vector power. The error vector is the difference between the equalized received pilot or preamble value and the known correct pilot value or preamble value. As a defining test case, for an ideal AWGN channel, an OFDM block containing a mix of QAM constellations, with 35 dB CNR on the QAM subcarriers, will yield an RxMER measurement of nominally 35 dB for all subcarrier locations. If some subcarriers (such as exclusion bands) cannot be measured by the CM, the CM indicates that condition in the measurement data for those subcarriers.

RxMER may be more clearly defined in mathematical notation in accordance with figure 9.2, which shows an ideal transmit and receive model, with no intent to imply an implementation. Let p = scattered pilot (or PLC preamble) symbol before transmit IFFT, H = channel coefficient for a given subcarrier frequency, n = noise, y = Hp + n = unequalized received symbol after receive FFT. The receiver computes G = estimate of H, and computes the equalized received symbol as r = y/G. Using the known modulation value of the pilot or preamble symbol p, the receiver computes the equalized error vector as e = r - p. All the above quantities are complex scalars for a given subcarrier. To compute RxMER, the receiver computes E = time average of |e|^2 over many visits of the scattered pilot to the given subcarrier (or PLC preamble symbol as applicable), and E_dB = 10*$\log_{10}$(E). Let S_dB = average power of ideal QAM constellation (not including pilots) expressed in dB. According to Annex A.2, QAM Constellation Scaling, all QAM constellations have the same average power. The CM reports RxMER_dB = S_dB - E_dB.

The CM shall be capable of providing measurements of RxMER for all active subcarrier locations for a single OFDM downstream channel, using pilots and PLC preamble symbols.

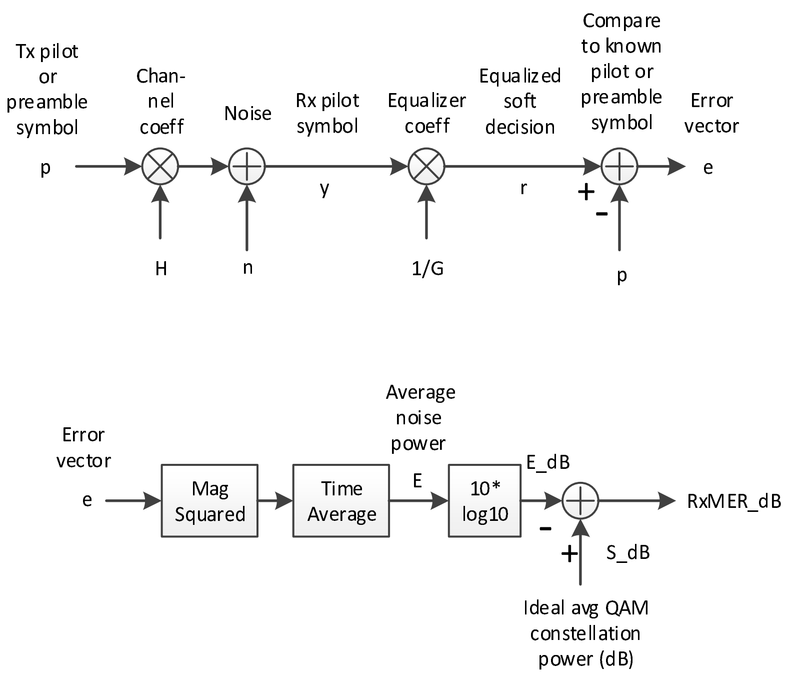

Figure 9.2: Computation of Received Modulation Error Ratio (RxMER) for a Given Subcarrier

9.3.6.1 Signal-to-Noise Ratio (SNR) Margin for Candidate Profile

The purpose of this item is to provide an estimate of the SNR margin available on the downstream data channel with respect to a candidate modulation profile. The CMTS has the capability described in [4] clause 7.8.1, CM and CMTS Profile Support, in which it sends test data to the CM to measure the performance of a transition profile. In addition, the CM shall implement an algorithm to estimate the SNR margin available on the downstream data channel for a candidate profile. Annex J suggests an algorithm that the CM can use to compute this estimate. The CM only performs this computation upon request from the CMTS via management message.

9.3.7 Downstream FEC Statistics

The purpose of this item is to monitor downstream link quality via FEC and related statistics. Statistics are taken on FEC codeword error events, taking into account both the inner LDPC code and outer BCH code. Statistics are provided on each OFDM channel and for each profile being received by the CM. That is, if the CM is receiving 4 downstream profiles, there will be 4 sets of FEC counters plus a set of counters for the transition profile used for OFDM Downstream Profile Test (OPT) (see MULPIv3.1). For profiles 1-4, statistics for data codewords include all codewords. For profile 5 (transition profile), statistics for data codewords include either all codewords, if Codeword Tagging [DOCSIS MULPIv3.1] is disabled; or only codewords marked with T bit = 1 in the NCP, if Codeword Tagging is enabled. Similar statistics are taken on the NCP and PLC, and on MAC frames.

The CM shall provide the start and end time of the measurement period for all measurements. The measurements are timestamped using bits 21-52 of the 64-bit extended timestamp, where bit 0 is the LSB, which provides a 32-bit timestamp value with resolution of 0,4 ms and range of 20 days. Timestamping is done with nominal accuracy of 100 ms or better.

The CM shall be capable of providing the following downstream performance metrics on data codewords for each profile:

- Uncorrectables: Number of codewords that failed BCH decoding.
- Correctables: Number of codewords that failed pre-decoding LDPC syndrome check and passed BCH decoding.
- Total number of FEC codewords.

The CM shall be capable of providing the following downstream performance metrics on Next Codeword Pointer (NCP) codewords and data fields:

- Unreliable NCP Codewords: Number of NCP codewords that failed LDPC post-decoding syndrome check.
- NCP CRC failures: Number of NCP fields that failed CRC check.
- Total number of NCP codewords.
- Total number of NCP fields.

The CM shall be capable of providing the following downstream performance metrics on PHY Link Channel (PLC) codewords:

- Unreliable PLC Codewords: Number of PLC codewords that failed LDPC post-decoding syndrome check.
- Total number of PLC codewords.

The CM shall be capable of providing the following downstream performance metrics on Media Access Control (MAC) frames addressed to the CM for each profile excluding the transition profile:

- MAC frame failures: Number of frames that failed MAC CRC check.
- Total number of MAC frames.

The CM shall be capable of providing the following downstream FEC summaries for data codewords on each OFDM channel for each profile being received by the CM:

- Codeword error ratio vs. time (seconds): Ratio of number of uncorrectable codewords to total number of codewords in each one-second interval for a rolling 10-minute period (600 values).
- Codeword error ratio vs. time (minutes): Ratio of number of uncorrectable codewords to total number of full-length codewords in each one-minute interval for a rolling 24-hour period (1 440 values).
- Start and end time of rolling period.
- Red/yellow/green summary link status (colors defined in [i.3]).

The CM shall provide two collection and reporting methods for each error-count metric:

- Long-term statistics. The CM always collects metrics in the background for each profile being received. The codeword (or frame) and error counters are long (e.g. 64-bit) integers, so that overflow is not an issue. ' To perform a measurement over a particular time interval, the user reads the counters, waits a period of time, reads the counters again, and computes the difference in the counter values.
- Short-term statistics. The CM performs a one-shot measurement with two configured parameters, N_e and N_c. The CM reports the results when at least N_e errors have occurred or at least N_c codewords have been processed, whichever comes first. This measurement is particularly useful for OPT testing of the transition profile. To perform this measurement, the CM reads the long-term counters, waits a short time, reads the counters again, and computes the difference in the counter values.

9.3.8 Downstream Histogram

The purpose of the downstream histogram is to provide a measurement of nonlinear effects in the channel such as amplifier compression and laser clipping. For example, laser clipping causes one tail of the histogram to be truncated and replaced with a spike.

The CM shall be capable of capturing the histogram of time domain samples at the wideband front end of the receiver (full downstream band). When a CM creates a downstream histogram, the CM shall create it such that it is two-sided; that is, it encompasses values from far-negative to far-positive values of the samples. When the CM creates a downstream histogram, the CM shall create it such that it has either 256 equally spaced bins with even symmetry about the origin, or 255 equally spaced bins with odd symmetry about the origin. These bins typically correspond to the 8 MSBs of the wideband analog-to-digital converter (ADC).

The histogram dwell count, a 32-bit unsigned integer, is the number of samples observed while counting hits for a given bin, and may have the same value for all bins. The histogram hit count, a 32-bit unsigned integer, is the number of samples falling in a given bin. The CM shall be capable of reporting the dwell count per bin and the hit count per bin. When enabled, the CM shall compute a histogram with a dwell of at least 10 million samples in 30 seconds or less. With this many samples, the histogram can reliably measure a probability density per bin as low as 10^{-6} with at least 10 hits in each bin. The CM shall continue accumulating histogram samples until it is restarted, disabled, approaches its 32-bit overflow value, or times out. The CM shall report the start and end time of the histogram measurement using bits 21 - 52 of the extended timestamp, which provides a 32-bit timestamp value with resolution of 0,4 ms and range of 20 days.

9.3.9 Downstream Received Power

The purpose of the downstream received power metric is to measure the average received downstream power in a set of non-overlapping 6 MHz bands for any DOCSIS 3.0 and 3.1 signals in the receive channel set (RCS) of the CM including the DOCSIS 3.1 PLC.

While digital power measurements are inherently accurate, the measurement referred to the analog signal at the input F connector depends on the measurement conditions and available calibration accuracy. The measurements are made along a contiguous set of defined 6 MHz bands. The measurement bands are designed to align with DOCSIS 3.0 channel locations, so that the total power of a DOCSIS 3.0 single carrier QAM signal in a 6 MHz bandwidth is measured.

For DOCSIS 3.1 OFDM signals, the measurement bands will also align with the edges of the occupied spectrum of the OFDM channel, although the measurement bands may not align with the edges of the modulated spectrum of the OFDM channel. In general, a DOCSIS 3.1 OFDM signal may contain excluded subcarriers within a given 6 MHz measurement band. However, the 6 MHz band containing the PLC at its centre is a special case which contains no excluded subcarriers and contains extra pilots, properties which make it useful for reliable power measurements. The PLC measurement band may be offset from the 6 MHz measurement bands due to the location of subcarriers in the PLC placement where the centre frequency of the lowest frequency subcarrier of the 6 MHz encompassed spectrum containing the PLC is an integer when the frequency is measured in units of MHz.

The CM shall provide an estimate of the average power for any 6 MHz band in the RCS referenced to the F connector input of the CM under the following specified received signal conditions:

The measurement band is defined as either of the following:

- Any 6 MHz bandwidth with a centre frequency of 111 + 6(n-1) MHz for n = 1, …, 185 (i.e. 111, 117, …, 1 215 MHz) contained in the RCS of the CM.
- The 6 MHz bandwidth containing the PLC with a lower channel edge (centre frequency of lowest subcarrier) frequency of 108 + m MHz for m = 0, 1, …, 1 104 (i.e. 108, 109, …, 1 212 MHz).

All measurements are made under the following conditions:

- The measurement bands shall contain only DOCSIS signals.
- The measured bands do not contain any gaps (regions with no signal present) greater than 24 MHz wide each from a 108 MHz or 258 MHz lower band edge to a 1 002 MHz or 1 218 MHz upper band edge.
- A constant temperature is maintained during measurements within a range of 20 ºC ± 2 ºC.
- A minimum warm up time of 30 minutes occurs before CM power measurements are made.
- The measured 6 MHz band does not contain gaps totalling more than 20 percent of that band.
- A maximum 8 dB range of signal input power variation can be measured without allowing recalibration via CM re-initialization. (Per [15], the maximum signal amplitude level variation at the end of the drop connected to the subscriber tap over any 6 month interval is limited to 8 dB.)
- The signal power variation (i.e. the deviation from nominal relative carrier power level) between the defined 6 MHz bands containing DOCSIS signals varies up to a maximum of ±3 dB over the full downstream band.
- The signal tilt and signal power variation are determined using 6 MHz channel power levels as specified in clause 6.5 of [20] - Nominal Relative Carrier Power Levels and Carrier Level Variations.
- The total spectrum of all the gaps is to be less than 20 % of the total encompassed spectrum.

The CM shall provide an average power estimate in any defined 6 MHz measurement bandwidth within ± 3 dB of the actual power at the F connector input under the following conditions:

- The power in the defined 6 MHz bands has a maximum tilt of ±1 dB over the entire downstream spectrum.
- The power in the defined 6 MHz bands (other than gaps) is in the range of -12 dBmV to +12 dBmV.

The CM shall provide an average power estimate in any defined 6 MHz measurement band within ± 5 dB of the actual power at the F connector input under the following relaxed signal conditions:

- The power in the defined 6 MHz measurement bands has a maximum upward tilt of +4 dB and a maximum downward tilt of -9 dB over the entire downstream band.
- The power in the defined 6 MHz bands (other than gaps) is in the range of -15 dBmV to +15 dBmV.

9.4 Upstream PNM Requirements

9.4.1 Upstream Capture for Active and Quiet Probe

The purpose of upstream capture is to measure plant response and view the underlying noise floor, by capturing at least one OFDM symbol during a scheduled active or quiet probe. An active probe provides the partial functionality of a network analyzer, since the input is known and the output is captured. This permits full characterization of the linear and nonlinear response of the upstream cable plant. A quiet probe provides an opportunity to view the underlying noise and ingress while no traffic is being transmitted in the OFDMA band being measured.

The PNM server selects an active CM to analyze by specifying its MAC address, or requests a quiet probe measurement. The CMTS shall be capable of selecting a specified transmitting CM, or quiet period when no CMs are transmitting, for the capture. The CMTS sets up the capture as described in [4], selecting either an active SID corresponding to the specified MAC address or the idle SID, and defining an active or quiet probe. The active probe symbol for this capture normally includes all non-excluded subcarriers across the upstream OFDMA channel, and normally has pre-equalization off. The quiet probe symbol normally includes all subcarriers, that is, during the quiet probe time there are no transmissions in the given upstream OFDMA channel. The CMTS shall capture samples of a full OFDMA symbol including the guard interval. The CMTS shall begin the capture with the first symbol of the specified probe. The sample rate is the FFT sample rate (102,4 Msym/s).

The CMTS shall report the list of excluded subcarriers, cyclic prefix length, and transmit window roll-off period in order to fully define the transmitted waveform. The CMTS shall report the index of the starting sample used by the receiver for its FFT. The CMTS shall report the timestamp corresponding to the beginning of the probe. In the case where the P-MAPs for the OFDMA upstream being analyzed are being sent in an OFDM downstream, the timestamp reported is the extended timestamp, while in a case with OFDMA upstream channels but no OFDM downstream channels, the reported timestamp is the DOCSIS 3.0 timestamp. For an active probe, the CMTS shall report the contents of the Probe Information Element (PIE) message describing that probe.

9.4.2 Upstream Triggered Spectrum Analysis

The upstream triggered spectrum analysis measurement provides a wideband spectrum analyzer function in the CMTS which can be triggered to examine desired upstream transmissions as well as underlying noise/interference during a quiet period.

The CMTS shall provide wideband upstream spectrum analysis capability covering the full upstream spectrum of the cable plant. The CMTS shall provide 100 kHz or better resolution (bin spacing) in the wideband upstream spectrum measurement. Free-running capture is done at a rate supported by the CMTS, and does not imply real-time operation.

The CMTS should provide the capability to average the FFT bin power of the spectrum over multiple captures.

The CMTS should provide a variable upstream spectrum analysis span.

The CMTS should be capable of providing the time-domain input samples as an alternative to the frequency-domain upstream spectrum results.

In pre-DOCSIS-3.1 mode, the CMTS shall provide the ability to trigger the spectrum sample capture and perform spectrum analysis using the following modes:

- Free running
- Trigger on minislot count
- Trigger on SID (service identifier)
- Trigger during quiet period (idle SID)

In DOCSIS 3.1 mode, the CMTS shall provide the ability to trigger spectrum sample capture and perform spectrum analysis using the following modes:

- Free running
- A specified timestamp value

- A specified MAC address, triggering at the beginning of the first minislot granted to any SID corresponding to the specified MAC address
- The idle SID, triggering at the beginning of the first minislot granted to that SID
- A specified active or quiet probe symbol, triggering at the beginning of the probe symbol

9.4.3 Upstream Impulse Noise Statistics

Upstream Impulse noise statistics gather statistics of burst/impulse noise occurring in a selected narrow band as defined in [i.2]. A bandpass filter is positioned in an unoccupied upstream band. A threshold is set, energy exceeding the threshold triggers the measurement of an event, and energy falling below the threshold ends the event. The CMTS may allow the threshold to be set to zero, in which case the average power in the band will be measured. The measurement is time-stamped using the DOCSIS 3.0 field of the 64-bit extended timestamp (bits 9-40, where bit 0 is the LSB), which provides a resolution of 98 ns and a range of 7 minutes.

The CMTS shall provide the capability to capture the following statistics in a selected band up to 5,12 MHz wide:

- Timestamp of event
- Duration of event
- Average power of event

The CMTS shall provide a time history buffer of up to 1 024 events.

9.4.4 Upstream Equalizer Coefficients

This item provides access to CM upstream pre-equalizer coefficients, and CMTS upstream adaptive equalizer coefficients, which taken together describe the linear response of the upstream cable plant for a given CM. The [i.3] specification defines summary metrics to avoid having to send all equalizer coefficients on every query. During the ranging process, the CMTS computes adaptive equalizer coefficients based on upstream probes; these coefficients describe the residual channel remaining after any pre-equalization. The CMTS sends these equalizer coefficients to the CM as a set of Transmit Equalization Adjust coefficients as part of the ranging process.

The CM shall provide the capability to report its upstream pre-equalizer coefficients (full set or summary) upon request. The CM shall provide the capability to also report the most recent set of Transmit Equalization Adjust coefficients which were applied to produce the reported set of upstream pre-equalizer coefficients. The CM shall report a condition in which it modified or did not apply the Transmit Equalization Adjust coefficients sent to it by the CMTS.

The CMTS shall provide a capability for reporting its upstream adaptive equalizer coefficients associated with probes from a CM upon request.

9.4.5 Upstream FEC Statistics

Upstream FEC statistics provide for monitoring upstream link quality via FEC and related statistics. Statistics are taken on codeword error events. The measurement is time-stamped using bits 21-52 of the extended timestamp, which provides a 32-bit timestamp value with resolution of 0,4 ms and range of 20 days. An LDPC codeword that fails post-decoding syndrome check will be labelled "unreliable", but the data portion of the codeword may not contain bit errors; hence the "unreliable codeword" count will tend to be pessimistic. All codewords, whether full-length or shortened, are included in the measurements. The codeword (or frame) and error counters are long (e.g. 64-bit) integers, so that overflow is not an issue.

The CMTS shall be capable of providing the following FEC statistics for any specified single upstream user:

- Pre-FEC Error-Free Codewords: Number of codewords that passed pre-decoding syndrome check.
- Unreliable Codewords: Number of codewords that failed post-decoding syndrome check.
- Corrected Codewords: Number of codewords that failed pre-decoding syndrome check, but passed post-decoding syndrome check.
- MAC CRC failures: Number of frames that failed MAC CRC check.

- Total number of FEC codewords.
- Total number of MAC frames.
- Start and stop time of analysis period, or time that snapshot of counters was taken.
- SID corresponding to upstream user being measured.

The CMTS shall be capable of providing the following FEC summaries over a rolling 10 minute period for any single upstream user:

- Total number of seconds.
- Number of errored seconds (seconds during which at least one unreliable codeword occurred).
- Count of codeword errors (unreliable codewords) in each 1-second interval (600 values over 10 minutes).
- Start and stop time of summary period.

9.4.6 Upstream Histogram

The purpose of the upstream histogram is to provide a measurement of nonlinear effects in the channel such as amplifier compression and laser clipping. For example, laser clipping causes one tail of the histogram to be truncated and replaced with a spike.

The CMTS shall be capable of capturing the histogram of time domain samples at the wideband front end of the receiver (full upstream band). When the CMTS creates an upstream histogram, the CMTS shall create it such that it is a two-sided histogram; that is, it encompasses values from far-negative to far-positive values of the samples. When a CMTS creates an upstream histogram, the CMTS shall create it such that is has either 256 equally spaced bins with even symmetry about the origin, or 255 equally spaced bins with odd symmetry about the origin. These bins typically correspond to the 8 MSBs of the wideband analog-to-digital converter (ADC).

The histogram dwell count, a 32-bit unsigned integer, is the number of samples observed while counting hits for a given bin, and may have the same value for all bins. The histogram hit count, a 32-bit unsigned integer, is the number of samples falling in a given bin. The CMTS shall be capable of reporting the dwell count per bin and the hit count per bin. When enabled, the CMTS shall compute a histogram with a dwell of at least 10 million samples 30 seconds or less. With this many samples, the histogram can reliably measure a probability density per bin as low as 10^{-6} with at least 10 hits in each bin. The CMTS shall continue accumulating histogram samples until it is restarted, disabled, approaches its 32-bit overflow value, or times out. The CMTS shall report the start and end time of the histogram measurement using bits 21 - 52 of the extended timestamp, which provides a 32-bit timestamp value with resolution of 0,4 ms and range of 20 days.

9.4.7 Upstream Channel Power

The purpose of the upstream channel power metric is to provide an estimate of the total received power in a specified OFDMA channel at the F connector input of the CMTS line card, or other agreed measurement point, for a given user. The measurement is based on upstream probes, which are typically the same probes used for pre-equalization adjustment. While digital power measurements are inherently accurate, the measurement referred to the analog input depends on available calibration accuracy.

The CMTS shall provide an estimate of total received power in a specified OFDMA channel at a reference input point, for a single specified upstream user. The CMTS shall provide configurable averaging over a range at least including 1 to 32 probes.

9.4.8 Upstream Receive Modulation Error Ratio (RxMER) Per Subcarrier

This item provides measurements of the upstream receive modulation error ratio (RxMER) for each subcarrier. The CMTS measures the RxMER using an upstream probe, which is not subject to symbol errors as data subcarriers would be. The probes used for RxMER measurement are typically distinct from the probes used for pre-equalization adjustment. For the purposes of this measurement, RxMER is defined as the ratio of the average power of the ideal QAM constellation to the average error-vector power. The error vector is the difference between the equalized received probe value and the known correct probe value. If some subcarriers (such as exclusion bands) cannot be measured by the CMTS, the CMTS indicates that condition in the measurement data for those subcarriers.

The CMTS shall be capable of providing measurements of RxMER for all active subcarriers for any single specified user in a specified OFDMA upstream channel, using probe symbols.

Annex A (normative): QAM Constellation Mappings

A.0 QAM Constellation Mappings Overview

The CMTS shall use the QAM constellation mappings given in this clause for all downstream transmissions. Downstream transmissions do not contain BPSK, 8-QAM and 32-QAM constellations.

The CM shall use the QAM constellation mappings given in this clause for all upstream transmissions. Upstream transmissions do not contain 8192-QAM and 16384-QAM constellations.

Sample code showing the bit to QAM constellation mapping is provided in [i.6].

A.1 QAM Constellations

Figures A.1 through A.7 below show the constellation mapping of an m-tuple onto a (Real, Imaginary) point in the complex plane. The horizontal axis is the real axis and the vertical axis is the imaginary axis.

The m-tuple is represented by:

$$\{y_0\ y_1\ \dots\ y_{m-1}\}$$

Mapping of the FEC encoded bit streams to the m-tuples is described in the sections detailing downstream and upstream transmissions.

Each m-tuple is represented as a hexadecimal number in all the constellation diagrams given below.

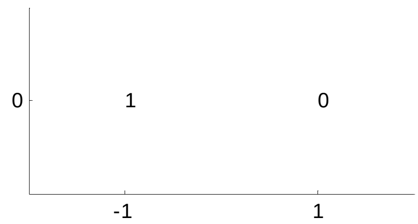

Figure A.1: BPSK Constellation Mapping of $\{y_0\}$

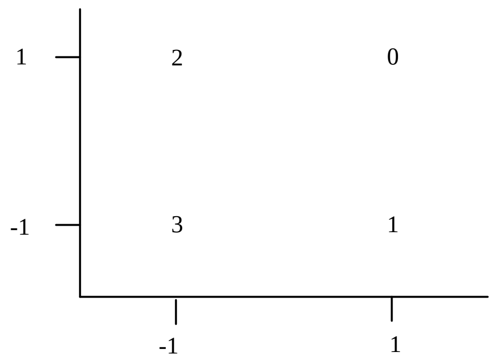

Figure A.2: QPSK Constellation Mapping of $\{y_0 y_1\}$

Figure A.3: 8-QAM Constellation Mapping of $\{y_0\ y_1\ y_2\}$

Figure A.4: 16-QAM Constellation Mapping of $\{y_0\ y_1\ y_2\ y_3\}$

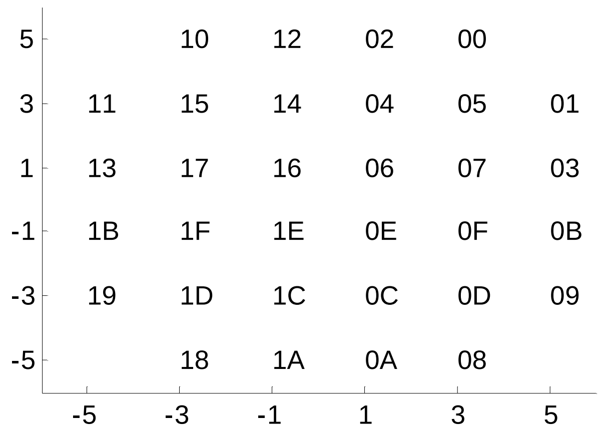

Figure A.5: 32-QAM Constellation Mapping of $\{y_0 y_1 y_2 y_3 y_4\}$

	-7	-5	-3	-1	1	3	5	7
7	20	22	2A	28	08	0A	02	00
5	21	23	2B	29	09	0B	03	01
3	25	27	2F	2D	0D	0F	07	05
1	24	26	2E	2C	0C	0E	06	04
-1	34	36	3E	3C	1C	1E	16	14
-3	35	37	3F	3D	1D	1F	17	15
-5	31	33	3B	39	19	1B	13	11
-7	30	32	3A	38	18	1A	12	10

Figure A.6: 64-QAM Constellation Mapping of $\{y_0 y_1 y_2 y_3 y_4 y_5\}$

	-11	-9	-7	-5	-3	-1	1	3	5	7	9	11
11			42	43	4B	4A	0A	0B	03	02		
9			40	41	49	48	08	09	01	00		
7	45	44	54	55	51	50	10	11	15	14	04	05
5	47	46	56	57	53	52	12	13	17	16	06	07
3	4F	4E	5E	5F	5B	5A	1A	1B	1F	1E	0E	0F
1	4D	4C	5C	5D	59	58	18	19	1D	1C	0C	0D
-1	6D	6C	7C	7D	79	78	38	39	3D	3C	2C	2D
-3	6F	6E	7E	7F	7B	7A	3A	3B	3F	3E	2E	2F
-5	67	66	76	77	73	72	32	33	37	36	26	27
-7	65	64	74	75	71	70	30	31	35	34	24	25
-9			60	61	69	68	28	29	21	20		
-11			62	63	6B	6A	2A	2B	23	22		

Figure A.7: 128-QAM Constellation Mapping of $\{y_0 y_1 y_2 y_3 y_4 y_5 y_6\}$

In order to reduce the size of the diagrams, only the first quadrant is shown for the larger constellations, namely, 256-QAM, 512-QAM, 1024-QAM, 2048-QAM, 4096-QAM, 8192-QAM and 16384-QAM. The mapping of the two bits $\{y_0\ y_1\}$ of $\{y_0\ y_1\ \dots\ y_{m-1}\}$ is the same for all these QAM constellations, as illustrated in figure A.8.

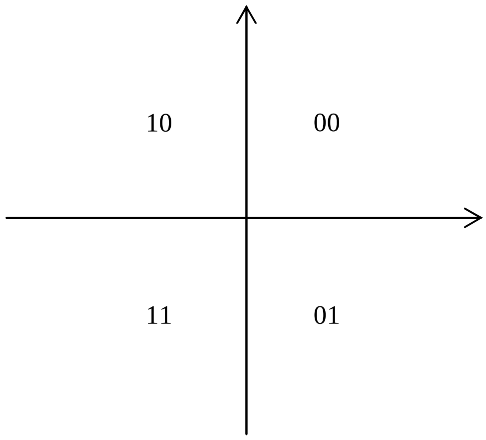

Figure A.8: Mapping of Bits $\{y_0\ y_1\}$ of $\{y_0\ y_1, \dots, y_{m-1}\}$ for Constellations with only One Quadrant Defined

The mapping of the bits of $\{y_2\ y_3, \dots, y_{m-1}\}$ to the first quadrant of the constellation is given in Figures A.10 through A.16 below for 256-QAM, 512-QAM, 1024-QAM, 2048-QAM, 4096-QAM, 8192-QAM and 16384-QAM. The mappings for the other three quadrants are obtained by mirroring the first quadrant about the horizontal and vertical axis as illustrated in Figure A.9 below. This figure shows only a (3x3) grid of points in each quadrant for illustration of the above mentioned reflective property. However, this reflective mapping is applicable to any number of points in each quadrant. Quadrant 1 is reflected about the vertical axis to get quadrant 2. Quadrant 1 is reflected about the horizontal axis to get quadrant 4. Quadrant 2 is reflected about the horizontal axis to get quadrant 3.

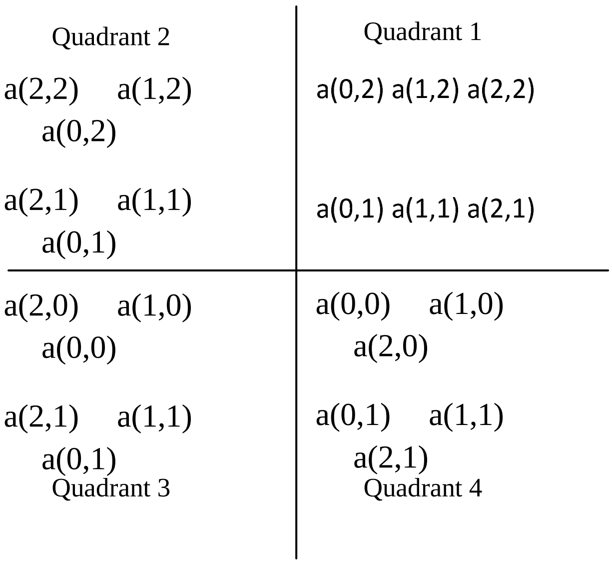

Figure A.9: Reflective Mapping of Bits $\{y_2\ y_3, \dots, y_{m-1}\}$ for All Constellations (Except BPSK)

15	20	22	2A	28	08	0A	02	00
13	21	23	2B	29	09	0B	03	01
11	25	27	2F	2D	0D	0F	07	05
9	24	26	2E	2C	0C	0E	06	04
7	34	36	3E	3C	1C	1E	16	14
5	35	37	3F	3D	1D	1F	17	15
3	31	33	3B	39	19	1B	13	11
1	30	32	3A	38	18	1A	12	10
	1	3	5	7	9	11	13	15

Figure A.10: 256-QAM Constellation Mapping of $\{y_2 y_3 y_4 y_5 y_6 y_7\}$ on to Quadrant 1

23	028	029	02D	02C	00C	00D	009	008				
21	02A	02B	02F	02E	00E	00F	00B	00A				
19	022	023	027	026	006	007	003	002				
17	020	021	025	024	004	005	001	000				
15	040	041	045	044	054	055	051	050	010	011	015	014
13	042	043	047	046	056	057	053	052	012	013	017	016
11	04A	04B	04F	04E	05E	05F	05B	05A	01A	01B	01F	01E
9	048	049	04D	04C	05C	05D	059	058	018	019	01D	01C
7	068	069	06D	06C	07C	07D	079	078	038	039	03D	03C
5	06A	06B	06F	06E	07E	07F	07B	07A	03A	03B	03F	03E
3	062	063	067	066	076	077	073	072	032	033	037	036
1	060	061	065	064	074	075	071	070	030	031	035	034
	1	3	5	7	9	11	13	15	17	19	21	23

Figure A.11: 512-QAM Constellation Mapping of $\{y_2 y_3 y_4 y_5 y_6 y_7 y_8\}$ on to Quadrant 1

	1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31
31	080	082	08A	088	0A8	0AA	0A2	0A0	020	022	02A	028	008	00A	002	000
29	081	083	08B	089	0A9	0AB	0A3	0A1	021	023	02B	029	009	00B	003	001
27	085	087	08F	08D	0AD	0AF	0A7	0A5	025	027	02F	02D	00D	00F	007	005
25	084	086	08E	08C	0AC	0AE	0A6	0A4	024	026	02E	02C	00C	00E	006	004
23	094	096	09E	09C	0BC	0BE	0B6	0B4	034	036	03E	03C	01C	01E	016	014
21	095	097	09F	09D	0BD	0BF	0B7	0B5	035	037	03F	03D	01D	01F	017	015
19	091	093	09B	099	0B9	0BB	0B3	0B1	031	033	03B	039	019	01B	013	011
17	090	092	09A	098	0B8	0BA	0B2	0B0	030	032	03A	038	018	01A	012	010
15	0D0	0D2	0DA	0D8	0F8	0FA	0F2	0F0	070	072	07A	078	058	05A	052	050
13	0D1	0D3	0DB	0D9	0F9	0FB	0F3	0F1	071	073	07B	079	059	05B	053	051
11	0D5	0D7	0DF	0DD	0FD	0FF	0F7	0F5	075	077	07F	07D	05D	05F	057	055
9	0D4	0D6	0DE	0DC	0FC	0FE	0F6	0F4	074	076	07E	07C	05C	05E	056	054
7	0C4	0C6	0CE	0CC	0EC	0EE	0E6	0E4	064	066	06E	06C	04C	04E	046	044
5	0C5	0C7	0CF	0CD	0ED	0EF	0E7	0E5	065	067	06F	06D	04D	04F	047	045
3	0C1	0C3	0CB	0C9	0E9	0EB	0E3	0E1	061	063	06B	069	049	04B	043	041
1	0C0	0C2	0CA	0C8	0E8	0EA	0E2	0E0	060	062	06A	068	048	04A	042	040

Figure A.12: 1024-QAM Constellation Mapping of $\{y_2 y_3 y_4 y_5 y_6 y_7 y_8 y_9\}$ on to Quadrant 1

	1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35	37	39	41	43	45	47
47	0A0	0A1	0A5	0A4	0B4	0B5	0B1	0B0	030	031	035	034	024	025	021	020								
45	0A2	0A3	0A7	0A6	0B6	0B7	0B3	0B2	032	033	037	036	026	027	023	022								
43	0AA	0AB	0AF	0AE	0BE	0BF	0BB	0BA	03A	03B	03F	03E	02E	02F	02B	02A								
41	0A8	0A9	0AD	0AC	0BC	0BD	0B9	0B8	038	039	03D	03C	02C	02D	029	028								
39	088	089	08D	08C	09C	09D	099	098	018	019	01D	01C	00C	00D	009	008								
37	08A	08B	08F	08E	09E	09F	09B	09A	01A	01B	01F	01E	00E	00F	00B	00A								
35	082	083	087	086	096	097	093	092	012	013	017	016	006	007	003	002								
33	080	081	085	084	094	095	091	090	010	011	015	014	004	005	001	000								
31	100	101	105	104	114	115	111	110	150	151	155	154	144	145	141	140	040	041	045	044	054	055	051	050
29	102	103	107	106	116	117	113	112	152	153	157	156	146	147	143	142	042	043	047	046	056	057	053	052
27	10A	10B	10F	10E	11E	11F	11B	11A	15A	15B	15F	15E	14E	14F	14B	14A	04A	04B	04F	04E	05E	05F	05B	05A
25	108	109	10D	10C	11C	11D	119	118	158	159	15D	15C	14C	14D	149	148	048	049	04D	04C	05C	05D	059	058
23	128	129	12D	12C	13C	13D	139	138	178	179	17D	17C	16C	16D	169	168	068	069	06D	06C	07C	07D	079	078
21	12A	12B	12F	12E	13E	13F	13B	13A	17A	17B	17F	17E	16E	16F	16B	16A	06A	06B	06F	06E	07E	07F	07B	07A
19	122	123	127	126	136	137	133	132	172	173	177	176	166	167	163	162	062	063	067	066	076	077	073	072
17	120	121	125	124	134	135	131	130	170	171	175	174	164	165	161	160	060	061	065	064	074	075	071	070
15	1A0	1A1	1A5	1A4	1B4	1B5	1B1	1B0	1F0	1F1	1F5	1F4	1E4	1E5	1E1	1E0	0E0	0E1	0E5	0E4	0F4	0F5	0F1	0F0
13	1A2	1A3	1A7	1A6	1B6	1B7	1B3	1B2	1F2	1F3	1F7	1F6	1E6	1E7	1E3	1E2	0E2	0E3	0E7	0E6	0F6	0F7	0F3	0F2
11	1AA	1AB	1AF	1AE	1BE	1BF	1BB	1BA	1FA	1FB	1FF	1FE	1EE	1EF	1EB	1EA	0EA	0EB	0EF	0EE	0FE	0FF	0FB	0FA
9	1A8	1A9	1AD	1AC	1BC	1BD	1B9	1B8	1F8	1F9	1FD	1FC	1EC	1ED	1E9	1E8	0E8	0E9	0ED	0EC	0FC	0FD	0F9	0F8
7	188	189	18D	18C	19C	19D	199	198	1D8	1D9	1DD	1DC	1CC	1CD	1C9	1C8	0C8	0C9	0CD	0CC	0DC	0DD	0D9	0D8
5	18A	18B	18F	18E	19E	19F	19B	19A	1DA	1DB	1DF	1DE	1CE	1CF	1CB	1CA	0CA	0CB	0CF	0CE	0DE	0DF	0DB	0DA
3	182	183	187	186	196	197	193	192	1D2	1D3	1D7	1D6	1C6	1C7	1C3	1C2	0C2	0C3	0C7	0C6	0D6	0D7	0D3	0D2
1	180	181	185	184	194	195	191	190	1D0	1D1	1D5	1D4	1C4	1C5	1C1	1C0	0C0	0C1	0C5	0C4	0D4	0D5	0D1	0D0

Figure A.13: 2048-QAM Constellation Mapping of $\{y_2 y_3 y_4 y_5 y_6 y_7 y_8 y_9 y_{10}\}$ on to Quadrant 1

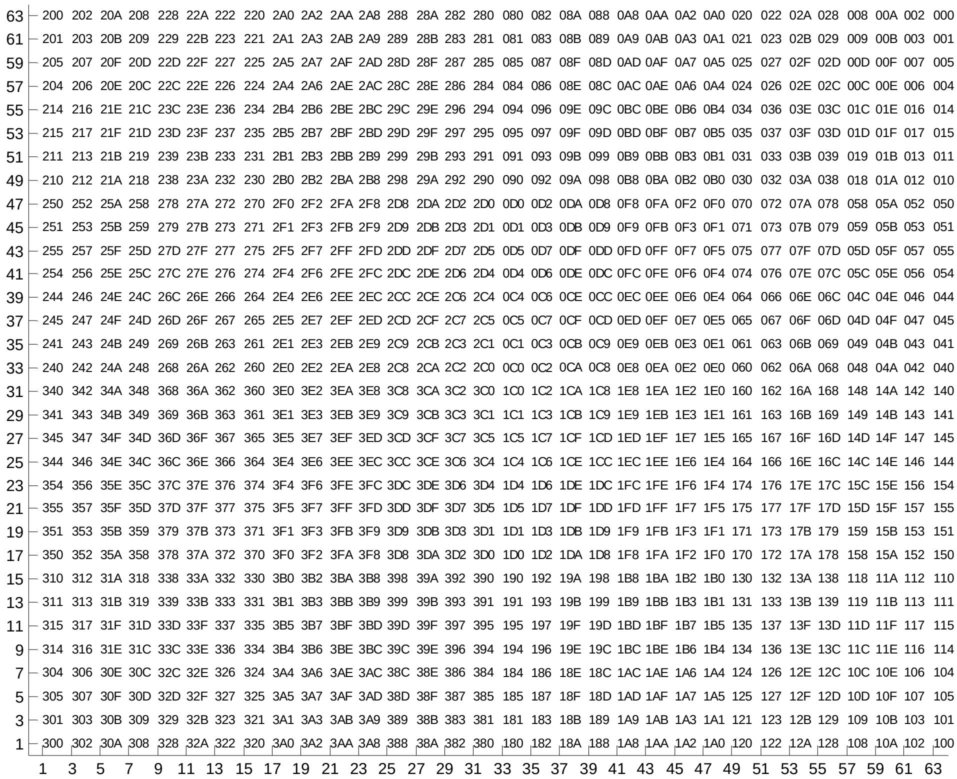

Figure A.14: 4096-QAM Constellation Mapping of $\{y_2 y_3 y_4 y_5 y_6 y_7 y_8 y_9 y_{10} y_{11}\}$ on to Quadrant 1

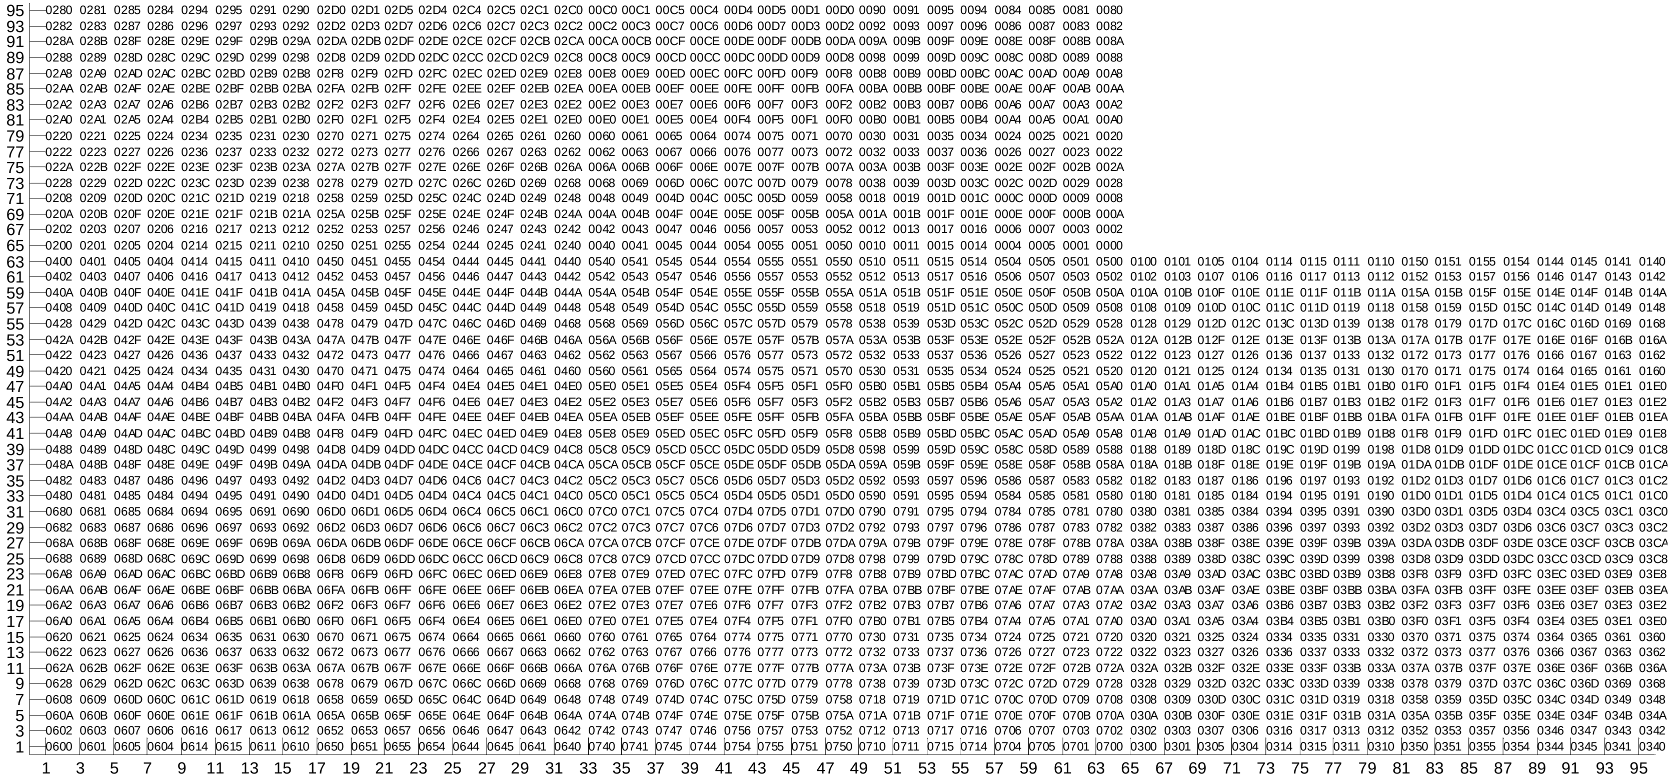

Figure A.15: 8192-QAM Constellation Mapping of $\{y_2 y_3 y_4 y_5 y_6 y_7 y_8 y_9 y_{10} y_{11} y_{12}\}$ on to Quadrant 1

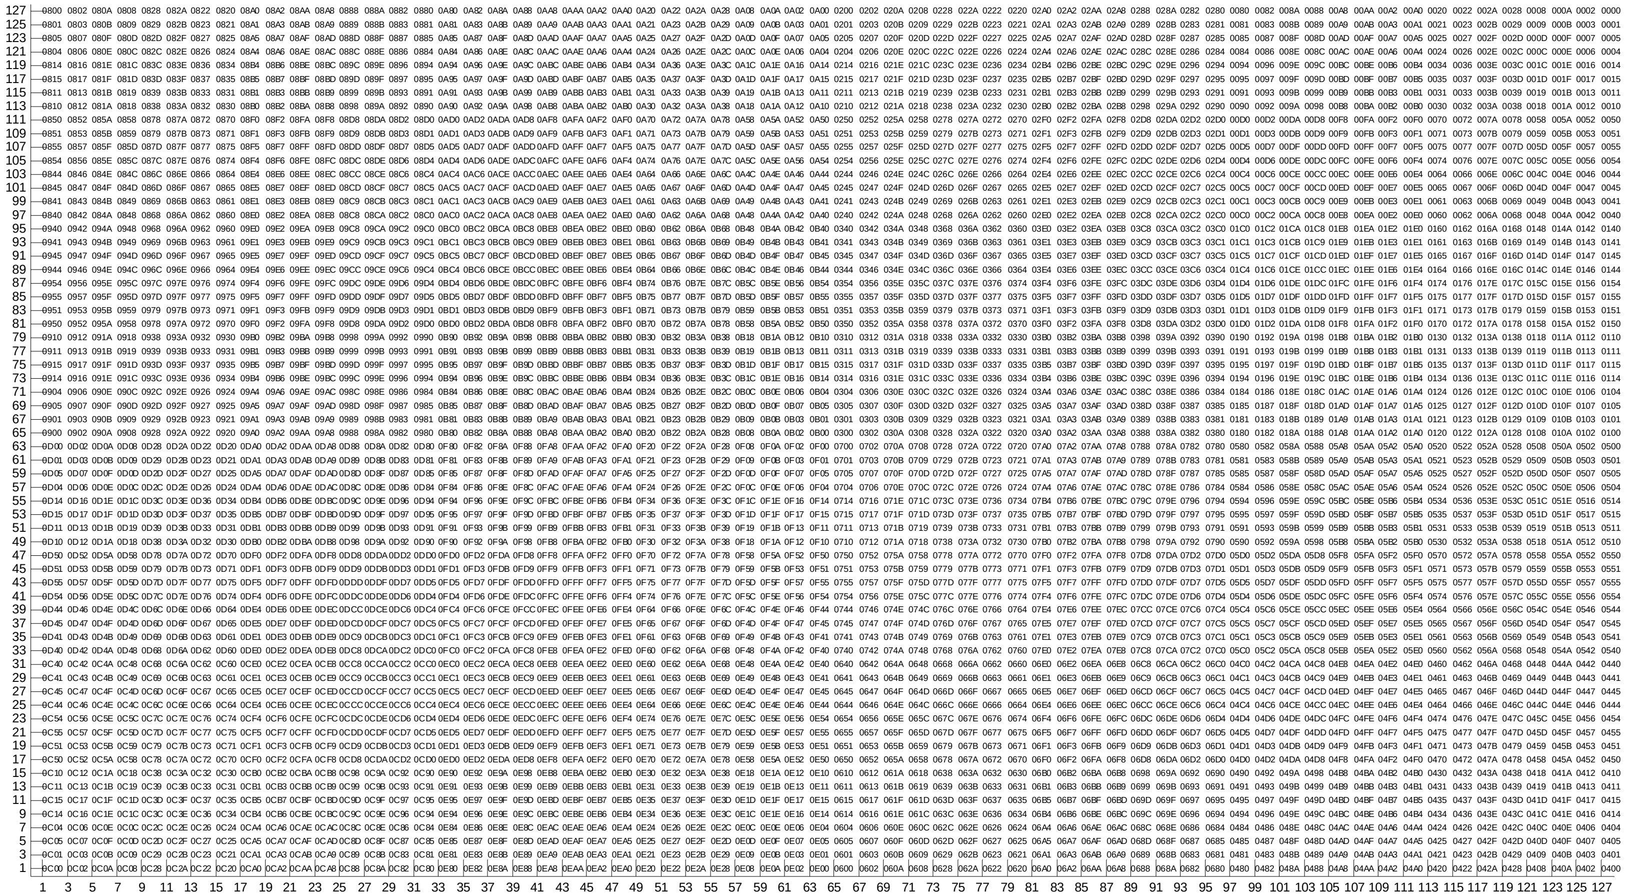

Figure A.16: 16384-QAM Constellation Mapping of $\{y_2y_3y_4y_5y_6y_7y_8y_9y_{10}y_{11}y_{12}y_{13}\}$ on to Quadrant 1

A.2 QAM Constellation Scaling

The CM shall scale real and imaginary axes of the constellations by the scaling factors given in column 3 of the table A.2, to ensure that the mean square value of all QAM constellations are equal to 1.0.

The CMTS shall scale real and imaginary axes of the constellations by the scaling factors given in column 3 of the table A.2, to ensure that the mean square value of all QAM constellations are equal to 1.0.

Table A.1: QAM Constellation Scaling Factors

QAM Constellation	m Number of bits	Scaling Factor
BPSK	1	1
QPSK	2	$1/\sqrt{2}$
8-QAM	3	$1/\sqrt{10}$
16-QAM	4	$1/\sqrt{10}$
32-QAM	5	$1/\sqrt{20}$
64-QAM	6	$1/\sqrt{42}$
128-QAM	7	$1/\sqrt{82}$
256-QAM	8	$1/\sqrt{170}$
512-QAM	9	$1/\sqrt{330}$
1024-QAM	10	$1/\sqrt{682}$
2048-QAM	11	$1/\sqrt{1322}$
4096-QAM	12	$1/\sqrt{2730}$
8192-QAM	13	$1/\sqrt{5290}$
16384-QAM	14	$1/\sqrt{10922}$

Annex B (normative): RFoG Operating Mode

The CMTS shall support the ability to limit the number of simultaneous US transmitters to a single transmitter at a time.

Annex C (normative): Additions and Modifications for European Specification with SC-QAM Operation

This clause applies to cases where a DOCSIS 3.1 CM or CMTS is operating with Single Carrier QAM (SC-QAM) operation only, with no OFDM operation. As such, it represents backward compatibility requirements when operating with DOCSIS 3.0 systems or with the DOCSIS 3.1 PHY disabled. It also applies only to the second technology option referred to in clause 1.1; for the first option refer to clause 6, and for the third option refer to Annex D.

As the requirements for a DOCSIS 3.1 CM and CMTS are largely unchanged relative to DOCSIS 3.0 devices for SC-QAM operation, the requirements for operating with this technology option and in this mode are addressed via reference to the PHYv3.0 and DRFI specifications, with the exception that the minimum requirement for upstream and downstream channels has been changed for DOCSIS 3.1 devices.

A DOCSIS 3.1 CM shall support the CM requirements in Annex B of [5], with the exception that the minimum requirement for upstream channels is 8, and the minimum requirement for downstream channels is 24.

A DOCSIS 3.1 CMTS shall support the CMTS requirements in Annex B of [5] with the exception that the minimum requirement for upstream channels is 8. A DOCSIS 3.1 CMTS shall support the CMTS requirements in Annex B of [3], with the addition that the minimum requirement for downstream channels is 24.

Annex D: Void

This annex may be added in a subsequent revision of the present document. It is left in the present document to reserve the sequence in annexes.

Annex E (normative): 24-bit Cyclic Redundancy Check (CRC) Code

This clause contains a 24-bits CRC code encoding, which is used for NCPs as specified in clause 7.5.14 and Initial Ranging as specified in clause 7.4.16.1.

The CRC encoder generates the 24 bits parity bits denoted by $p_0, p_1, p_2, p_3, ..., p_{23}$ for the input bit stream $b_0, b_1, ..., b_{k-1}$ .using the following generator polynomial:

$$g_{CRC24}(x) = x^{24} + x^{22} + x^{20} + x^{19} + x^{18} + x^{16} + x^{14} + x^{13} + x^{11} + x^{10} + x^8 + x^7 + x^6 + x^3 + x + 1$$

(127266713 in octal representation) [1], which means in GF(2) the following equation holds

$$b_0 x^{k+23} + b_1 x^{k+22} + ... + b_{k-1} x^{24} + p_0 x^{23} + p_1 x^{22} + ... + p_{22} x^1 + p_{23} = 0 \bmod g_{CRC24}(x).$$

This 24-bit CRC polynomial is optimized by G. Castangnoli, S. Bräuer and M. Hermann in [i.1].

CRC-24-D is displayed most significant byte first, most significant bit first. CRC is shown in brackets.

Example for 7 byte frame such as O-INIT-RNG-REQ:

7 byte frame:

```
01 02 03 04 05 06 07 [ cd ef 27 ]
```

Example for 255 byte frame MSB first:

```
01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f 10
11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f 20
21 22 23 24 25 26 27 28 29 2a 2b 2c 2d 2e 2f 30
31 32 33 34 35 36 37 38 39 3a 3b 3c 3d 3e 3f 40
41 42 43 44 45 46 47 48 49 4a 4b 4c 4d 4e 4f 50
51 52 53 54 55 56 57 58 59 5a 5b 5c 5d 5e 5f 60
61 62 63 64 65 66 67 68 69 6a 6b 6c 6d 6e 6f 70
71 72 73 74 75 76 77 78 79 7a 7b 7c 7d 7e 7f 80
81 82 83 84 85 86 87 88 89 8a 8b 8c 8d 8e 8f 90
91 92 93 94 95 96 97 98 99 9a 9b 9c 9d 9e 9f a0
a1 a2 a3 a4 a5 a6 a7 a8 a9 aa ab ac ad ae af b0
b1 b2 b3 b4 b5 b6 b7 b8 b9 ba bb bc bd be bf c0
c1 c2 c3 c4 c5 c6 c7 c8 c9 ca cb cc cd ce cf d0
d1 d2 d3 d4 d5 d6 d7 d8 d9 da db dc dd de df e0
e1 e2 e3 e4 e5 e6 e7 e8 e9 ea eb ec ed ee ef f0
f1 f2 f3 f4 f5 f6 f7 f8 f9 fa fb fc fd fe ff [ 2c a8 8b ]
```

Annex F (informative): Downstream Frequency Interleaver Sample C Code

In the downstream frequency interleaver C code given below, X is the input data array, Y is the output data array, and N is the size of each array. It has been written to illustrate each operation clearly, and as such it may not necessarily be the most efficient implementation.

```
void Docsis_3_1_Freqeuncy_Interleaver ( int *X, int *Y, int N )
{
    int store[128][64];
    int i, j, K, k1, base;
    int RowCount, RowCount_BR;
    int ColumnCount, ColumnCount_BR;
    int Last_Column_Size;
    int column_rotate[64], rotated_column[128];

    // Number of rows is 128
    // Number of columns is K
    K = (int) (ceil (((double) N)/128.0));
    Last_Column_Size = N - ((K-1)*128);

    // Generate the column rotation array using the 6-bit LFSR
    column_rotate[0]=17;
    for (i=0; i<63; i++)
    {
             int lsb;
             lsb = column_rotate[i]&1;
            column_rotate[i+1] = (column_rotate[i] >> 1) ^ (lsb << 5) ^ (lsb << 4);
    }

    base = 0;
     // Save data in 2-D store with rows addressed in bit-reversed order
    for (RowCount=0; RowCount< 128; RowCount++)
    {
            RowCount_BR = 0;
            for (j=0; j<7; j++)
            {
                    RowCount_BR = RowCount_BR << 1;;
                    if (((RowCount >> j) & 1) == 1)
                    RowCount_BR++;
             }
            if (RowCount_BR < Last_Column_Size)
            {
                    for (j=0; j<K; j++)
                        store[RowCount_BR][j] = X[base+j];
                    base += K;
            }
            else
            {
                    for (j=0; j<K-1; j++)
                        store[RowCount_BR][j] = X[base+j];
                    base += (K-1);
             }
      }

    // Rotate columns 0 to K-2
    // Last column that could be partially filled is not rotated
    for (j=0; j<(K-1); j++)
    {
            for (i=0; i<128; i++)
                    rotated_column[(i + column_rotate[j]) & 0x7F] = store[i][j];
            for (i=0; i<128; i++)
                    store[i][j] = rotated_column[i];
    }

    // Determine the number of bits k1 in ColumnCount
    k1 = 0;
    while ((1 << k1) < K)
            k1++;

    // Address columns in bit-reversed order, but only if bit-reversed number is within range
    // Otherwise use non-bit-reversed version
    base = 0;
```

```
    for (ColumnCount=0; ColumnCount<K; ColumnCount++)
    {
            int ReadColumn;
            // Bit Reverse ColumnCount
            ColumnCount_BR = 0;
            for (i=0; i<k1; i++)
            {
                    ColumnCount_BR = ColumnCount_BR << 1;
                    if (((ColumnCount >> i) & 1) == 1)
                        ColumnCount_BR++;
            }
            if (ColumnCount_BR > (K-1)) // Read from ColumnCount
                    ReadColumn = ColumnCount;
            else
                    ReadColumn = ColumnCount_BR;
            if (ReadColumn == (K-1)) // Last column could be a partially-filled column
            {
                    for (i=0; i<Last_Column_Size; i++)
                        Y[base+i] = store[i][ReadColumn];
                    base = base+Last_Column_Size;
            }
            else
            {
                    for (i=0; i<128; i++)
                        Y[base+i] = store[i][ReadColumn];
                    base = base+128;
            }
    }
    return;
}
```

Annex G (informative): Use Cases: Maximum Number of Simultaneous Transmitters

This annex will be added in a subsequent revision to the present document.

Annex H (informative): Upstream Time and Frequency Interleaver Sample C Code

The algorithm for generating the sequence of addresses *(t, f)* is described using the following C code segment:

```
Count_t = 0;
Count_f = 0;
Count_diagonal = 0;
for (idx_t=0;   idx_t<K;   idx_t++)
{
    for (idx_f=0;   idx_f<L;   idx_f++)
    {
        Address.t  = Bit_Reverse_Count(&Count_t, K);
        Address.f  = Bit_Reverse_Count(&Count_f, L);
        Count_t = (Count_t + 1) % K1;
        Count_f = (Count_f + 1) % L1;
    }
    Count_diagonal = (Count_diagonal + 1) % K1;
    d = Bit_Reverse_Count(&Count_diagonal, K);
    Count_t = Count_diagonal;
    Count_f = 0;
}
```

Annex I (informative): FEC Codeword Selection Algorithm Upstream Time and Frequency Interleaver Sample C Code

If the CMTS scheduler wishes to grant a certain number of information bits, it needs to perform the opposite (forward) calculation in order to determine how many codewords of what sizes are necessary to hold the desired number of information bits. This is part of the process of determining the grant size. For informative purposes, the script forward_calc.m is provided to show how the CMTS could perform this calculation. The variable finfo_size is the input to this script.

```
% forward_calc script

% set values for codeword sizes
% total bits = size including parity
% info bits = information bits only
% thresholds - if more bits than threshold, shorten this cw instead of
% using a smaller one

% short codeword
SHORT_TOTAL_BITS = 1120;
SHORT_INFO_BITS = 840;
SHORT_PARITY_BITS = SHORT_TOTAL_BITS - SHORT_INFO_BITS;

SHORT_TOTAL_THRESH_BITS = SHORT_PARITY_BITS + 1;
SHORT_MIN_INFO_BITS = SHORT_INFO_BITS / 2;

% medium codeword
MED_TOTAL_BITS = 5940;
MED_INFO_BITS = 5040;
MED_PARITY_BITS = MED_TOTAL_BITS - MED_INFO_BITS;

MED_INFO_THRESH_BITS = 2521;

% long codeword
LONG_TOTAL_BITS = 16200;
LONG_INFO_BITS = 14400;
LONG_PARITY_BITS = LONG_TOTAL_BITS - LONG_INFO_BITS;

LONG_TOTAL_THRESH_BITS = 11881;
LONG_INFO_THRESH_BITS = 10081;


% variable finfo_size is input
% set finfo_size to desired input value in workspace

% initialize output variables
flong_cws = 0;
fshortened_long_cws = 0;
fmed_cws = 0;
fshortened_med_cws = 0;
fshort_cws = 0;
fshortened_short_cws = 0;
fother_shortened_cw_bits = 0;
fshortened_cw_bits = 0;
mac_padding = 0;

% intermediate variable to track type of last full codeword
flast_full_cw = '';

% now begin calculation


  bits_remaining = finfo_size;

  % if there are no bits at all, we don't want to give any grant - just let
  % everything fall through to zero codewords.

  % However, if we have a nonzero number of bits that is not enough to make
  % a min-size short codeword, we do want to give a grant but we are to
  % allow space for the CM to make a segment that will fill the min-size
  % short codeword.
```

```
if (bits_remaining > 0 && bits_remaining < SHORT_MIN_INFO_BITS)
    mac_padding = SHORT_MIN_INFO_BITS - bits_remaining;
    bits_remaining = SHORT_MIN_INFO_BITS;
end

  % make as many long cws as possible
  while bits_remaining >= LONG_INFO_BITS
      flong_cws = flong_cws + 1;
      bits_remaining = bits_remaining - LONG_INFO_BITS;
      flast_full_cw = 'Long';
  end

  % if remaining bits can make a shortened long cw, do so
  if bits_remaining >= LONG_INFO_THRESH_BITS
      fshortened_long_cws = 1;
      fshortened_cw_bits = bits_remaining;
      bits_remaining = 0;
  end

  % now make as many medium cws as possible
  while bits_remaining >= MED_INFO_BITS
      fmed_cws = fmed_cws + 1;
      bits_remaining = bits_remaining - MED_INFO_BITS;
      flast_full_cw = 'Medium';
  end

  % if remaining bits can make a shortened med cw, do so
  if bits_remaining >= MED_INFO_THRESH_BITS
      fshortened_med_cws = 1;
      fshortened_cw_bits = bits_remaining;
      bits_remaining = 0;
  end

  % now make as many short cws as possible
  while bits_remaining >= SHORT_INFO_BITS
      fshort_cws = fshort_cws + 1;
      bits_remaining = bits_remaining - SHORT_INFO_BITS;
      flast_full_cw = 'Short';
  end

  % if there are any bits left, finish with a shortened short cw
  if bits_remaining >= 1
      fshortened_short_cws = 1;

      % we need at least SHORT_MIN_INFO_BITS in a shortened short
      % codeword; if we don't have enough, we will borrow from the
      % immediately preceding full codeword, which will become another
      % shortened codeword.
      % Note that we will always borrow SHORT_MIN_INFO_BITS.

      if bits_remaining >= SHORT_MIN_INFO_BITS
          % no need to borrow bits
          fshortened_cw_bits = bits_remaining;
          bits_remaining = 0;
      else
          % identify type/size of last full codeword
          switch flast_full_cw
              case 'Long'
                  % change last full cw to a shortened cw
                  flong_cws = flong_cws - 1;
                  fshortened_long_cws = fshortened_long_cws + 1;
                  % number of bits in that cw is reduced by
                  % SHORT_MIN_INFO_BITS
                  fother_shortened_cw_bits = LONG_INFO_BITS - ...
                      SHORT_MIN_INFO_BITS;
                  % put those bits plus bits_remaining into the last
                  % shortened cw
            fshortened_cw_bits = SHORT_MIN_INFO_BITS +bits_remaining;
                  bits_remaining = 0;
              case 'Medium'
                  % same steps as with long
                  fmed_cws = fmed_cws - 1;
                  fshortened_med_cws = fshortened_med_cws + 1;
                  fother_shortened_cw_bits = MED_INFO_BITS - ...
                      SHORT_MIN_INFO_BITS;
                  fshortened_cw_bits = SHORT_MIN_INFO_BITS +bits_remaining;
                  bits_remaining = 0;
```

```
            case 'Short'
                % also same as long
                fshort_cws = fshort_cws - 1;
                fshortened_short_cws = fshortened_short_cws + 1;
                fother_shortened_cw_bits = SHORT_INFO_BITS - ...
                    SHORT_MIN_INFO_BITS;
                fshortened_cw_bits = SHORT_MIN_INFO_BITS +bits_remaining;
                bits_remaining = 0;
        end
    end
end
```

Tables showing the number and size of codewords to be used for grant sizes from 1 bit up to two full long codewords are provided in [i.5].

Annex J (informative): CMTS Proposed Configuration Parameters

Table J.1: CMTS Proposed Configuration Parameters

FFT	Cyclic Prefix Samples (N_{cp})	Roll-Off Period Samples (N_{rp})	Bandedge Exclusion Sub-band (MHz)	Lower Edge Exclusion Sub-band (Subcarriers)	Upper Edge Exclusion Sub-band (Subcarriers)
4K	192	64	3,650	201	200
		128	2,000	168	167
	256	64	3,650	201	200
		128	2,000	168	167
		192	1,450	157	156
	512	64	3,650	201	200
		128	1,950	167	166
		192	1,400	156	155
		256	1,100	150	149
	768	64	3,600	200	199
		128	1,950	167	166
		192	1,400	156	155
		256	1,100	150	149
	1 024	64	3,600	200	199
		128	1,900	166	165
		192	1,350	155	154
		256	1,050	149	148
8K	192	64	3,400	392	391
		128	1,750	326	325
	256	64	3,400	392	391
		128	1,750	326	325
		192	1,225	305	304
	512	64	3,375	391	390
		128	1,750	326	325
		192	1,200	304	303
		256	0,925	293	292
	768	64	3,375	391	390
		128	1,750	326	325
		192	1,200	304	303
		256	0,925	293	292
	1 024	64	3,350	390	389
		128	1,725	325	324
		192	1,200	304	303
		256	0,925	293	292

Annex K (informative): Suggested Algorithm to Compute Signal-to-Noise Ratio (SNR) Margin for Candidate Profile

The CM measures the RxMER value for each data subcarrier as specified in clause 9.3.6. From these measurements it calculates the average RxMER over all data subcarriers, MER1.

The CM accepts as an input the required RxMER delivering a defined threshold of CER = 1e-5 under ideal AWGN conditions for each bit loading. The CM computes the difference of the measured RxMER values from the required RxMER values. The CM computes the required average RxMER, denoted MER2, over all data subcarriers for the candidate profile.

The averaging computations for MER1 and MER2 use values in the log (dB) domain. The SNR margin is defined as MER1 - MER2, where all quantities are in dB. As an example, if the CM measures MER1 = 33 dB, and the candidate profile requires MER2 = 30 dB, the CM reports an SNR margin of 3 dB. In addition, the CM reports the number of subcarriers whose RxMER is at least x dB below the defined threshold for the bit loading of the given subcarrier, where x is a configurable parameter with default value = 1.

Annex L (informative): Bibliography

- Digital Transmission Characterization of Cable Television Systems, Cable Television Laboratories, Inc., November 1994.

NOTE: See https://apps.cablelabs.com/specification/CL-TR-DTCCTS-V01-941101.

History

Document history		
V1.1.1	February 2017	Publication