

Superseded **by a later version of this document**

Data Over Cable Service Interface Specifications DOCSIS 3.0

Physical Layer Specification

CM-SP-PHYv3.0-I01-060804

ISSUED

Notice

This DOCSIS specification is a cooperative effort undertaken at the direction of Cable Television Laboratories, Inc. (CableLabs®) for the benefit of the cable industry. Neither CableLabs, nor any other entity participating in the creation of this document, is responsible for any liability of any nature whatsoever resulting from or arising out of use or reliance upon this document by any party. This document is furnished on an AS-IS basis and neither CableLabs, nor other participating entity, provides any representation or warranty, express or implied, regarding its accuracy, completeness, or fitness for a particular purpose.

© Copyright 2006 Cable Television Laboratories, Inc.

All rights reserved.

Document Status Sheet

Document Control Number:	CM-SP-PHYv3.0-I01-060804			
Document Title:	Physical Layer Specification			
Revision History:	I01 – Released 8/4/06			
Date:	August 4, 2006			
Status:	Work in Progress	Draft	Issued	Closed
Distribution Restrictions:	Author Only	CL/Member	CL/Member/Vendor	Public

Key to Document Status Codes:

- Work in Progress** An incomplete document, designed to guide discussion and generate feedback that may include several alternative requirements for consideration.

- Draft** A document in specification format considered largely complete, but lacking review by Members and vendors. Drafts are susceptible to substantial change during the review process.

- Issued** A stable document, which has undergone rigorous member and vendor review and is suitable for product design and development, cross-vendor interoperability, and for certification testing.

- Closed** A static document, reviewed, tested, validated, and closed to further engineering change requests to the specification through CableLabs.

Trademarks:

DOCSIS[®], eDOCSIS[™], PacketCable[™], CableHome[®], CableOffice[™], OpenCable[™], CableCARD[™], OCAP[™], M-CMTS[™], and CableLabs[®] are trademarks of Cable Television Laboratories, Inc.

Contents

1	SCOPE.....	1
1.1	INTRODUCTION AND PURPOSE.....	1
1.2	BACKGROUND.....	1
1.2.1	<i>Broadband Access Network.....</i>	<i>1</i>
1.2.2	<i>Network and System Architecture.....</i>	<i>2</i>
1.2.3	<i>Service Goals.....</i>	<i>3</i>
1.2.4	<i>Statement of Compatibility.....</i>	<i>3</i>
1.2.5	<i>Reference Architecture.....</i>	<i>4</i>
1.2.6	<i>DOCSIS 3.0 Documents.....</i>	<i>4</i>
1.3	REQUIREMENTS.....	5
1.4	CONVENTIONS.....	5
1.5	ORGANIZATION OF DOCUMENT.....	5
2	REFERENCES.....	7
2.1	NORMATIVE REFERENCES.....	7
2.2	INFORMATIVE REFERENCES.....	8
2.3	REFERENCE ACQUISITION.....	8
3	TERMS AND DEFINITIONS.....	9
4	ABBREVIATIONS AND ACRONYMS.....	16
5	FUNCTIONAL ASSUMPTIONS.....	18
5.1	EQUIPMENT ASSUMPTIONS.....	18
5.1.1	<i>Frequency Plan.....</i>	<i>18</i>
5.1.2	<i>Compatibility with Other Services.....</i>	<i>18</i>
5.1.3	<i>Fault Isolation Impact on Other Users.....</i>	<i>18</i>
5.1.4	<i>Cable System Terminal Devices.....</i>	<i>19</i>
5.2	RF CHANNEL ASSUMPTIONS.....	19
5.2.1	<i>Transmission Downstream.....</i>	<i>19</i>
5.2.2	<i>Transmission Upstream.....</i>	<i>20</i>
5.3	TRANSMISSION LEVELS.....	21
5.4	FREQUENCY INVERSION.....	21
6	PHYSICAL MEDIA DEPENDENT SUBLAYER SPECIFICATION.....	22
6.1	SCOPE.....	22
6.2	UPSTREAM.....	22
6.2.1	<i>Overview.....</i>	<i>22</i>
6.2.2	<i>Signal Processing Requirements.....</i>	<i>24</i>
6.2.3	<i>Modulation Formats.....</i>	<i>27</i>
6.2.4	<i>R-S Encode.....</i>	<i>27</i>
6.2.5	<i>Upstream R-S Frame Structure for DOCSIS 3.0 Multiple Transmit Channel mode Enabled.....</i>	<i>28</i>
6.2.6	<i>Upstream R-S Frame Structure for DOCSIS 3.0 Multiple Transmit Channel mode Not Enabled.....</i>	<i>30</i>
6.2.7	<i>TDMA Byte Interleaver.....</i>	<i>30</i>
6.2.8	<i>Scrambler (Randomizer).....</i>	<i>33</i>
6.2.9	<i>TCM Encoder.....</i>	<i>34</i>
6.2.10	<i>Preamble Prepend.....</i>	<i>37</i>
6.2.11	<i>Modulation Rates.....</i>	<i>38</i>
6.2.12	<i>S-CDMA Framer and Interleaver.....</i>	<i>38</i>
6.2.13	<i>S-CDMA Framer.....</i>	<i>45</i>
6.2.14	<i>Symbol Mapping.....</i>	<i>48</i>

6.2.15	<i>S-CDMA Spreader</i>	56
6.2.16	<i>Transmit Pre-Equalizer</i>	62
6.2.17	<i>Spectral Shaping</i>	64
6.2.18	<i>Relative Processing Delays</i>	65
6.2.19	<i>Transmit Power Requirements</i>	66
6.2.20	<i>Burst Profiles</i>	74
6.2.21	<i>Burst Timing Convention</i>	80
6.2.22	<i>Fidelity Requirements</i>	81
6.2.23	<i>Upstream Demodulator Input Power Characteristics</i>	90
6.2.24	<i>Upstream Electrical Output from the CM</i>	91
6.2.25	<i>Upstream CM Transmitter Capabilities</i>	92
6.3	DOWNSTREAM	93
6.3.1	<i>Downstream Protocol and Interleaving Support</i>	93
6.3.2	<i>Downstream Electrical Input to CM</i>	93
6.3.3	<i>CM BER Performance</i>	94
6.3.4	<i>Downstream Multiple Receiver Capabilities</i>	95
6.3.5	<i>Non-Synchronous DS Channel Support</i>	99
ANNEX A	TIMING REQUIREMENTS FOR SUPPORTING BUSINESS SERVICES OVER	
DOCSIS	101	
A.1	CMTS	101
A.2	CM	101
ANNEX B	ADDITIONS AND MODIFICATIONS FOR EUROPEAN SPECIFICATION	102
B.1	SCOPE	102
B.1.1	<i>Introduction and Purpose</i>	102
B.1.2	<i>Background</i>	102
B.1.3	<i>Requirements</i>	103
B.1.4	<i>Conventions</i>	103
B.1.5	<i>Organization of Document</i>	103
B.2	REFERENCES	104
B.2.1	<i>Normative References</i>	104
B.2.2	<i>Informative References</i>	104
B.2.3	<i>Reference Acquisition</i>	104
B.3	TERMS AND DEFINITIONS	104
B.4	ABBREVIATIONS AND ACRONYMS	104
B.5	FUNCTIONAL ASSUMPTIONS	104
B.5.1	<i>Equipment Assumptions</i>	104
B.5.2	<i>RF Channel Assumptions</i>	105
B.5.3	<i>Transmission Levels</i>	107
B.5.4	<i>Frequency Inversion</i>	107
B.6	PHYSICAL MEDIA DEPENDENT SUBLAYER SPECIFICATION	107
B.6.1	<i>Scope</i>	107
B.6.2	<i>Upstream</i>	107
B.6.3	<i>Downstream</i>	135
ANNEX C	MPEG HEADER SYNCHRONIZATION AND RECOVERY	140
C.1	MPEG HEADER SYNCHRONIZATION AND RECOVERY IN THE NORTH AMERICAN TECHNOLOGY	
	OPTION	140
C.2	MPEG HEADER SYNCHRONIZATION AND RECOVERY IN THE EUROPEAN TECHNOLOGY OPTION	140
APPENDIX I	EXAMPLE PREAMBLE SEQUENCE	141
I.1	INTRODUCTION	141
I.2	EXAMPLE PREAMBLE SEQUENCE	141

APPENDIX II	S-CDMA FRAMING	143
II.1	CODED SUBSYMBOL NUMBERING	143
II.2	UNCODED SUBSYMBOL NUMBERING	143
II.3	FRAMER OUTPUT NUMBERING	144
II.4	COMMENTS.....	144
APPENDIX III	AMBIENT TEMPERATURE AND WIND LOADING EFFECTS	145
III.1	SYNCHRONIZATION TOLERANCES TO PLANT DELAY VARIATIONS	145
III.2	CHANGE IN PROPAGATION DELAY DUE TO TEMPERATURE CHANGES	146
III.2.1	<i>Fiber Delay Changes Due to Temperature</i>	146
III.2.2	<i>Coaxial Cable Delay Changes Due to Temperature</i>	147
III.2.3	<i>Delay Change Due to Wind</i>	147
APPENDIX IV	DESCRIPTION OF UPSTREAM TRANSMIT CHANNEL SET CAPABILITY: EXAMPLE CALCULATIONS FOR REPORTING AND FIGURING THE NUMBER OF ACTIVE CHANNELS SUPPORTED	148
APPENDIX V	DESCRIPTION OF UPSTREAM CHANNEL POWER CONTROL WITH MULTIPLE UPSTREAM CHANNELS	150
V.1	DOCSIS 2.0 PARAMETERS EXTENDED TO MULTIPLE TRANSMIT CHANNEL MODE.....	150
V.2	NEW PARAMETERS IN DOCSIS 3.0 UPSTREAM POWER CONTROL ("LOADING," PLOAD_N, PLOAD_MIN_ SET, DYNAMIC RANGE WINDOW, PLOW_MULTI).....	150
V.3	EXAMPLE UPSTREAM POWER CONTROL WITH MULTIPLE TRANSMIT CHANNEL MODE ENABLED.....	151
V.4	EXAMPLES REGARDING CONCURRENT AND CONSECUTIVE CHANGES IN P _{R_N} AND P _{LOAD_MIN_SET}	153
APPENDIX VI	EXAMPLE SPURIOUS EMISSIONS NOISE POWER LIMITS WITH MULTIPLE CHANNELS BURSTING	155
APPENDIX VII	ACKNOWLEDGEMENTS	156

Figures

Figure 1-1 - The DOCSIS Network.....	2
Figure 1-2 - Transparent IP Traffic Through the Data-Over-Cable System.....	3
Figure 1-3 - Data-over-Cable Reference Architecture.....	4
Figure 6-1 - Upstream Signal-Processing Sequence	24
Figure 6-2 - TDMA Upstream Transmission Processing	25
Figure 6-3 - S-CDMA Upstream Transmission Processing	26
Figure 6-4 - Example Frame Structures with Flexible Burst Length Mode, DOCSIS 3.0 Operation	28
Figure 6-5 - Byte Interleaver Operation	31
Figure 6-6 - Interleaver Operation for Last Interleaver Block (with Shortened Last Codeword)	32
Figure 6-7 - T Mode Calculations	33
Figure 6-8 - Scrambler Structure	34
Figure 6-9 - Convolutional encoder	34
Figure 6-10 - Repetitive Patterns of Byte Mapping to Symbol Map Bits for TCM	35
Figure 6-11 - Example Byte to Bit Assignment for 64 QAM	36
Figure 6-12 - Example of return to zero bits followed by "0".....	37
Figure 6-13 - Timestamp Snapshot	39
Figure 6-14 - Mini-slot Mapping with Two Codes per mini-slot, 128 Active Codes	42
Figure 6-15 - Mini-slot Mapping with Three Codes per mini-slot, 126 Active Codes	42
Figure 6-16 - Mini-slot Mapping with Four Codes per Mini-slot, 124 Active Codes, Codes 0, 1, 5, and 125 Unused; Selectable Active Codes Mode 2.....	43

Figure 6-17 - S-CDMA Spreader-on and Spreader-off Intervals	44
Figure 6-18 - Subframe structure	46
Figure 6-19 - Symbol Numbering With and Without TCM	48
Figure 6-20 - Symbol Constellations	51
Figure 6-21 - QPSK Gray and Differential Symbol Mapping	52
Figure 6-22 – 8 QAM Symbol Mapping	52
Figure 6-23 – 16 QAM Symbol Mapping	53
Figure 6-24 – 32 QAM Symbol Mapping	53
Figure 6-25 – 64 QAM Symbol Mapping	54
Figure 6-26 - QPSK and 8 QAM TCM Symbol Mapping	54
Figure 6-27 – 16 QAM and 32 QAM TCM Symbol Mapping	55
Figure 6-28 – 64 QAM and 128 QAM TCM Symbol Mapping	55
Figure 6-29 - Code Hopping Random Number Generator	61
Figure 6-30 - Transmit Pre-Equalizer Structure	62
Figure 6-31 - Nominal TDMA Burst Timing	80
Figure 6-32 - Worst-Case TDMA Burst Timing	81
Figure 6-33 - 6 MHz DOCSIS Standard Receive Channel Profile 01	98
Figure 6-34 - Example Manufacturer Receive Channel Profile	99
Figure B-1 - Transmit Pre-Equalizer Structure	113
Figure B-2 - 8 MHz DOCSIS Standard Receive Channel Profile 01	138

Tables

Table 1-1 - DOCSIS 3.0 Series of Specifications	4
Table 5-1 - Assumed Downstream RF Channel Transmission Characteristics	19
Table 5-2 - Assumed Upstream RF Channel Transmission Characteristics	20
Table 6-1 - Burst Size	29
Table 6-2 - Interleaver Operating Parameters	30
Table 6-3 - I/Q Mapping	49
Table 6-4 - Definition of Differential Quadrant Coding	50
Table 6-5 - Maximum Channel Width for Operation with a DOCSIS 3.0 CMTS	64
Table 6-6 - Maximum Channel Width for Operation with a pre-3.0-DOCSIS CMTS	64
Table 6-7 - Constellation Gains and Per Channel Power Limits with One Channel in the Transmit Channel Set	68
Table 6-8 - Constellation Gains and Per Channel Power Limits with Two Channels in the Transmit Channel Set	68
Table 6-9 - Constellation Gains and Per Channel Power Limits with Three or Four Channels in the Transmit Channel Set	68
Table 6-10 – Minimum Transmit Power P_{min} for Multiple Transmit Channel mode Not Enabled	74
Table 6-11 - Burst Profile Attributes	75
Table 6-12 - User Unique Burst Parameters	76
Table 6-13 - Timing Offset for Modulation Rate Changes	79
Table 6-14 - Spurious Emissions	83
Table 6-15 - Adjacent Channel Spurious Emissions Requirements Relative to the Per Channel Transmitted Burst Power Level for Each Channel	84
Table 6-16 - Adjacent Channel Spurious Emissions Goals Relative to the Per Channel Transmitted Burst Power Level for Each Channel	84
Table 6-17 - Spurious Emissions Requirements in the Upstream Frequency Range Relative to the Per Channel Transmitted Burst Power Level for Each Channel	85
Table 6-18 - Spurious Emissions Goals in the Upstream Frequency Range Relative to the Per Channel Transmitted Burst Power Level for Each Channel	85
Table 6-19 – Single Channel Filter Amplitude Distortion	88
Table 6-20 - Upstream Channel Demodulator Input Power Characteristics	90
Table 6-21 –Electrical Output from CM	91
Table 6-22 - Electrical Input to CM	94
Table B-1 - DOCSIS 3.0 Series of Specifications	103
Table B-2 - Assumed Downstream RF Channel Transmission Characteristics	105
Table B-3 - Assumed Upstream RF Channel Transmission Characteristics	106
Table B-4 - Constellation Gains and Per Channel Power Limits with One Channel in the Transmit Channel Set	117
Table B-5 - Constellation Gains and Per Channel Power Limits with Two Channels in the Transmit Channel Set	117
Table B-6 - Constellation Gains and Per Channel Power Limits with Three or Four Channels in the Transmit Channel Set	118
Table B-7 - Minimum Transmit Power P_{min} for Multiple Transmit Channel mode Not Enabled	123
Table B-8 - Burst Profile Attributes	124
Table B-9 - User Unique Burst Parameters	125
Table B-10 - Spurious Emissions	130
Table B-11 - Adjacent Channel Spurious Emissions Requirements Relative to the Per Channel Transmitted Burst Power Level for Each Channel	131
Table B-12 - Adjacent Channel Spurious Emissions Goals Relative to the Per Channel Transmitted Burst Power Level for Each Channel	131

Table B–13 - Spurious Emissions Requirements in the Upstream Operating Frequency Range Relative to the Per Channel Transmitted Burst Power Level for Each Channel 131

Table B–14 - Spurious Emissions Goals in the Upstream Operating Frequency Range Relative to the Per Channel Transmitted Burst Power Level for Each Channel..... 132

Table B–15 - Single Channel Electrical Output from CM 134

Table B–16 - Electrical Input to CM 136

Table III–1 - Allowable Plant Timing Drift 145

1 SCOPE

1.1 Introduction and Purpose

This specification is part of the DOCSIS[®] family of specifications developed by Cable Television Laboratories (CableLabs). In particular, this specification is part of a series of specifications that defines the third generation of high-speed data-over-cable systems. This specification was developed for the benefit of the cable industry, and includes contributions by operators and vendors from North America, Europe, and other regions.

There are differences in the cable spectrum planning practices adopted for different networks in the world. Therefore, two options for physical layer technology are included, which have equal priority and are not required to be interoperable. One technology option is based on the downstream multi-programme television distribution that is deployed in North America using 6 MHz channelling. The other technology option is based on the corresponding European multi-programme television distribution. Both options have the same status, notwithstanding that the document structure does not reflect this equal priority. The first of these options is defined in Sections 5 and 6, whereas the second is defined by replacing the content of those sections with the content of Annex B. Correspondingly, [ITU-T J.83-B] and [CEA-542-B] apply only to the first option, and [EN 300 429] only to the second. Compliance with this document requires compliance with the one or the other of these implementations, not with both. It is not required that equipment built to one option shall interoperate with equipment built to the other.

These optional physical-layer technologies allow operators flexibility in mandated areas of operation, including any frequency planning, EMC (electromagnetic compatibility), and safety requirements. For example, the 6 MHz downstream based option defined in Sections 5 and 6 might be deployable within an 8 MHz channel plan. Compliance with frequency planning and EMC requirements is not covered by this specification and remains the operators' responsibility. In this respect, [FCC15] and [FCC76] are relevant to North America and [EG 201 212], [EN 50083-1], [EN 50083-2], [EN 50083-7], [EN 50083-10], [EN 60950-1], [EN 61000-6-1], and [EN 61000-6-3] are relevant to the European Union.

Backwards compatibility with earlier versions of that technology [DOCSIS RFI 2.0] is only ensured within the same technology options referred to above and not between the two options.

1.2 Background

1.2.1 Broadband Access Network

A coaxial-based broadband access network is assumed. This may take the form of either an all-coax or hybrid-fiber/coax (HFC) network. The generic term "cable network" is used here to cover all cases.

A cable network uses a tree-and-branch architecture with analog transmission. The key functional characteristics assumed in this document are the following:

- Two-way transmission.
- A maximum optical/electrical spacing between the CMTS and the most distant CM of 100 miles in each direction, although typical maximum separation may be 10-15 miles.
- A maximum differential optical/electrical spacing between the CMTS and the closest and most distant modems of 100 miles in each direction, although this would typically be limited to 15 miles.

At a propagation velocity in fiber of approximately 1.5 ns/ft, 100 miles of fiber in each direction results in a round-trip delay of approximately 1.6 ms.

1.2.2 Network and System Architecture

1.2.2.1 The DOCSIS Network

The elements that participate in the provisioning of DOCSIS services are shown in the following figure:

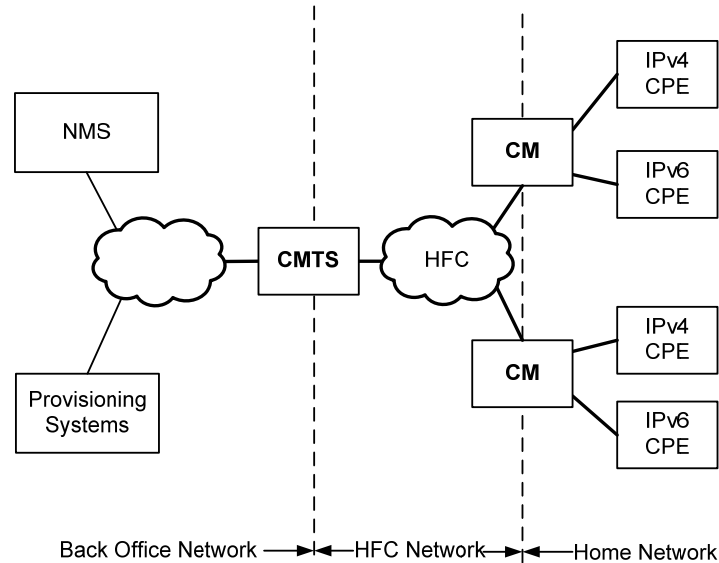


Figure 1–1 - The DOCSIS Network

The CM connects to the operator's HFC network and to a home network, bridging packets between them. Many CPEs' devices can connect to the CMs' LAN interfaces. CPE devices can be embedded with the CM in a single device, or they can be separate standalone devices, as shown in Figure 1–1. CPE devices may use IPv4, IPv6 or both forms of IP addressing. Examples of typical CPE devices are home routers, set-top devices, personal computers, *etc.*

The CMTS connects the operator's back office and core network with the HFC network. Its main function is to forward packets between these two domains, and between upstream and downstream channels on the HFC network.

Various applications are used to provide back office configuration and other support to the devices on the DOCSIS network. These applications use IPv4 and/or IPv6 as appropriate to the particular operator's deployment. The following applications include:

- Provisioning Systems
 - The DHCP servers provide the CM with initial configuration information, including the device IP address(es), when the CM boots.
 - The Config File server is used to download configuration files to CMs when they boot. Configuration files are in binary format and permit the configuration of the CM's parameters. The Config File server is also used to download software upgrades to the CM.
 - The Time Protocol server provides Time Protocol clients, typically CMs, with the current time of day.
 - Certificate Revocation server provides certificate status.

- NMS
 - The SNMP Manager allows the operator to configure and monitor SNMP Agents, typically the CM and the CMTS.
 - The Syslog server collects messages pertaining to the operation of devices.
 - The IPDR Collector server allows the operator to collect bulk statistics in an efficient manner.

1.2.3 Service Goals

As cable operators have widely deployed high-speed data services on cable television systems, the demand for bandwidth has increased. Additionally, networks have scaled to such a degree that IPv4 address constraints are becoming a burden on network operations. To this end, CableLabs' member companies have decided to add new features to the DOCSIS specification for the purpose of increasing channel capacity, enhancing network security, expanding addressability of network elements, and deploying new service offerings.

The DOCSIS system allows transparent bi-directional transfer of Internet Protocol (IP) traffic, between the cable system head-end and customer locations, over an all-coaxial or hybrid-fiber/coax (HFC) cable network. This is shown in simplified form in Figure 1–2.




Figure 1–2 - Transparent IP Traffic Through the Data-Over-Cable System

1.2.4 Statement of Compatibility

This document specifies an interface, commonly referred to as DOCSIS 3.0, which is the third generation of the interface, commonly referred to as DOCSIS 1.x and 2.0. DOCSIS 3.0 MUST be backward- and forward-compatible with equipment built to the previous specifications. DOCSIS 3.0-compliant CMs MUST interoperate seamlessly with DOCSIS 2.0 and DOCSIS 1.x CMTSs, albeit in the 2.0 and 1.x modes, as the case may be. DOCSIS 3.0-compliant CMTSs MUST seamlessly support DOCSIS 2.0 and DOCSIS 1.x CMs.

1.2.5 Reference Architecture

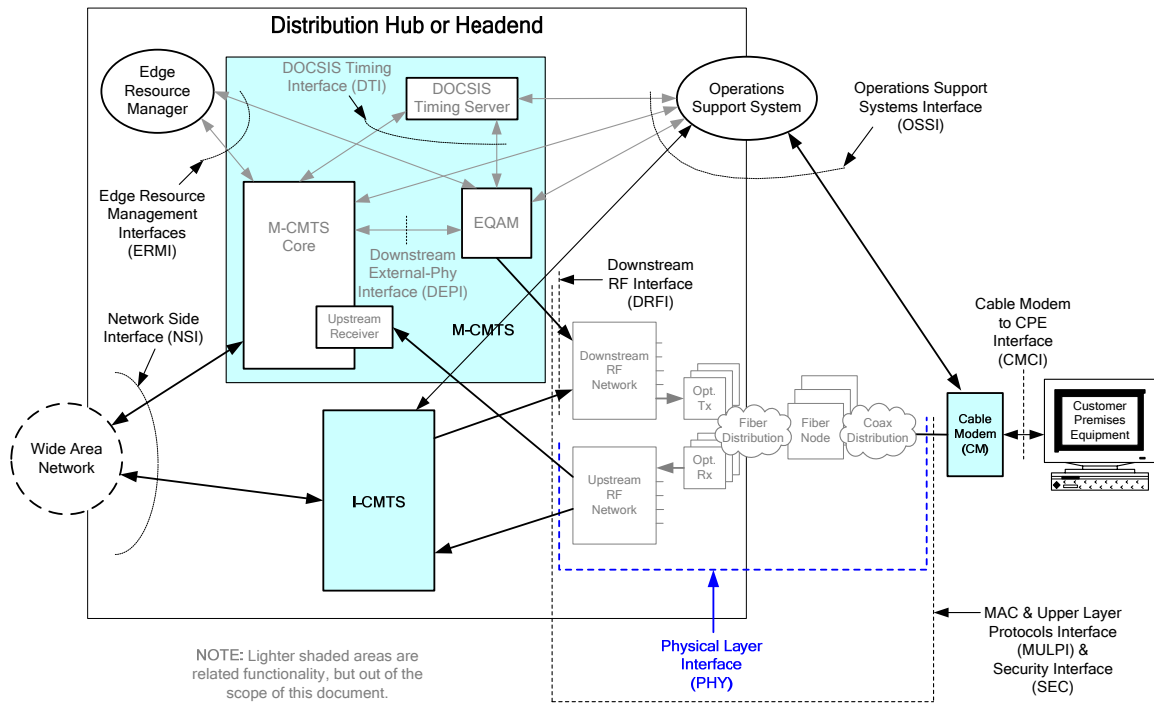


Figure 1-3 - Data-over-Cable Reference Architecture

The reference architecture for data-over-cable services and interfaces is shown in Figure 1-3.

1.2.6 DOCSIS 3.0 Documents

A list of the specifications in the DOCSIS 3.0 series is provided in Table 1-1. For further information, please refer to <http://www.cablemodem.com>.

Table 1-1 - DOCSIS 3.0 Series of Specifications

Designation	Title
CM-SP-PHYv3.0	Physical Layer Specification
CM-SP-MULPIv3.0	Media Access Control and Upper Layer Protocols Interface Specification
CM-SP-OSSIV3.0	Operations Support System Interface Specification
CM-SP-SECv3.0	Security Specification

This specification defines the interface for the physical layer.

1.3 Requirements

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

"MUST"	This word means that the item is an absolute requirement of this specification.
"MUST NOT"	This phrase means that the item is an absolute prohibition of this specification.
"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.

This 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 must comply with all mandatory (MUST and MUST NOT) requirements to be considered compliant with this specification. Support of non-mandatory features and parameter values is optional.

1.4 Conventions

In this specification 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 MSB being the first bit read and the LSB being the last bit read.

1.5 Organization of Document

Section 1 provides an overview of the DOCSIS 3.0 series of specifications including the DOCSIS reference architecture and statement of compatibility.

Section 2 includes a list of normative and informative references used within this specification.

Section 3 defines the terms used throughout this specification.

Section 4 defines the acronyms used throughout this specification.

Section 5 provides a technical overview and lists the DOCSIS 3.0 key features for the functional area of this specification.

Section 6 defines the interface requirements and the performance requirements for the CM downstream and upstream physical layer, and for the CMTS upstream physical layer.

Annex A describes the timing requirements for the CM and CMTS for supporting business services with cable systems using DOCSIS.

Annex B contains the PHY requirements for the European technology option for DOCSIS.

Annex C contains the MPEG header synchronization and recovery requirements.

Appendix I presents an illustrative example of the DOCSIS upstream programmable preamble superstring.

Appendix II presents an algorithmic description of the subsymbol mapping for S-CDMA framing.

Appendix III describes the impact of temperature and wind loading on timing variation of the signaling across a cable plant, and the tolerances and impacts on the DOCSIS communications system to such variations.

Appendix IV contains example calculations for reporting from the CM to the CMTS the number of upstream active channels of various bandwidths supported by the CM.

Appendix V provides an explanation of the power control algorithm with multiple upstream channels.

Appendix VI provides examples illustrating the calculation of the CM noise power limits for the upstream when more than one channel is bursting.

Appendix VII contains acknowledgements to contributors of this specification.

2 REFERENCES

2.1 Normative References

In order to claim compliance with this specification, it is necessary to conform to the following standards and other works as indicated, in addition to the other requirements of this specification. Notwithstanding, intellectual property rights may be required to use or implement such normative references.

- [CEA-542-B] CEA-542-B: CEA Standard: Cable Television Channel Identification Plan," July 2003.
- [DRFI] Downstream Radio Frequency Interface Specification, CM-SP-DRFI-I03-060106, January 6, 2006, Cable Television Laboratories, Inc.
- [EG 201 212] ETSI EG 201 212 V1.2.1: Electrical safety; Classification of interfaces for equipment to be connected to telecommunication networks, November 1998.
- [EN 300 429] ETSI EN 300 429 V1.2.1: Digital Video Broadcasting (DVB); Framing structure, channel coding and modulation for cable systems, April 1998.
- [EN 50083-1] CENELEC EN 50083-1: Cable networks for television signals, sound signals and interactive services -- Part 1: Safety requirements, 2002.
- [EN 50083-10] CENELEC EN 50083-10: Cable networks for television signals, sound signals and interactive services -- Part 10: System performance for return paths, March 2002.
- [EN 50083-2] CENELEC EN 50083-2: Cable networks for television signals, sound signals and interactive services -- Part 2: Electromagnetic compatibility for equipment, 2005.
- [EN 50083-7] CENELEC EN 50083-7: Cable networks for television signals, sound signals and interactive services -- Part 7: System performance, April 1996.
- [EN 60950-1] CENELEC EN 60950-1: Information technology equipment - Safety -- Part 1: General requirements, December 2001.
- [EN 61000-6-1] CENELEC EN 61000-6-4: Electromagnetic compatibility (EMC) -- Part 6-1: Generic standards - Immunity for residential, commercial and light-industrial environments, October 2001.
- [EN 61000-6-3] CENELEC EN 61000-6-3: Electromagnetic compatibility (EMC) -- Part 6-3: Generic standards - Emission standard for residential, commercial and light-industrial environments, 2003.
- [FCC15] Code of Federal Regulations, Title 47, Part 15, October 2005.
- [FCC76] Code of Federal Regulations, Title 47, Part 76, October 2005.
- [ISO 13818] ISO/IEC 13818-1, Information Technology – Generic Coding Of Moving Pictures And Associated Audio Systems Recommendation H.222.0, February 2000.
- [ISO/IEC-61169-24] ISO/IEC-61169-24, 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), 2001.
- [ITU-T J.83-B] Annex B to ITU-T Rec. J.83 (4/97), Digital multi-program systems for television sound and data services for cable distribution, April 1997.
- [MULPI] Media Access Control and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.0-D04-060721, July 21, 2006, Cable Television Laboratories, Inc.
- [OSSI3.0] DOCSIS 3.0 Operations Support System Interface Specification, CM-SP-OSSIv3.0-D02-060504, May 4, 2006, Cable Television Laboratories, Inc.
- [DOCSIS RFI 2.0] Data-Over-Cable Service Interface Specifications, Radio Frequency Interface Specification v2.0, CM-SP-RFIv2.0-I11-060602, June 2, 2006, Cable Television Laboratories, Inc.

- [SCTE 02] ANSI/SCTE 02, Specification for "F" Port, Female Indoor, 2006.
- [TEI] Business Services over DOCSIS TDM Emulation Interface Specification, CM-SP-TEI-I01-060512, May 12, 2006, Cable Television Laboratories, Inc.

2.2 Informative References

- [CableLabs1] Digital Transmission Characterization of Cable Television Systems, Cable Television Laboratories, Inc., November 1994.
- [NCTA] NCTA Recommended Practices for measurements on Cable Television Systems – National Cable Television Association, Washington DC, 2nd Edition, revised October 1993.

2.3 Reference Acquisition

Cable Television Laboratories, Inc., <http://www.cablelabs.com/>

CENELEC: European Committee for Electro-technical Standardization, <http://www.cenelec.org>

EIA: Electronic Industries Alliance, http://www.eia.org/new_contact/

ETSI: European Telecommunications Standards Institute,
http://www.etsi.org/services_products/freestandard/home.htm

Internet Engineering Task Force (IETF), <http://www.ietf.org/>

ISO: International Organization for Standardization (ISO), <http://www.iso.org/iso/en/xsite/contact/contact.html>

ITU: International Telecommunications Union (ITU), <http://www.itu.int/home/contact/index.html>

3 TERMS AND DEFINITIONS

This specification uses the following terms:

Active Codes	The set of spreading codes which carry information in an S-CDMA upstream. The complementary set, the unused codes, are idle and are not transmitted. Reducing the number of active codes below the maximum value of 128 may provide advantages including more robust operation in the presence of colored noise.
Allocation	A group of contiguous mini-slots in a MAP which constitutes a single transmit opportunity.
Availability	In cable television systems, availability is the long-term ratio of the actual RF channel operation time to scheduled RF channel operation time (expressed as a percent value) and is based on a bit error rate (BER) assumption.
Bandwidth Allocation Map (MAP)	The MAC Management Message that the CMTS uses to allocate transmission opportunities to cable modems (MAP).
Bit Error Rate (BER)	The percentage of bits that have errors relative to the total number of bits received in a transmission, usually expressed as ten to a negative power.
Burst	A single continuous RF signal from the upstream transmitter, from transmitter on to transmitter off.
Cable Modem (CM)	A modulator-demodulator at subscriber locations intended for use in conveying data communications on a cable television system. (CM).
Cable Modem Termination System (CMTS)	Cable modem termination system, located at the cable television system head-end or distribution hub, which provides complementary functionality to the cable modems to enable data connectivity to a wide-area network.
Capture Bandwidth (CBW)	The sum of the Tuning Bands in the TB List in MHz.
Carrier Hum Modulation	The peak-to-peak magnitude of the amplitude distortion relative to the RF carrier signal level due to the fundamental and low-order harmonics of the power-supply frequency.
Carrier-to-Noise Ratio (C/N)	The ratio of signal power to noise power in the defined measurement bandwidth. For digital modulation, $C/N = E_s/N_o$, the energy-per-symbol to noise-density ratio; the signal power is measured in the occupied bandwidth, and the noise power is normalized to the modulation-rate bandwidth. For video, the measurement bandwidth is 4 MHz (C/N).
Channel	(See RF Channel.)
Channel Bonding	A logical process that combines the data packets received on multiple independent channels into one higher-speed data stream. Channel bonding can be implemented independently on upstream channels or downstream channels.
Chip	Each of the 128 bits comprising the S-CDMA spreading codes.
Chip Rate	The rate at which individual chips of the S-CDMA spreading codes are transmitted. (1280 to 5120 kHz). E_s/N_o
Codeword	An element of an error-correcting code used to detect and correct transmission errors.
Codeword Error Rate	The ratio of the number of uncorrectable code words to the total number of code words sent without errors, with corrected errors and with uncorrectable errors.

Composite Second Order Beat (CSO)	The peak of the average level of distortion products due to second order nonlinearities in cable system equipment.
Composite Triple Beat (CTB)	The peak of the average level of distortion components due to third-order nonlinearities in cable system equipment.
Cross-modulation	A form of television signal distortion where modulation from one or more television channels is imposed on another channel or channels.
Customer Premises Equipment (CPE)	Equipment at the end user's premises; may be provided by the end user or the service provider.
Decibel-Millivolt (dBmV)	A dB measurement system wherein 0 dBmV is defined as 1 millivolt over 75 ohms.
Decibels (dB)	A unit to measure the relative levels of current, voltage or power. An increase of 3 dB indicates a doubling of power, an increase of 10 dB indicates a 10x increase in power, and an increase of 20 dB indicates a 100x increase in power.
Demodulator Module	A physical entity in the CM that demodulates a block of one or more contiguous channels of a single bandwidth (6 MHz or 8 MHz) within the output from a single tuner.
DOCSIS 1.x	Abbreviation for "DOCSIS 1.0 or 1.1." DOCSIS stands for Data-Over-Cable Service Interface Specifications.
DOCSIS 2.0 Mode	A CM operates in this mode when: 1) Multiple Transmit Channel (MTC) Mode is disabled; 2) the Enable 2.0 Mode configuration setting in the REG-RSP is set to 1 (Enable) explicitly or by default; and 3) it operates on at least one upstream channel using the burst descriptors associated with IUC 9, 10, and 11 as opposed to IUC 5 and 6. A CM is enabled for DOCSIS 2.0 Mode when the Enable 2.0 Mode configuration setting in the REG-RSP is set to 1 (Enable). A CM may be enabled for DOCSIS 2.0 Mode but may not be operating in DOCSIS 2.0 Mode. When a CM has MTC Mode enabled, the CM is not considered to be in DOCSIS 2.0 Mode even if some of the upstream channels it is using are operating with post-1.1 DOCSIS physical layer mechanisms. Therefore, "DOCSIS 2.0 Mode" does not have relevance for a CM operating in MTC Mode.
Downstream	In cable television, the direction of transmission from the head-end to the subscriber.
Downstream Channel	Physical layer characteristics and MAC layer parameters and functions associated to a DOCSIS forward channel.
Dynamic Host Configuration Protocol (DHCP)	An Internet protocol used for assigning network-layer (IP) addresses.
Dynamic Range	The ratio between the greatest signal power that can be transmitted over a multichannel analog transmission system without exceeding distortion or other performance limits, and the least signal power that can be utilized without exceeding noise, error rate or other performance limits.
Dynamic Range Window (DRW)	A 12 dB range defining the maximum power difference between multiple transmitters in a CM in Multiple Transmit Channel mode.
Electronic Industries Alliance (EIA)	A voluntary body of manufacturers which, among other activities, prepares and publishes standards.
F Connector (F conn)	A male F-connector is the final piece of hardware (familiar to subscribers) on a drop cable. It is cylindrical with a center pin sticking out that plugs into the female F-connector on a set-top box, cable ready TV or VCR.

Floor	The Floor function returns the first integer that is less than or equal to a given value.
Forward Channel	The direction of RF signal flow away from the head-end toward the end user; equivalent to Downstream.
Forward Error Correction (FEC)	FEC enables the receiver to detect and fix errors to packets without the need for the transmitter to retransmit packets.
Frame	See MAC frame, S-CDMA frame, and MPEG frame.
Frequency Division Multiple Access (FDMA)	A multiple access technology that separates users by putting each traffic channel on a discrete frequency band.
Group Delay	The difference in transmission time between the highest and lowest of several frequencies through a device, circuit or system.
Guard Band	Minimum time, measured in modulation symbols, allocated between bursts in the upstream referenced from the symbol center of the last symbol of a burst to the symbol center of the first symbol of the following burst. The guard band should be at least the duration of five symbols plus the maximum system timing error. The guard band should be at least the duration of five symbols plus the maximum system timing error.
Guard Time	Guard time, measured in modulation symbols, is similar to the guard band, except that it is measured from the end of the last symbol of one burst to the beginning of the first symbol of the preamble of an immediately following burst. Thus, the guard time is equal to the guard band – 1.
Harmonic Related Carrier (HRC)	A method of spacing television channels on a cable television system in exact 6 MHz increments, with all carrier frequencies harmonically related to a common reference.
Head-end	The central location on the cable network that is responsible for injecting broadcast video and other signals in the downstream direction. See also Master Head-End, Distribution Hub.
Header	Protocol control information located at the beginning of a protocol data unit.
Hertz (Hz)	A unit of frequency equivalent to one cycle per second. See also kilohertz (kHz) and megahertz (MHz).
Hum Modulation	Undesired modulation of the television visual carrier by the fundamental or low-order harmonics of the power supply frequency or other low-frequency disturbances.
Hybrid Fiber/Coaxial System (HFC)	A broadband bidirectional shared-media transmission system using fiber trunks between the head-end and the fiber nodes, and coaxial distribution from the fiber nodes to the customer locations.
Impulse Noise	Noise characterized by non-overlapping transient disturbances.
Incremental Related Carriers (IRC)	A method of spacing NTSC television channels on a cable television system in which all channels except 5 and 6 correspond to the standard channel plan, used to reduce composite triple beat distortions.
International Electrotechnical Commission (IEC)	An international standards body.
International Organization for Standardization (ISO)	An international standards body, commonly known as the International Standards Organization.
Internet Engineering Task Force (IETF)	A body responsible, among other things, for developing standards used in the Internet.

Internet Protocol (IP)	The computer network protocol (analogous to written and verbal languages) that all machines on the Internet must know so that they can communicate with one another. IP is a layer 3 (network layer) protocol in the OSI model. The vast majority of IP devices today support IP version 4 (IPv4) defined in RFC-791, although support for IP version 6 (IPv6, RFC-2460) is increasing.
Interval Usage Code (IUC)	A field in MAPs and UCDs to link burst profiles to grants.
Latency	The time, expressed in quantity of symbols, taken for a signal element to pass through a device.
Layer	A subdivision of the Open System Interconnection (OSI) architecture, constituted by subsystems of the same rank.
Local Area Network (LAN)	A non-public data network in which serial transmission is used for direct data communication among data stations located on the user's premises.
Logical (Upstream) Channel	A MAC entity identified by a unique channel ID and for which bandwidth is allocated by an associated MAP message. A physical upstream channel may support multiple logical upstream channels. The associated UCD and MAP messages completely describe the logical channel.
MAC Frame	MAC header plus optional protocol data unit
Maximum Downstream Bonded Channels (MDBC)	Maximum number of downstream bonded channels supported by the cable modem. (See Channel Bonding)
Media Access Control (MAC)	The MAC sublayer is the part of the data link layer that supports topology-dependent functions and uses the services of the Physical Layer to provide services to the Logical Link Control (LLC) sublayer. It can also be a component of a networking software stack. In the OSI 7-layer model, the Media Access Control is a part of layer 2, the data link layer.
Megahertz (MHz)	One million cycles per second.
Micro-reflections	Echoes in the forward or reverse transmission path due to impedance mismatches between the physical plant components. Micro-reflections are distinguished from discrete echoes by having a time difference (between the main signal and the echo) on the order of one microsecond. Micro-reflections cause departures from ideal amplitude and phase characteristics for the transmission channel.
Microsecond (μs)	One millionth of a second.
Millisecond (ms)	One thousandth of a second.
Millivolt (mV)	One thousandth of a volt.
Mini-slot	A mini-slot is an integer multiple of 6.25 μ s increments.
Modulation Error Ratio (MER)	MER measures the cluster variance in dB caused by the transmit waveform. It includes the effects of ISI, spurious, phase noise, and all other transmitter degradations.
Modulation Rate	The signaling rate of the upstream modulator (1280 to 5120 kHz). In S-CDMA it is the chip rate. In TDMA, the channel symbol rate.
National Cable Telecommunications Association (NCTA)	A voluntary association of cable television operators which, among other things, provides guidance on measurements and objectives for cable television systems in the USA.
National Television Systems Committee (NTSC)	Committee which defined the analog color television broadcast standard used today in North America.

Number of Allocated Codes	The total number of codes which a single CM uses in a single S-CDMA frame. This number is determined by the size of the grants in mini-slots and the mapping of these mini-slots to S-CDMA frames (note that a CM may receive multiple grants which are mapped to a single S-CDMA frame). The number of allocated codes can be in the range of the number of Codes per Mini-slot to the number of active codes, and may vary from frame to frame, but is constant within an S-CDMA frame.
Phase Noise	Rapid, short-term, random fluctuations in the phase of a wave, caused by time domain instabilities.
Physical Layer	Layer 1 in the Open System Interconnection (OSI) 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 (PHY).
Physical Media Dependent Sublayer (PMD)	A sublayer of the Physical Layer which is concerned with transmitting bits or groups of bits over particular types of transmission link between open systems and which entails electrical, mechanical and handshaking procedures (PMD).
Picosecond (ps)	One trillionth of a second
Primary Channel	For CMTS/CM systems using downstream bonded channels, the CMTS MUST designate a single channel in the bonding group as the primary channel with the timing information such as timestamps or timestamp snapshots.
Primary Downstream Channel	The downstream channel from which a CM derives CMTS master clock timing for upstream transmission. All other concurrently received channels are called "secondary downstream channels."
Protocol	A set of rules and formats that determines the communication behavior of layer entities in the performance of the layer functions.
Quadrature Amplitude Modulation (QAM)	A method of modulating digital signals onto a radio-frequency carrier signal involving both amplitude and phase coding.
Quadrature Phase Shift Keying (QPSK)	A method of modulating digital signals onto a radio-frequency carrier signal using four phase states to code two digital bits.
Radio Frequency (RF)	In cable television systems, electromagnetic signals in the range 5 to 1000 MHz (RF).
Radio Frequency Channel (RFC)	The frequency spectrum occupied by a signal. Usually specified by center frequency and bandwidth parameters.
Return Loss	The parameter describing the attenuation of a guided wave signal (e.g., via a coaxial cable) returned to a source by a device or medium resulting from reflections of the signal generated by the source.
Reverse Channel	The direction of signal flow towards the head-end, away from the subscriber; equivalent to Upstream.
Root Mean Square (RMS)	A mathematical method of computing an "average" magnitude of n elements by taking the square root of the mean of the elements squared.
S-CDMA Frame	A two dimensional representation of mini-slots, where the dimensions are codes and time. An S-CDMA frame is composed of p active codes in the code dimension and K spreading intervals in the time dimension. Within the S-CDMA frame, the number of mini-slots is determined by the number of codes per mini-slot (c) and p, the number of active codes in the S-CDMA frame. Each S-CDMA frame thus contains s mini-slots, where $s=p/c$, and each mini-slot contains $c*K$ information (QAM) symbols.

S-CDMA Subframe	A subframe is a vertically-smaller subset of an S-CDMA frame over which interleaving is performed, where the vertical dimension is R' codes, where $R' \leq p$ (the number of active codes). A subframe is generally used to constrain the interleaving region to be of a similar size to the Reed-Solomon codeword in order to provide protection from impulse noise.
Selectable Active Codes (SAC)	A methodology to determine the set of active codes and its complement, the set of unused codes. In SAC mode 1, a consecutive set of codes starting with code 0 are unused. In SAC mode 2, the active codes are selectable via a 128-bit string.
Service Identifier (SID)	A Service Flow Identifier assigned by the CMTS (in addition to a Service Flow Identifier) to an Active or Admitted Upstream Service Flow. [14 bits] (SID).
Spread Symbol	At the output of the spreader, a group of 128 chips which comprises a single S-CDMA spreading code, and is the result of spreading a single information (QAM constellation) symbol, which is referred to as a "spread symbol."
Spreader-Off S-CDMA Burst	A transmission from a single CM in a spreader-off frame on an S-CDMA channel defined by the time in which the cable modem's transmitter turns on to the time it turns off. There will generally be several spreader off bursts in a spreader-off frame.
Spreader-Off S-CDMA Frame	TDMA mini-slots on an S-CDMA channel in which the spreader is turned off. These are differentiated from TDMA bursts on a TDMA channel in that, for example, the number of mini-slots per spreader-off S-CDMA burst frame is constrained to be the same as the number of mini-slots in a spreader-on S-CDMA frame (s). This number of mini-slots will be less than the number of TDMA mini-slots in a TDMA channel over the same time interval if the number of active codes is significantly less than 128.
Spreading Codes	A family of orthogonal digital code words used in S-CDMA direct-sequence spread-spectrum modulation
Spreading Interval	The period of a spread symbol (128 chips) is called a "spreading interval."
Sublayer	A subdivision of a layer in the Open System Interconnection (OSI) reference model.
Subscriber	(See End User).
Synchronous-Code Division Multiple Access (S-CDMA)	A multiple access physical layer technology in which different transmitters can share a channel simultaneously. 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.
Tick	6.25 μ s time intervals that are the reference for upstream mini-slot definition and upstream transmission times.
Time Division Multiple Access (TDMA)	A digital technology that enables a large number of users to access, in sequence, a single radio frequency channel without interference by allocating unique time slots to each user within each channel.
Transit Delay	The time difference between the instant at which the first bit of a PDU crosses one designated boundary, and the instant at which the last bit of the same PDU crosses a second designated boundary.
Tuner Module	A physical entity in the CM that converts a block of one or more contiguous channels of a single bandwidth (6 MHz or 8 MHz) from the RF input from the cable plant to an intermediate frequency suitable for distribution to a Demodulator Module.

Tuning Band (TB)	Defines a single continuous frequency interval, in MHz, located anywhere in the downstream band (108 MHz to 870 MHz).
Tuning Band List (TB List)	A list of one or more Tuning Bands supported by the CM that defines the cable modem tuning capabilities
Type/Length/Value (TLV)	An encoding of three fields, in which the first field indicates the type of element, the second the length of the element, and the third field the value of the element.
Upstream	The direction from the subscriber location toward the head-end.
Upstream Channel	Physical layer characteristics and MAC layer parameters and functions associated to a DOCSIS reverse channel.
Upstream Channel Descriptor (UCD)	The MAC Management Message used to communicate the characteristics of the upstream physical layer to the cable modems.

4 ABBREVIATIONS AND ACRONYMS

This specification uses the following abbreviations:

BER	Bit Error Rate
CM	Cable Modem
CMTS	Cable Modem Termination System
C/N	Carrier to Noise Ratio
CPE	Customer Premises Equipment
DHCP	Dynamic Host Configuration Protocol
DOCSIS	Data-Over-Cable Service Interface Specifications
EMC	Electromagnetic compatibility
EIA	Electronic Industries Alliance
FDMA	Frequency Division Multiple Access
FEC	Forward Error Correction
HFC	Hybrid Fiber/Coax System
HRC	Harmonic Related Carriers
Hz	Hertz
IEC	International Electrotechnical Commission
IETF	Internet Engineering Task Force
IP	Internet Protocol
IRC	Incremental Related Carriers
ISI	Inter-Symbol Interference
ISO	International Standards Organization
IUC	Interval Usage Code
LAN	Local Area Network
LSB	Least Significant Bit
LFSR	Linear Feedback Shift Register
MAC	Media Access Control
MAP	Bandwidth Allocation Map
MDBC	Maximum Downstream Bonded Channels
MSB	Most Significant Bit
MSC	Maximum Scheduled Codes

MER	Modulation Error Ratio
MPEG	Moving Picture Experts Group
MTC	Multiple Transmit Channel
NCTA	National Cable Telecommunications Association
NTSC	National Television Systems Committee
OSI	Open Systems Interconnection
PHY	Physical Layer
PLL	Phase Locked Loop
PMD	Physical Media Dependent sublayer
QAM	Quadrature Amplitude Modulation
QPSK	Quadrature Phase-Shift Keying
RF	Radio Frequency
RFC	Radio Frequency Channel
RMS	Root Mean Square
R-S	Reed Solomon
SAC	Selectable Active Codes
S-CDMA	Synchronous-Code Division Multiple Access
SID	Service Identifier
SNMP	Simple Network Management Protocol
TB	Tuning Band
TCM	Trellis Code Modulation
TCS	Transmit Channel Set
TDMA	Time Division Multiple Access
TLV	Type/Length/Value
UCD	Upstream Channel Descriptor

5 FUNCTIONAL ASSUMPTIONS

This section describes the characteristics of a cable television plant, assumed to be for the purpose of operating a data-over-cable system. It is not a description of CMTS or CM parameters. The data-over-cable system **MUST** be interoperable within the environment described in this section.

Whenever a reference in this section 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 NTSC analog signals in 6 MHz channels does not imply that such signals are physically present.

5.1 Equipment Assumptions

5.1.1 Frequency Plan

In the downstream direction, the cable system is assumed to have a pass band with a lower edge of either 54 MHz or 108 MHz, and an upper edge that is implementation-dependent but is typically in the range of 300 to 1002 MHz. Within that pass band, NTSC analog television signals in 6 MHz channels are assumed present on the standard, HRC or IRC frequency plans of [CEA-542-B], as well as other narrowband and wideband digital signals.

In the upstream direction, the cable system may have a 5-30 MHz, 5-42 MHz or 5-85 MHz pass band. NTSC analog television signals in 6 MHz channels may be present, as well as other signals.

5.1.2 Compatibility with Other Services

The CM and CMTS **MUST** coexist with any services on the cable network.

In particular,

CM and CMTS **MUST** 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; and

CM and CMTS **MUST NOT** cause harmful interference to any other services that are assigned to the cable network in spectrum outside of that allocated to the CMTS.

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).

5.1.3 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 Section 5.1.2 above.

5.1.4 Cable System Terminal Devices

The CM MUST meet and preferably exceed all applicable regulations for Cable System Termination Devices and Cable Ready Consumer Equipment as defined in FCC Part 15 [FCC15] and Part 76 [FCC76]. None of these specific requirements may be used to relax any of the specifications contained elsewhere within this document.

5.2 RF Channel Assumptions

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 this specification, MUST be interoperable on cable networks having characteristics defined in this section. 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.

5.2.1 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 6 MHz channel bandwidth for carrier 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 a cable TV system. 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. No combination of the following parameters will exceed any stated interface limit defined elsewhere in this specification.

Table 5-1 - Assumed Downstream RF Channel Transmission Characteristics

Parameter	Value
Frequency range	Cable system normal downstream operating range is from 50 MHz to 1002 MHz. However, the values in this table apply only at frequencies ≥ 108 MHz (including Pre-3.0 DOCSIS modes)
RF channel spacing (design bandwidth)	6 MHz
Transit delay from head-end to most distant customer	≤ 0.800 ms (typically much less)
Carrier-to-noise ratio in a 6 MHz band	Not less than 35 dB ^{1,2}
Carrier-to-Composite triple beat distortion ratio	Not less than 41 dB ^{1,2}
Carrier-to-Composite second order distortion ratio	Not less than 41 dB ^{1,2}
Carrier-to-Cross-modulation ratio	Not less than 41 dB ^{1,2}
Carrier-to-any other discrete interference (ingress)	Not less than 41 dB ^{1,2}
Amplitude ripple	3 dB within the design bandwidth ¹
Group delay ripple in the spectrum occupied by the CMTS	75 ns within the design bandwidth ¹
Micro-reflections bound for dominant echo	-10 dBc @ ≤ 0.5 μ s -15 dBc @ ≤ 1.0 μ s -20 dBc @ ≤ 1.5 μ s -30 dBc @ > 1.5 μ s ¹
Carrier hum modulation	Not greater than -26 dBc (5%) ¹
Burst noise	Not longer than 25 μ s at a 10 Hz average rate ¹

Parameter	Value
Maximum analog video carrier level at the CM input	17 dBmV
Maximum number of analog carriers	121
¹ . Measurement methods defined in [NCTA] or [CableLabs1]. ² . Measured relative to a QAM signal that is equal to the nominal video level in the plant.	

5.2.2 Transmission Upstream

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 this specification. Transmission is from the CM output at the customer location to the head-end.

Table 5-2 - Assumed Upstream RF Channel Transmission Characteristics

Parameter	Value
Frequency range	5 to 42 MHz edge to edge or 5 to 85 MHz edge to edge
Transit delay from head-end to most distant customer	≤ 0.800 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 ¹ .
Carrier hum modulation	Not greater than -23 dBc (7.0%)
Burst noise	Not longer than 10 μs at a 1 KHz average rate for most cases ^{2,3}
Amplitude ripple across upstream operating frequency range	0.5 dB/MHz
Group delay ripple across upstream operating frequency range	200 ns/MHz
Micro-reflections – single echo	-10 dBc @ ≤ 0.5 μs -20 dBc @ ≤ 1.0 μs -30 dBc @ > 1.0 μs
Seasonal and diurnal reverse gain (loss) variation	Not greater than 14 dB min to max
¹ . Ingress avoidance or tolerance techniques may be used to ensure operation in the presence of time-varying discrete ingress signals that could be as high as 10 dBc. The ratios are guaranteed only within the digital carrier channels. ² . Amplitude and frequency characteristics sufficiently strong to partially or wholly mask the data carrier. ³ . Impulse noise levels more prevalent at lower frequencies (<15 MHz)	

5.2.2.1 Availability

Typical cable network availability is considerably greater than 99%.

5.3 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.

5.4 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 PHYSICAL MEDIA DEPENDENT SUBLAYER SPECIFICATION

6.1 Scope

This section applies to the first technology option referred to in Section 1.1. For the second option, refer to Annex B.

This specification defines the electrical characteristics and signal processing operations for a cable modem (CM) and Cable Modem Termination System (CMTS). It is the intent of this specification 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 this specification to imply any specific implementation.

6.2 Upstream

6.2.1 Overview

The upstream Physical Media Dependent (PMD) sublayer uses an FDMA/TDMA (herein called TDMA mode) or FDMA/TDMA/S-CDMA (herein called S-CDMA mode) burst type format, which provides six modulation rates and multiple modulation formats. The use of TDMA or S-CDMA mode is configured by the CMTS via MAC messaging.

FDMA (frequency division multiple access) indicates that multiple RF channels are assigned in the upstream band. A CM transmits on one or more RF channels and may be reconfigured to change channels.

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

The CM reports its maximum number of upstream channels capability, and certain other of its capability characteristics, to the CMTS (Section 6.2.25).

The CM MUST be able to operate each channel in the Transmit Channel Set, simultaneously, anywhere in the upstream band, subject to restrictions on transmit power across the channels and on reconfiguration of certain transmit properties (see Sections 6.2.19 and 6.2.20, and subsections of each). The CMTS MUST be capable of assigning and receiving each RF channel anywhere in the upstream band. The CMTS MUST set the number of assigned channels and assigned channels' center frequency, and all other channel attributes. The CMTS MAY change the number of assigned channels and the channel attributes. Each RF channel has its own set of UCD parameters as defined in [MULPI] Section 6.4.3.

TDMA (time division multiple access) indicates that upstream transmissions have a burst nature. A given RF channel is shared by multiple CMs via the dynamic assignment of time slots. S-CDMA (synchronous code division multiple access) indicates that multiple CMs can transmit simultaneously on the same RF channel and during the same TDMA time slot, while being separated by different orthogonal codes.

In this document, the following naming conventions are used. For TDMA, the term "modulation rate" refers to the RF channel symbol rate (160 to 5120 ksym/s). For S-CDMA, the term "modulation rate" refers to the "chip rate," which is the rate (1280 to 5120 kHz) of the individual elements (chips) of the S-CDMA spreading code. Modulation rates are represented in units of "Hz" denoting the number of symbols per second in TDMA mode or the number of chips per second in S-CDMA mode. The "modulation interval" is the symbol period (TDMA mode) or chip period (S-CDMA mode) and is the reciprocal of the modulation rate. At the output of the spreader, a group of 128 chips which comprise a single S-CDMA spreading code, and are the result of spreading a single information (QAM constellation) symbol is referred to as a "spread symbol." The period of a spread symbol (128 chips) is called a

"spreading interval." A "burst" is a physical RF transmission that contains a single preamble plus data, and (in the absence of preceding and following bursts) exhibits RF energy ramp-up and ramp-down.

In some cases logical zeros or logical ones are used to pad data blocks; this indicates data with zero-valued or one-valued binary bits, which result in non-zero transmitted RF energy. In other cases a numerical zero is used; this denotes, for example, symbols which result in zero transmitted RF energy (after ramp-up and ramp-down are taken into account).

The modulation format includes pulse shaping for spectral efficiency, is carrier-frequency agile, and has selectable output power level.

Each burst supports a flexible modulation order, modulation rate, preamble, randomization of the payload, and programmable FEC encoding.

All of the upstream transmission parameters associated with burst transmission outputs from the CM are configurable by the CMTS via MAC messaging. Many of the parameters are programmable on a burst-by-burst basis.

The PMD sublayer can support a near-continuous mode of transmission, wherein ramp-down of one burst MAY overlap the ramp-up of the following burst, so that the transmitted envelope is never zero. In TDMA mode, the system timing of the TDMA transmissions from the various CMs MUST provide that the center of the last symbol of one burst and the center of the first symbol of the preamble of an immediately following burst are separated by at least the duration of five symbols. The guard band MUST be greater than or equal to the duration of five symbols plus the maximum timing error. Timing error is contributed by both the CM and CMTS. CM timing performance is specified in Section 6.2.20.1. Maximum timing error and guard band may vary with CMTSs from different vendors. The term guard time is similar to the guard band, except that it is measured from the end of the last symbol of one burst to the beginning of the first symbol of the preamble of an immediately following burst. Thus, the guard time is equal to the guard band – 1.

The PMD sublayer also supports a synchronous mode of transmission when using S-CDMA, wherein ramp-down of one burst MAY completely overlap the ramp-up of the following burst, so that the transmitted envelope is never zero. There is no guard time for transmission on S-CDMA channels. The system timing of the S-CDMA transmissions from the various CMs MUST provide adequate timing accuracy so that different CMs do not appreciably interfere with each other. S-CDMA utilizes precise synchronization so that multiple CMs can transmit simultaneously.

The upstream modulator is part of the cable modem which interfaces with the cable network. The modulator contains the electrical-level modulation function and the digital signal-processing function; the latter provides the FEC, preamble prepend, symbol mapping, and other processing steps.

At the Demodulator, similar to the Modulator, there are two basic functional components: the demodulation function and the signal processing function. The Demodulator resides in the CMTS and there is one demodulation function (not necessarily an actual physical demodulator) for each carrier frequency in use. The demodulation function receives all bursts on a given frequency.

The demodulation function of the Demodulator accepts a varying-level signal centered around a commanded power level and performs symbol timing and carrier recovery and tracking, burst acquisition, and demodulation. Additionally, the demodulation function provides an estimate of burst timing relative to a reference edge, an estimate of received signal power, may provide an estimate of signal-to-noise ratio, and may engage adaptive equalization to mitigate the effects of a) echoes in the cable plant, b) narrowband ingress, and c) group delay. The signal-processing function of the Demodulator performs the inverse processing of the signal-processing function of the Modulator. This includes accepting the demodulated burst data stream and decoding, etc. The signal-processing function also provides the edge-timing reference and gating-enable signal to the demodulators to activate the burst acquisition for each assigned burst slot. The signal-processing function may also provide an indication of successful decoding, decoding error, or fail-to-decode for each codeword and the number of corrected Reed-Solomon symbols

in each codeword. For every upstream burst, the CMTS has a prior knowledge of the exact burst length in modulation intervals (see Sections 6.2.5.1, 6.2.5.2, 6.2.6, 6.2.20, and [MULPI], in the MAC Service ID section).

6.2.2 Signal Processing Requirements

The signal processing order for each burst packet type **MUST** be compatible with the sequence shown in Figure 6-1. For TDMA mode, the signal processing order for each burst packet type **MUST** follow the order of steps in Figure 6-2. For S-CDMA mode, the signal processing order for each burst packet type **MUST** follow the order of steps in Figure 6-3.

The blocks used only in S-CDMA consist of a TCM encoder, S-CDMA framer, and S-CDMA spreader. The TCM encoder provides trellis modulation encoding of data symbols and is described in Section 6.2.9. The S-CDMA framer maps mini-slots into code resources, provides interleaving of data symbols, and is described in Section 6.2.11.2. The S-CDMA spreader spreads S-CDMA framed symbols for transmission and is described in Section 6.2.15, "S-CDMA Spreader."

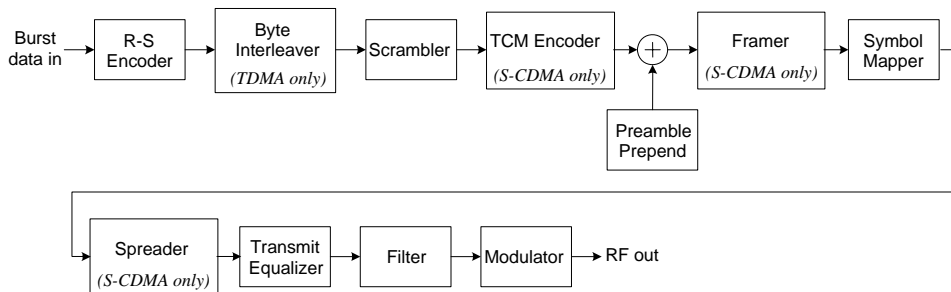


Figure 6-1 - Upstream Signal-Processing Sequence

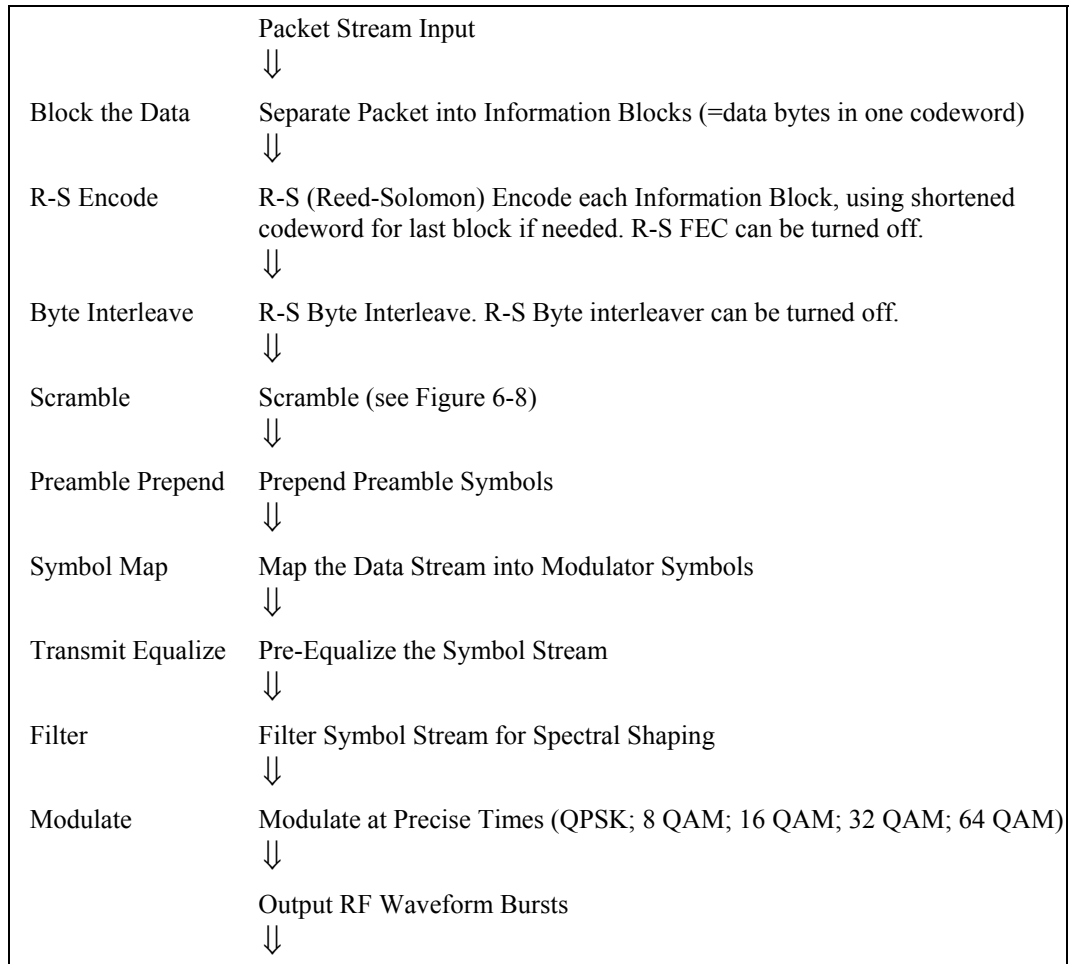


Figure 6-2 - TDMA Upstream Transmission Processing

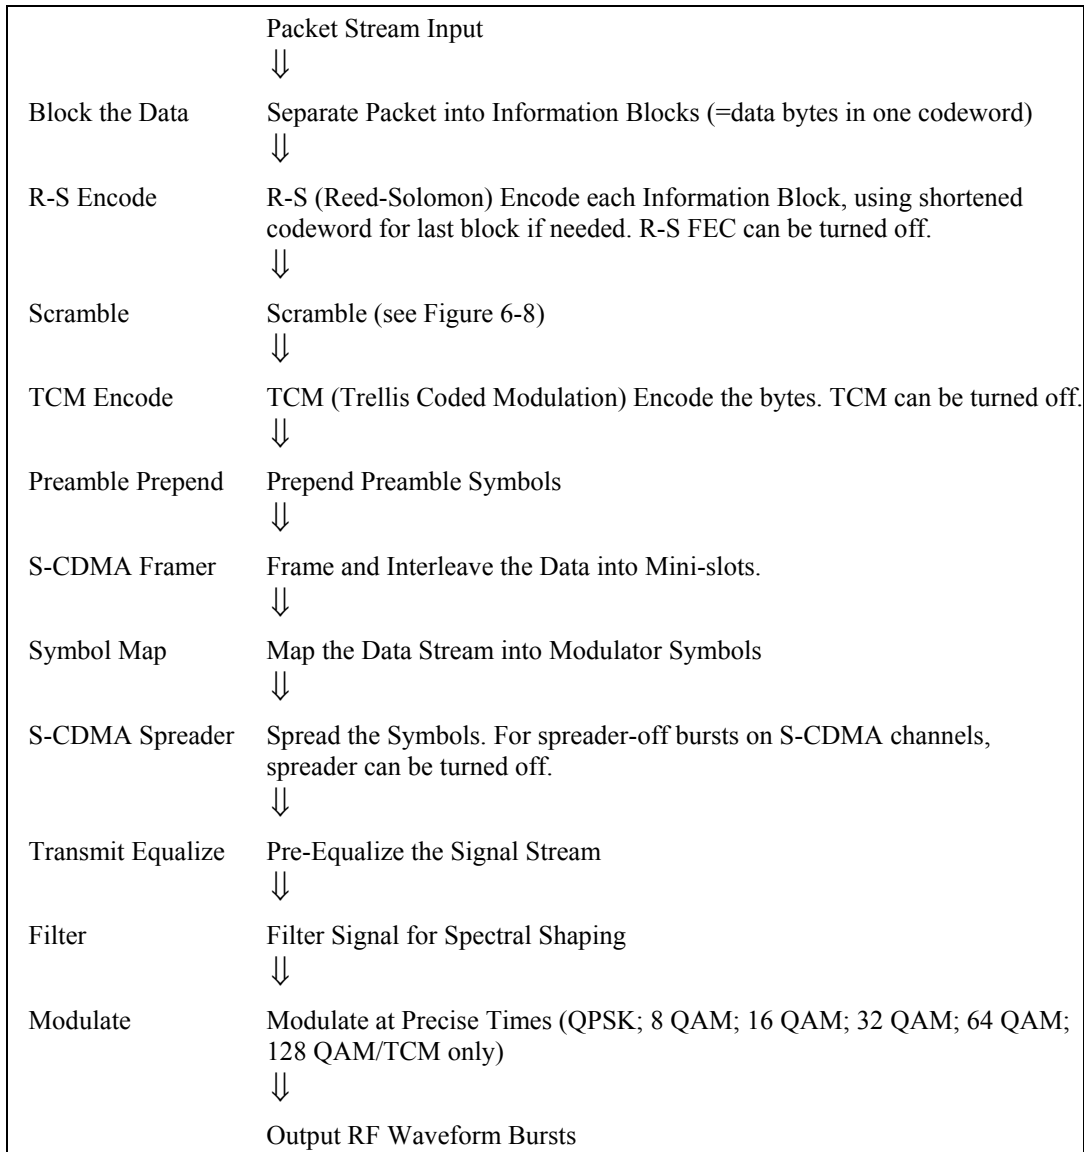


Figure 6-3 - S-CDMA Upstream Transmission Processing

6.2.3 Modulation Formats

The upstream modulator **MUST** provide QPSK and 16 QAM differential encoded modulations for TDMA.

The upstream modulator **MUST** provide QPSK, 8 QAM, 16 QAM, 32 QAM, and 64 QAM modulations for TDMA and S-CDMA channels.

The upstream modulator **MUST** provide QPSK, 8 QAM, 16 QAM, 32 QAM, 64 QAM, and 128 QAM TCM encoded modulations for S-CDMA channels.

The upstream demodulator **MAY** support QPSK, and 16 QAM differential modulation for TDMA.

The upstream demodulator **MUST** support QPSK, 16 QAM, and 64 QAM modulations for TDMA and S-CDMA channels.

The upstream demodulator **MAY** support 8 QAM and 32 QAM modulation for TDMA and S-CDMA channels.

The upstream demodulator **MAY** support QPSK, 8 QAM, 16 QAM, 32 QAM, 64 QAM, and 128 QAM TCM encoded modulations for S-CDMA channels.

6.2.4 R-S Encode

6.2.4.1 R-S Encode Modes

The upstream modulator **MUST** be able to provide the following selections: Reed-Solomon codes over GF(256) with T = 1 to 16 or no R-S coding.

The following Reed-Solomon generator polynomial **MUST** be supported:

$$g(x) = (x+\alpha^0) (x+\alpha^1) \dots (x+\alpha^{2T-1}) \text{ where the primitive element alpha is 0x02 hex}$$

The following Reed-Solomon primitive polynomial **MUST** be supported:

$$p(x) = x^8 + x^4 + x^3 + x^2 + 1$$

The upstream modulator **MUST** provide code words from a minimum size of 18 bytes (16 information bytes [k] plus two parity bytes for T = 1 error correction) to a maximum size of 255 bytes (k-bytes plus parity-bytes). The minimum uncoded word size **MUST** be one byte.

In Shortened Last Codeword mode, the CM **MUST** provide the last codeword of a burst shortened from the assigned length of k data bytes per codeword as described in Sections 6.2.5.1.3 and 6.2.6.

The value of T **MUST** be configured in response to the Upstream Channel Descriptor from the CMTS.

6.2.4.2 R-S Bit-to-Symbol Ordering

The input to the Reed-Solomon Encoder is logically a serial bit stream from the MAC layer of the CM, and the first bit of the stream MUST be mapped into the MSB of the first Reed-Solomon symbol into the encoder. The MSB of the first symbol out of the encoder MUST be mapped into the first bit of the serial bit stream fed to the Scrambler.


Note: The MAC byte-to-serial upstream convention calls for the byte LSB to be mapped into the first bit of the serial bit stream.

6.2.5 Upstream R-S Frame Structure for DOCSIS 3.0 Multiple Transmit Channel mode Enabled

This section applies to CMs operating in DOCSIS 3.0 Multiple Transmit Channel mode enabled in the upstream direction.

Figure 6-4 shows two examples of the R-S frame structure: one where the packet length equals the number of information bytes in a codeword, and another where the packet length is longer than the number of information bytes in one codeword, but less than in two code words. Example 1 illustrates the fixed codeword-length mode, and Example 2 illustrates the shortened last codeword mode. These modes are defined in Section 6.2.5.1.

Example 1. Packet length = number of information bytes in codeword = k



Example 2. Packet length = k + remaining information bytes in 2nd codeword = k + k' ≤ k + k'' ≤ 2k

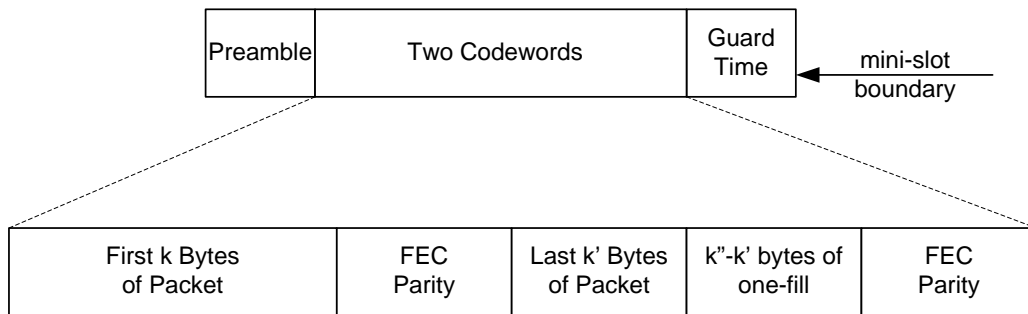


Figure 6-4 - Example Frame Structures with Flexible Burst Length Mode, DOCSIS 3.0 Operation

6.2.5.1 R-S Codeword Length

When R-S FEC is enabled, the CM operates in either fix-length codeword mode or in shortened-last codeword mode. The minimum number of information bytes in a codeword in either mode is 16. Shortened-last codeword mode only provides a benefit when the number of bytes in a codeword is greater than the minimum of 16 bytes.

The intent of the following sections is to define rules and conventions such that the CMTS PHY knows what to expect regarding the R-S FEC framing in both fixed codeword length and shortened last codeword modes. Shortened last codeword mode **MUST NOT** be used for Initial Maintenance (broadcast or unicast).

6.2.5.1.1 Burst Size

For an allocation of mini-slots (in both contention and non-contention regions), the requirements of Sections 6.2.5.1.2 and 6.2.5.1.3 apply to a burst transmitted in that allocation. Regardless of the size of the allocation, the size of the burst **MUST** be as specified in Table 6-1 below.

Table 6-1 - Burst Size

IUC	Burst Size
1, 3	Minimum number of mini-slots required for message transmission including burst overhead. Burst overhead includes pre-amble, R-S parity bytes, TCM return-to-zero bits, and guard time if applicable.
2	Number of mini-slots specified in the Well-Known Multicast SID.
4-6, 9-11	Number of mini-slots allocated.

6.2.5.1.2 Fixed Codeword Length

With the fixed-length code words, after all the data are encoded, one-fill **MUST** occur in this codeword if necessary to reach the assigned k data bytes per codeword. Additionally, one-fill **MUST** continue up to the point when no additional fixed-length code words can be inserted before the end of the burst specified in Table 6-1 above, accounting for preamble, FEC parity, return-to-zero bits, and guard-time symbols (if any).

6.2.5.1.3 Shortened Last Codeword

As shown in Figure 6-4, let k' = the number of information bytes that remain after partitioning the information bytes of the burst into full-length (k burst data bytes) code words. The value of k' is less than k . Given operation in a shortened last codeword mode, let k'' = the number of burst data bytes plus one-fill bytes in the shortened last codeword. In shortened codeword mode, the CM **MUST** encode the data bytes of the burst (including MAC Header) using the assigned codeword size (k information bytes per codeword) until 1) all the data are encoded, or 2) a remainder of data bytes is left over which is less than k . Shortened last code words **MUST NOT** have less than 16 information bytes, and this is to be considered when CMs make requests of mini-slots. In shortened last codeword mode, the CM **MUST** one-fill data if necessary up to the Burst Size specified in Table 6-1 above accounting for preamble, FEC parity, return-to-zero bits and guard-time symbols (if any). Therefore, in many cases, only $k'' - k'$ one-fill bytes are necessary with $16 \leq k'' \leq k$ and $k' \leq k''$.

Generally, the CM **MUST** one-fill data until the point when no additional fixed-length code words can be inserted before the end of the burst specified in Table 6-1 above, accounting for preamble, FEC parity, return-to-zero bits, and guard-time symbols (if any). Then, if possible, a shortened last codeword of one-fill **MUST** be inserted to fit into the last mini-slot.

If, after one-fill of additional code words with k information bytes, there are less than 16 bytes remaining before the end of the burst specified in Table 6-1 above, accounting for preamble, FEC parity, return-to-zero bits and guard-time symbols (if any), then the CM **MUST NOT** create this last shortened codeword.

6.2.5.2 R-S FEC Disabled

When $T = 0$ (no FEC parity bytes), the CM MUST one-fill in full bytes to the end of the burst specified in Section 6.2.5.1.1 above, accounting for preamble, return-to-zero bits, and guard-time symbols (if any).

6.2.6 Upstream R-S Frame Structure for DOCSIS 3.0 Multiple Transmit Channel mode Not Enabled

When Multiple Transmit Channel mode is not enabled, all requirements in Section 6.2.5 for R-S frame structure in Multiple Transmit Channel Operation apply, except that zero-fill MUST be used instead of one-fill in Sections 6.2.5.1.2, 6.2.5.1.3, 6.2.5.2, and in Figure 6-4.

6.2.7 TDMA Byte Interleaver

R-S codeword interleaving in a byte (R-S symbol) format MUST be performed after R-S encoding on a TDMA channel. The byte interleaver changes the order of the bytes at the R-S encoder output, i.e., it performs an operation of byte permutation. At the receiver side, the original order of bytes is restored prior to the R-S decoding. Therefore, if some consecutive bytes were corrupted by burst noise, they are spread between various R-S code words, averaging the number of erroneous bytes in each codeword. The interleaver is a block interleaver type, i.e., the permutation is achieved by filling a table row-wise (one row per R-S codeword), and reading it column-wise. The total memory size allocated for the table is 2048 bytes.

The byte interleaver is disabled when the R-S encoder is turned off ($T = 0$).

6.2.7.1 Byte Interleaver Parameters

The interleaver operating parameters described in Table 6-2 determine the operation of the interleaver for every burst.

Table 6-2 - Interleaver Operating Parameters

Parameter	Definition	Allowed Values
N_r	Interleaver Width (R-S Codeword Length, $k+2*T$)	18 to 255
I_r	Interleaver Depth	0 - Dynamic Mode 1 - No Interleaving 2 to $\text{floor}(2048/N_r)$ - Fixed Mode
B_r	Interleaver Block Size	$2*N_r$ to 2048
N_f	Packet Size (in bytes, including FEC)	≥ 18 bytes

The CMTS and CM MUST use the interleaver parameters within the allowed values in Table 6-2 with the following additional restrictions:

1. N_r and I_r MUST be chosen such that $N_r I_r \leq 2048$ (in other words, for a given N_r , the maximal value of I_r is $I_{r,\max} = \text{floor}(2048/N_r)$).
2. N_r MUST be identical to the R-S codeword length (i.e., $k+2T$).
3. B_r is effective only when $I_r=0$. This mode is called dynamic mode.

4. When $I_r=1$, interleaving is disabled.

N_r , I_r , and B_r are specified in the burst profile, and N_f is implied in the MAP message.

6.2.7.2 Interleaver Operating Modes

The interleaver MUST support both an operating mode in which the block size is fixed, as well as a dynamic mode in which the interleaver depth is determined based on the burst size.

6.2.7.2.1 Fixed Mode

The R-S encoded data bytes of the packet are first divided into interleaver blocks of $N_r I_r$ bytes (i.e., blocks of I_r R-S code words each). The size of the last interleaver block may be smaller when the packet length is not an integer multiple of $N_r I_r$. Each interleaver block is interleaved separately.

Each interleaver block is filled into a table with I_r rows and N_r columns. The data is written row by row (from left to right). Therefore, each row corresponds to one R-S codeword. The bytes are read column by column (from top to bottom). The interleaver operation is demonstrated in Figure 6-5.

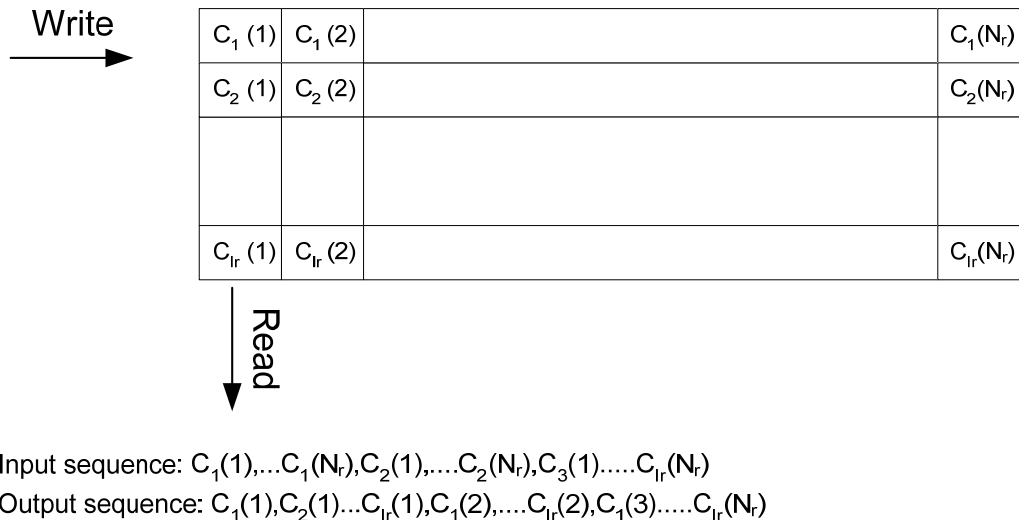
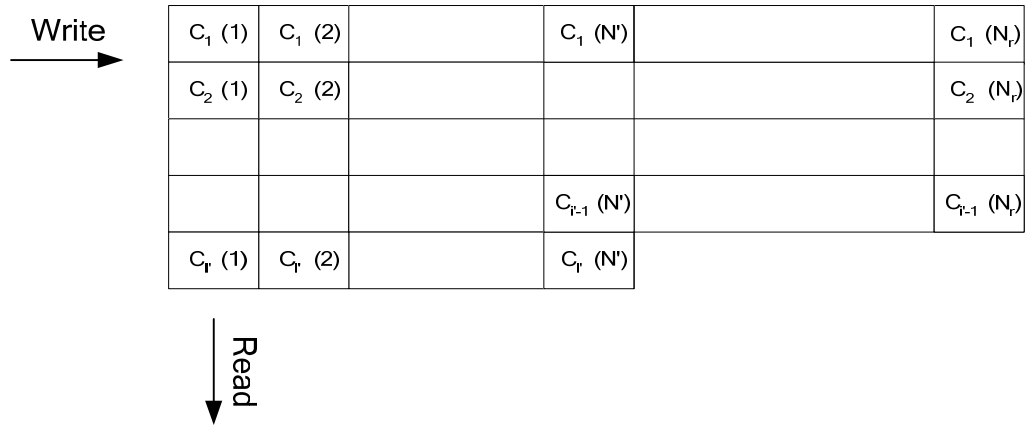


Figure 6-5 - Byte Interleaver Operation

The last interleaver block might have fewer rows than I_r . If the shortened last codeword mode is applied, then the last row might have fewer elements than N_r . In these cases, the interleaver table is read column by column, skipping the empty elements of the table. The interleaver operation for the last interleaver block is demonstrated in Figure 6-6.



Input sequence: $C_1(1), \dots, C_1(N_r), C_2(1), \dots, C_2(N_r), C_3(1), \dots, C_r(1), \dots, C_r(N')$
 Output sequence:
 $C_1(1), C_2(1), \dots, C_r(1), C_1(2), \dots, C_r(2), \dots, C_1(N'), \dots, C_r(N'), C_1(N'+1), \dots, C_{r-1}(N'+1), C_1(N'+2), \dots, C_{r-1}(N'+2), \dots, C_1(N_r), \dots, C_{r-1}(N_r)$

Figure 6-6 - Interleaver Operation for Last Interleaver Block (with Shortened Last Codeword)

6.2.7.2.2 Dynamic Mode

In the fixed mode, the interleaving depth of the last interleaving block of a packet (I' in Figure 6-6) may be as small as one, resulting in low burst noise robustness for this block. In the dynamic mode, the depths of the interleaver blocks are chosen such that all blocks have approximately the same depth to achieve nearly optimal burst noise robustness (for the given block size).

The R-S encoded data bytes of the packet are first divided into N_s^0 interleaver blocks. The size of the i^{th} interleaver block is $N_r \cdot I_r^{(i)}$ bytes (i.e., a block of $I_r^{(i)}$ R-S code words). The size of the last interleaver block may be smaller in the shortened last codeword mode. Each interleaver block is interleaved separately (see the equations for N_s^0 and $I_r^{(i)}$ in Section 6.2.7.2.2.1).

The i^{th} interleaver block is filled into a table with $I_r^{(i)}$ rows and N_r columns. The data is written row-wise (from left to right). Therefore, each row corresponds to one R-S codeword. The bytes are read column-wise (from top to bottom). The interleaver operation is demonstrated in Figure 6-5 (except that there are $I_r^{(i)}$ rows instead of I_r).

If the shortened last codeword mode is applied, then the last row might have fewer elements than N_r . In this case, the interleaver table is read column by column, skipping the empty elements of the table. The interleaver operation for the last interleaver block is demonstrated in Figure 6-6 (except that there are $I_r^{(N_s^0)}$ rows instead of I').

6.2.7.2.2.1 Dynamic Mode Calculations

Dynamic mode calculations are illustrated below.

N_s^0 and $I_r^{(i)}$ are determined by the following equations:

Total number of interleaver rows: $I_{tot}^0 = \text{ceil}(N_f / N_r)$.

Maximal number of rows per segment: $I_{r,max} = \text{floor}(B_r / N_r)$.

Number of segments: $N_s^0 = \text{ceil}(I_{tot}^0 / I_{r,max})$

Interleaver depth of first block: $I_r^1 = \text{floor}(I_{tot}^0 / N_s^0)$

No. of blocks with depth of I_r^1 : $M = N_s^0 \cdot (I_r^1 + 1) - I_{tot}^0$

Then for segment i , $I_r^{(i)}$ is calculated as follows ($i = 1 \dots N_s^0$):

$$I_r^{(i)} = \begin{cases} I_r^1, & i = 1, \dots, M \\ I_r^1 + 1, & i = M + 1, \dots, N_s^0 \end{cases}$$

Figure 6-7 - T Mode Calculations

6.2.8 Scrambler (Randomizer)

The upstream modulator MUST implement a scrambler (shown in Figure 6-8) where the 15-bit seed value MUST be arbitrarily programmable.

At the beginning of each burst, the register is cleared and the seed value is loaded. The seed value MUST be used to calculate the scrambler bit which is combined in an XOR with the first bit of data of each burst (which is the MSB of the first symbol following the last symbol of the preamble).

The scrambler seed value MUST be configured in response to the Upstream Channel Descriptor from the CMTS.

The polynomial MUST be $x^{15} + x^{14} + 1$.

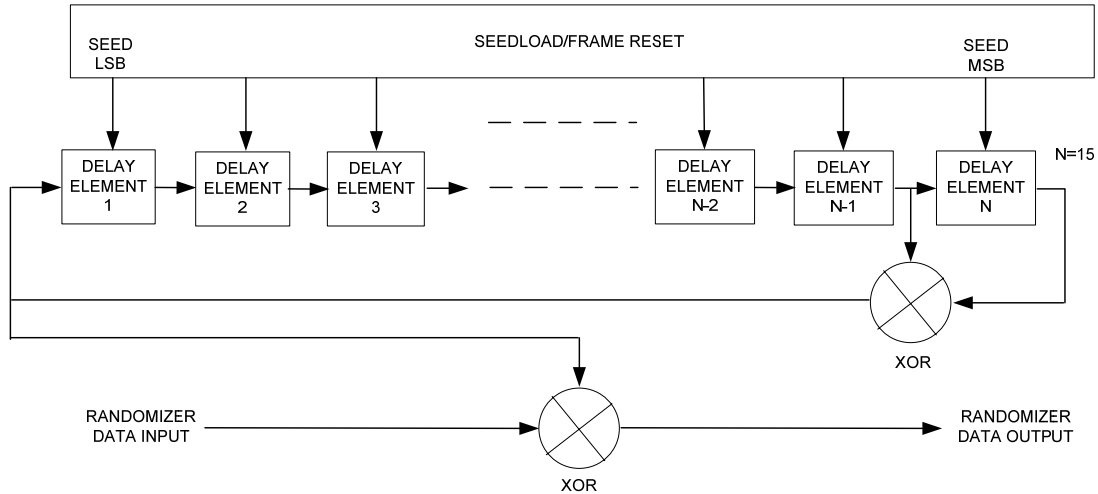


Figure 6-8 - Scrambler Structure

6.2.9 TCM Encoder

R-S symbol interleaving is commonly included between the TCM and R-S blocks to preserve coding gain in the presence of bursts of errors produced at the output of the TCM decoder. This interleaver was not included in the original baseline S-CDMA proposal to reduce memory requirements at the expense of coding gain.

In S-CDMA mode, the CM MUST support trellis coded modulation for transmission of $m = 1, 2, 3, 4, 5,$ and 6 bits per symbol with QPSK0, 8 QAM, 16 QAM, 32 QAM, 64 QAM, and 128 QAM constellations, respectively. Support of TCM in the CMTS is optional.

Figure 6-9 shows the employed 8-state TCM encoder. The encoding operation causes a mapping of m input bits into $m+1$ output bits for input into the symbol mapping block. The systematic convolutional encoder adds the coded bit $x^1 = s^0$ to the input bits i^m, i^3, i^2, i^1 . For $m = 1$, only input bit i^1 is used ($i^2 = 0$), and encoding is reduced to rate-1/2 coding.

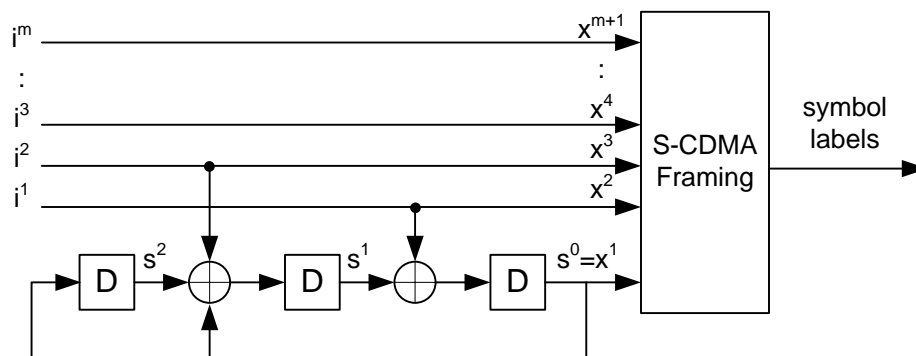


Figure 6-9 - Convolutional encoder

The initial state of the TCM encoder MUST be the zero state. The zero state MUST be reached again with the last encoded symbol.

To return to the zero state from all possible Trellis paths, if $m = 1$ (QPSK) three tail symbols ($n_t = 3$) MUST be generated with input bit i^1 set to $i^1 = s^1$. By inspection of Figure 6-9, after three symbols the state bits s^2, s^1 , and $s^0 = x^1$ will be zero. Tail symbols are extra symbols, which carry no information.

If $m = 2$, to return to the zero state from all possible Trellis paths, two tail symbols ($n_t=2$) MUST be generated. The input bits $i^2 i^1$ MUST be set such that the zero state is reached after two symbols. If the first symbol is set to $i^2 = 0, i^1 = s^1$, and the second (final) symbol to $i^2 = s^2, i^1 = s^1$ after these two symbols the state bits s^2, s^1 , and $s^0 = x^1$ will be zero.

If $m \geq 3$, the uncoded bits i^m, \dots, i^3 MUST be used for information encoding, when this is possible. Otherwise, uncoded bits MUST be set to zero. The number of tail symbols carrying no information depends on the ending conditions and can vary between zero and two ($0 \leq n_t \leq 2$).

6.2.9.1 Byte to TCM Symbol Mapping

The mapping of bytes to TCM symbols is done such that each byte is mapped entirely to the uncoded bits i^m, \dots, i^3 , or entirely to the convolutional encoder input bits $i^2 i^1$. The decision is made sequentially for each byte using the rule that the byte assignment should lead to the shortest packet of symbols including tail symbols, if the current byte were the last byte to be encoded. This rule results in the repetitive patterns of byte assignments to label bits shown in Figure 6-10 for $m=1$ to 6. In the figure bit i^m is at the top and bit i^1 is at the bottom.

The MSB (i^m) MUST be the first bit in the serial data fed into the uncoded input bits (i.e., i^m to i^3). The MSB (i^2) MUST be the first bit in the serial data fed into the coded input bits.

Figure 6-11 illustrates the byte assignments for Trellis-coded 64 QAM modulation by two examples. Notice that bytes are assigned in a repetitive pattern of five bytes. In the first example, N_f is divisible by five. In this case two tail symbols are appended. In the second example, N_f is not divisible by five and no tail symbols are required. The bits needed for returning to the zero state are available in symbols still carrying information.

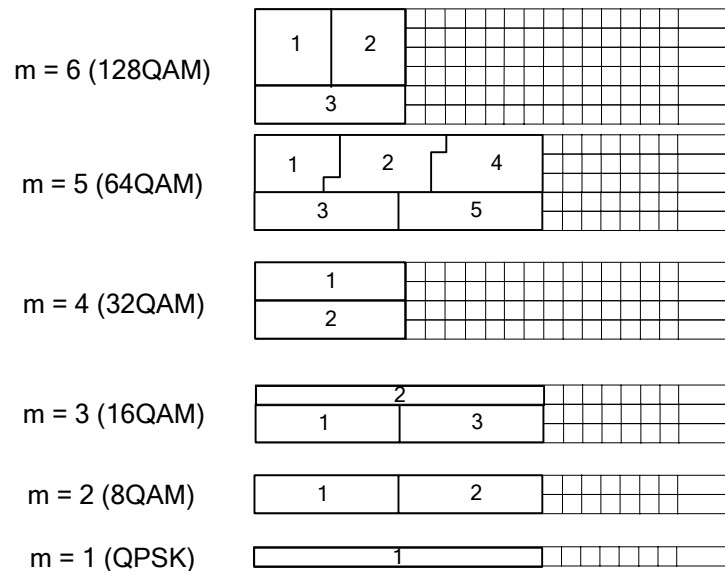


Figure 6-10 - Repetitive Patterns of Byte Mapping to Symbol Map Bits for TCM

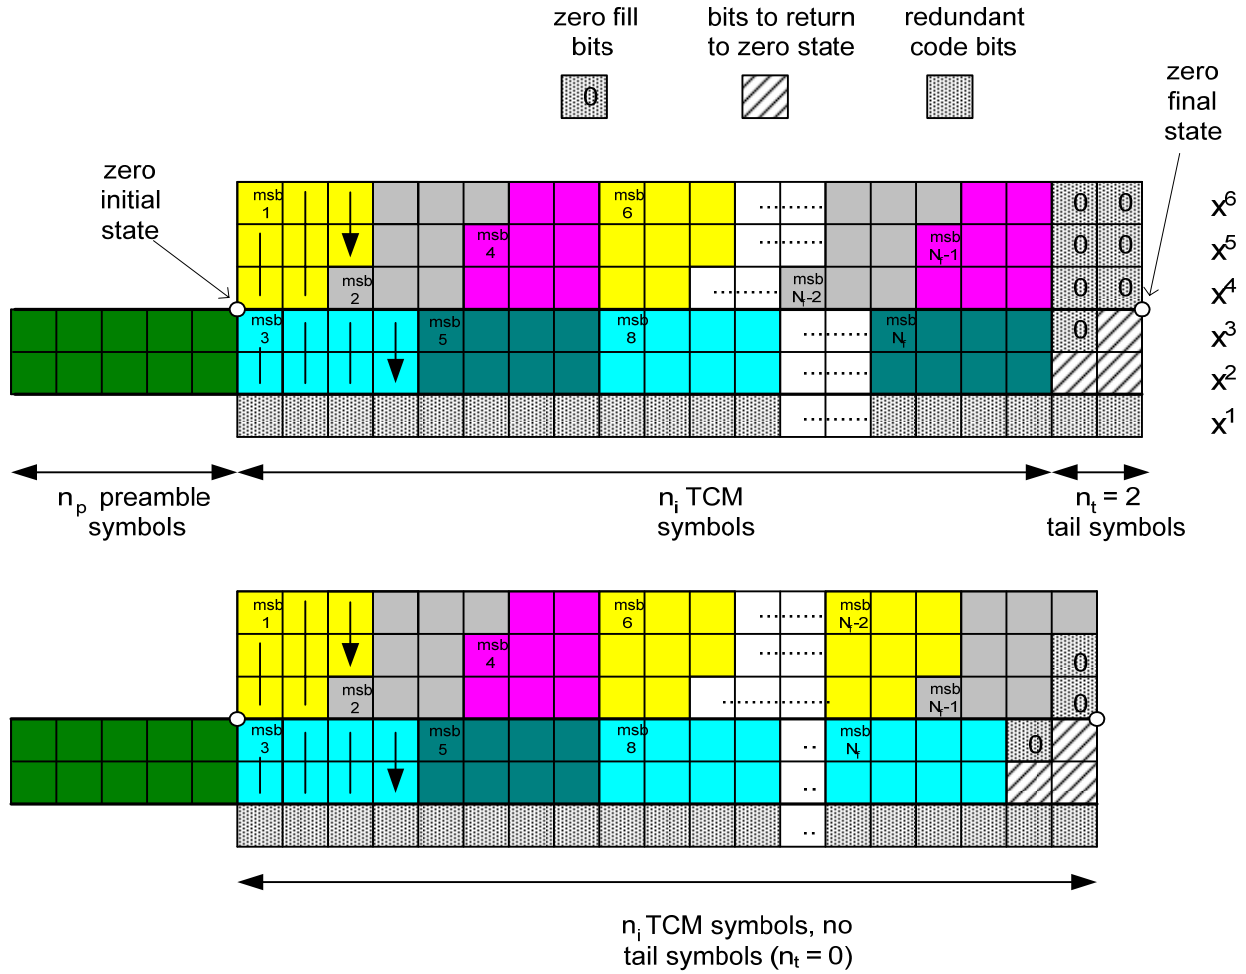


Figure 6-11 - Example Byte to Bit Assignment for 64 QAM

The CM MUST place the return-to-zero bits right after the last TCM-coded data sub-symbol, that is, the last coded sub-symbol corresponding to the parity bytes of the last shortened or fixed R-S codeword including any zero-filled (for operation not in MTC Mode) or one-filled (for operation in MTC Mode) R-S code words in the grant. The rest of the TCM-coded bits MUST be filled with zeros.

Figure 6-12 illustrates the placing of return-to-zero bits for 64 QAM when the last transmitted byte is #1. The first two pairs of x^2 and x^3 are the return to zero bits, and the last empty coded pair is zero-filled or one-filled.

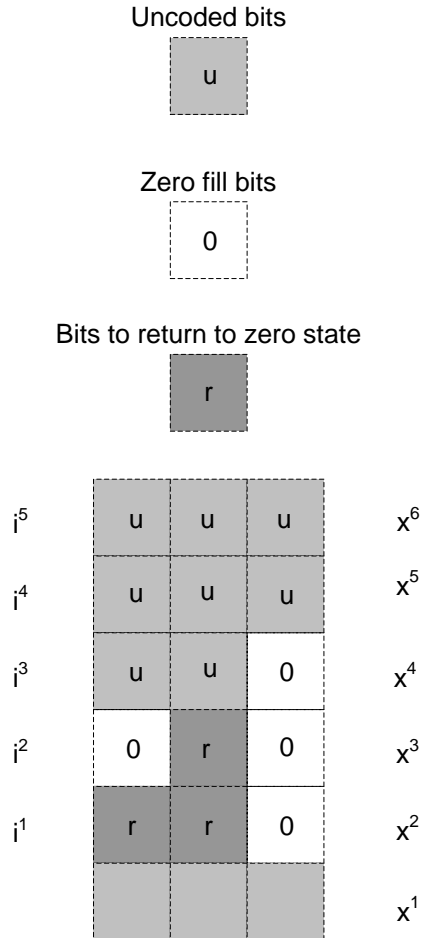


Figure 6-12 - Example of return to zero bits followed by "0"

6.2.10 Preamble Prepend

The upstream PMD sublayer **MUST** support a variable-length preamble field that is prepended to the data after it has been randomized, Reed-Solomon encoded, and TCM encoded.

The first bit of the Preamble Pattern is the first bit into the symbol mapper (see Section 6.2.14). The first bit of the Preamble Pattern is designated by the Preamble Value Offset. The preamble is interleaved by the framer in S-CDMA mode.

The preamble sequence **MUST** be programmable. For DOCSIS 2.0/3.0 bursts (bursts encoded using a Type 5 burst descriptor), the preamble **MUST** use the QPSK0 or QPSK1 constellation (per Figure 6-20 and Figure 6-21) with preamble length 0, 2, 4, 6, ..., or 1536 bits (maximum 768 QPSK symbols). For DOCSIS 1.x compatible bursts (Type 4 burst descriptor) that use QPSK modulation, the preamble and data **MUST** use the QPSK0 constellation with preamble length 0, 2, 4, 6, ..., or 1024 bits (maximum 512 QPSK symbols). For DOCSIS 1.x compatible bursts (Type 4 burst descriptor) that use 16 QAM modulation, the preamble and data **MUST** use the 16 QAM constellation with preamble length 0, 4, 8, 12, ..., or 1024 bits (maximum 256 16 QAM symbols).

The preamble length and value **MUST** be configured in response to the Upstream Channel Descriptor message transmitted by the CMTS.

6.2.11 Modulation Rates

6.2.11.1 DOCSIS 3.0 Modulation Rates

In TDMA and S-CDMA modes, the CM upstream modulator **MUST** provide all modulations at 1280, 2560 and 5120 kHz.

In TDMA and S-CDMA modes, the CMTS upstream demodulator **MUST** be able to support demodulation at 1280, 2560 and 5120 kHz.

This variety of modulation rates, and flexibility in setting upstream carrier frequencies, permits operators to position carriers in gaps in the pattern of narrowband ingress.

The modulation rate for each upstream channel is defined in an Upstream Channel Descriptor (UCD) MAC message. All CMs using that upstream channel **MUST** use the defined modulation rate for upstream transmissions.

6.2.11.2 Backward Compatibility Modulation Rates

Operating with a DOCSIS 1.x or DOCSIS 2.0 CMTS, or with a CMTS operating in one of these modes, in TDMA and S-CDMA modes, the CM upstream modulator **MUST** provide all modulations at 1280, 2560 and 5120 kHz.

In addition, with such pre-3.0-DOCSIS CMTS operation, for TDMA mode, the CM upstream modulator **MAY** provide all modulations at 160, 320 and 640 kHz.

In pre-3.0-DOCSIS CMTS operation, for both TDMA and S-CDMA modes, the CMTS upstream demodulator **MUST** be able to support demodulation at 1280, 2560, and 5120 kHz. In pre-3.0-DOCSIS CMTS operation, for TDMA mode, the CMTS upstream demodulator **MAY** support demodulation at 160, 320, and 640 kHz.

6.2.12 S-CDMA Framer and Interleaver

6.2.12.1 S-CDMA Framing Considerations

The S-CDMA mode of the PHY layer accepts data presented to it for transmission from the MAC layer. This data is presented as bursts of n mini-slots. These bursts are mapped within the PHY layer to a combination of spreading codes and time slots, in order to exploit the multi-dimensional spreading of information by the S-CDMA mode.

There are various adjustable parameters in the upstream channel parameters and upstream burst attributes that allow controlling the mini-slot to physical layer mapping, as well as tuning the channel to accommodate a variety of channel conditions, noise characteristics, capacities, reliability levels, and latency requirements.

When operating in S-CDMA mode, data is transmitted in two dimensions: codes and time. For this reason, data to be transmitted is grouped into two-dimensional rectangular frames prior to transmission.

At the physical layer, data is sent over an array of up to 128 spreading codes. There is a programmable number of *spreading intervals* per frame, as shown in Figure 6-13 below. A *spreading interval* is the time required to transmit one symbol per code across all 128 codes in S-CDMA mode. Note that the specific codes which are used and the details of the spreading operation are described in detail in Section 6.2.15.

A burst from a particular CM may be transmitted on two or more codes in one or more frames. A frame may contain bursts transmitted simultaneously from multiple CMs (each on a separate subset of the codes) as defined by the MAP message.

6.2.12.2 Mini-slot Numbering

In normal operation, the MAC will request the PHY to transmit a burst of length n mini-slots, starting at mini-slot m , as defined by the MAP. All CMs and the CMTS MUST have a common protocol of how mini-slots are numbered, and how they are mapped onto the physical layer framing structure. This common protocol is obtained from information in the SYNC and Upstream Channel Descriptor (UCD) messages.

Mini-slots are mapped onto frames starting at the first active code, are numbered sequentially through the remainder of the frame, and then wrap to the next sequential frame (see Section 6.2.12.2.2 for mini-slot numbering examples). Mini-slots are mapped onto a group of consecutive codes.

The CMTS and the CMs require a common protocol for mini-slot numbering. For operation on a TDMA channel, this is achieved solely through recovery of the timestamp. Since the time duration of an S-CDMA frame is not necessarily a power-of-2 multiple of the 10.24 MHz reference, the timestamp rollover (at 2^{32} counts) is not necessarily at an S-CDMA frame boundary. Therefore, an additional synchronization step is required.

The CMTS is required to identify frame boundaries relative to the timestamp counter on a periodic basis. This is called the *timestamp snapshot* and has to be sent in the UCD for each upstream S-CDMA channel.

The CMTS is required to maintain a frame counter and a mini-slot counter. The CMTS has to sample these values along with the timestamp, on a frame boundary, as shown in Figure 6-13 below. The CMTS is required to obtain a new sample prior to sending each UCD message.

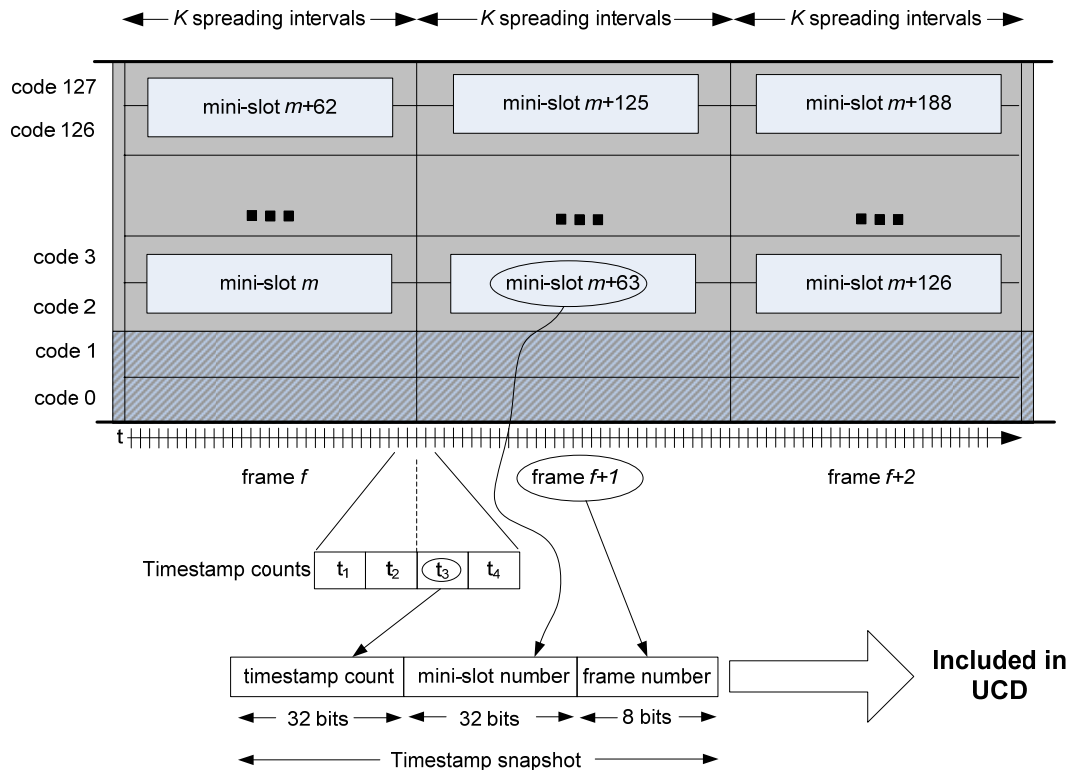


Figure 6-13 - Timestamp Snapshot

Each CM MUST maintain a timestamp counter, mini-slot counter, and frame counter functionally identical to the CMTS.

From the UCD message, the CM receives the CMTS timestamp snapshot and parameters from which it can calculate the number of time counts per S-CDMA frame. Using modulo arithmetic, the CM can then calculate accurate values for timestamp, mini-slot, and frame counters at any point into the future.

The CM can then update its local mini-slot and frame counters at an appropriate timestamp counter value. At this point, the CM representation of mini-slots and frames are aligned with those in the CMTS.

The CMTS and CM MUST implement a 32-bit timestamp counter, a 32-bit mini-slot counter, and an 8-bit frame counter, as follows:

- The mini-slot counter MUST contain the value of the first mini-slot of the frame when it is sampled. It MAY be incremented by the number of mini-slots per frame, once per frame interval. The mini-slot counter will use all 32 bits and mini-slot numbers will, therefore, range from 0 to $2^{32}-1$.
- The only specified function for the frame counter is to reset the code hopping sequence at the frame 0 (modulo-256) boundary, as defined in Section 6.2.15.1.

The frame structure above relates to the entire upstream and not necessarily to the transmission from a single CM. The codes are resources which are allocated to CMs over each S-CDMA frame. The assignment of codes to CMs is performed by the framer as it assigns a burst of symbols a particular order in the two-dimensional matrix of codes and time. This symbol sequencing is described in detail in Section 6.2.13.

6.2.12.2.1 Mini-slot Numbering Parameters in UCD

There are three parameters specified in the UCD that define mini-slot mapping: *spreading intervals per frame*, *codes per mini-slot*, and *number of active codes*.

Spreading intervals per frame

The number of *spreading intervals per frame*, K , (along with the signaling rate), $1/T_s$, define the time duration of an S-CDMA frame, T_{fr} :

$$T_{fr} = K * 128 * T_s$$

Note that the code length in the above equation is always 128, regardless of how many codes are currently active.

The valid range of the *spreading intervals per frame* parameter is 1 to 32.

Codes per mini-slot

In conjunction with the spreading intervals per frame parameter, the codes per mini-slot (C_{ms}) parameter defines the total number of symbols per mini-slot and, therefore, the mini-slot capacity. The mini-slot capacity, S_{ms} , is given in symbols by the following expression:

$$S_{ms} = K * C_{ms}$$

The lower limit on mini-slot capacity is 16 symbols. However, the mini-slot also has to be large enough to allow the transmission of the largest-sized data PDU (including physical layer overhead) in 255 mini-slots. The maximum mini-slot capacity in symbols is given by the product of the maximum number of spreading intervals per frame and the maximum number of codes per mini-slot ($32 \times 32 = 1024$ symbols). The valid range of the *codes per mini-slot* parameter is 2 to 32.

Selectable Active Codes (SAC)

The number of active codes parameter N_a allows the number of codes used to carry data to be less than or equal to 128. When $N_a < 128$, the active codes may be chosen using the following two modes.

Selectable Active Codes Mode 1: The low-numbered codes starting with code 0 are not used, as shown below in Figure 6-15 (an example of 126 active codes and code hopping off).

Selectable Active Codes Mode 2: The active codes are selectable via a 128-bit string. The first element in the string corresponds to code 0 (the all-ones code). A "1" element in the string indicates an active code, and a "0" indicates an unused code.

There are several reasons why it may be desirable to reduce the number of active codes:

- Code 0 does not have the same spreading properties as the other codes and, therefore, under certain colored noise conditions, will degrade performance.
- In extremely noisy plant conditions, a reduction in the number of active codes (along with the corresponding increase in power per code for the remaining codes) can allow reliable operation at reduced capacities. Reduction in active codes from 128 to 64 results in a 3 dB improvement in SNR.
- The number of mini-slots per S-CDMA frame MUST be an integer. Therefore the codes per mini-slot and number of active codes parameters MUST be chosen to result in an integral number of mini-slots per frame.

When $N_a \geq 64$, the S-CDMA frame consists of more than 1 mini-slot, since the number of codes per mini-slot is in the range 2 to 32. This implies that N_a is non-prime. The prime numbers between 64 and 128 are {67, 71, 73, 79, 83, 89, 97, 101, 103, 107, 109, 113, and 127}.

The CM MUST support any non-prime number of active codes in the range $64 \leq N_a \leq 128$. The CM MUST support Selectable Active Codes Mode 1. The CM SHOULD support Selectable Active Codes Mode 2.

The CMTS MUST support 126 and 128 active codes. The CMTS MUST support Selectable Active Codes Mode 1. The CMTS MAY support Selectable Active Codes Mode 2.

6.2.12.2.2 Mini-slot Numbering Examples

A typical mini-slot numbering example with $N_a = 128$ active codes and code hopping off is shown in Figure 6-14. In this example, there are two codes per mini-slot defined. The number of codes per mini-slot is an adjustable parameter (via the UCD) to allow flexibility in determining the effective capacity of each mini-slot.

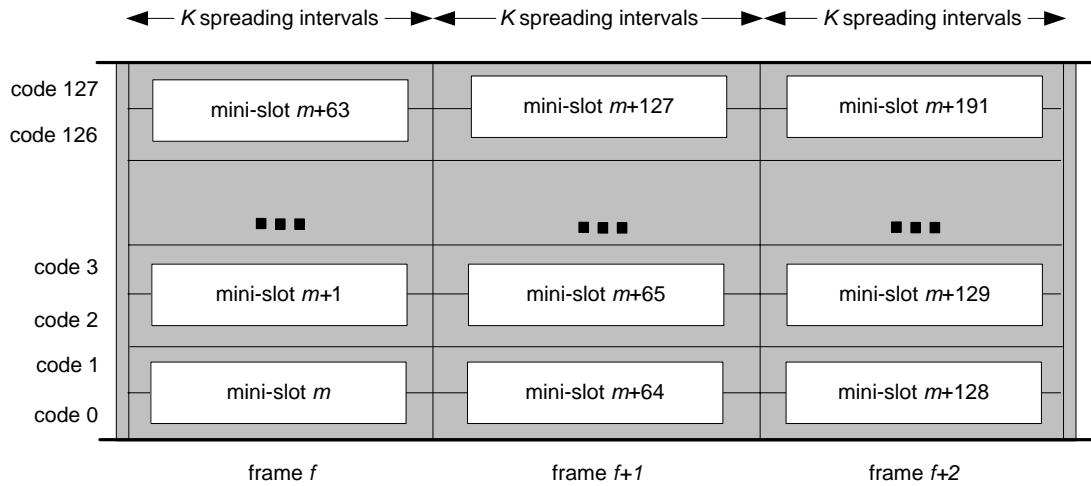


Figure 6-14 - Mini-slot Mapping with Two Codes per mini-slot, 128 Active Codes

A second example, using three codes per mini-slot, with code hopping off, is shown in Figure 6-15. Since it is required that there be an integral number of mini-slots per frame, the number of active codes N_a has been restricted to 126, via either Selective Active Codes Mode 1 or 2. In this example, a tradeoff has been made to increase mapping flexibility at the expense of a small reduction in channel capacity (the ratio 2/128).

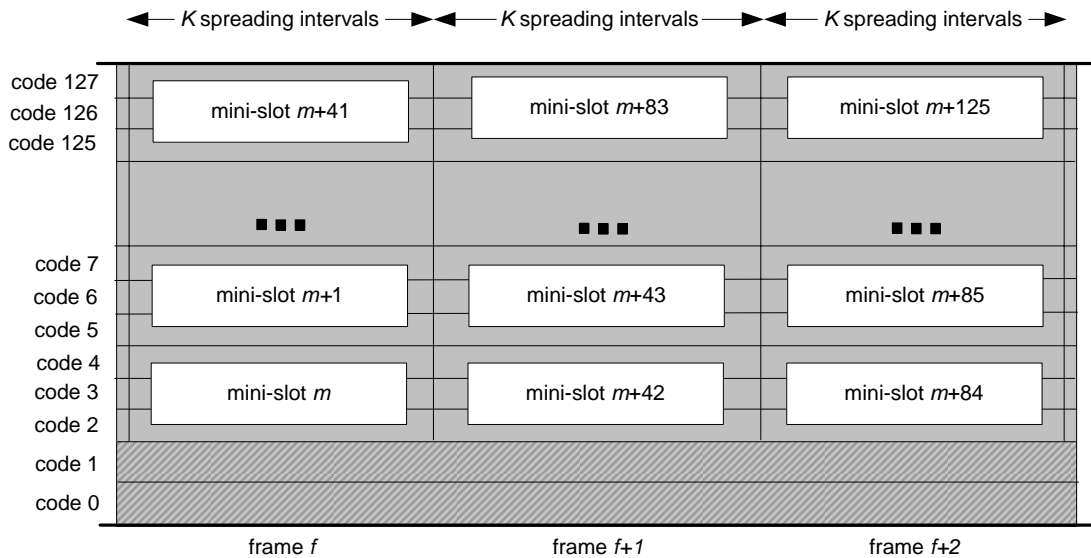


Figure 6-15 - Mini-slot Mapping with Three Codes per mini-slot, 126 Active Codes

Figure 6-16 shows an example of $N_a = 124$ active codes with codes 0, 1, 5, and 125 unused; Selectable Active Codes Mode 2 and code hopping off. The original spreading code matrix has been re-ordered such that the unused codes are moved to the bottom rows. The active codes are arranged in the matrix in increasing order from bottom to top, as are the unused codes. The mini-slots are defined on the active codes only, as shown in Figure 6-16.

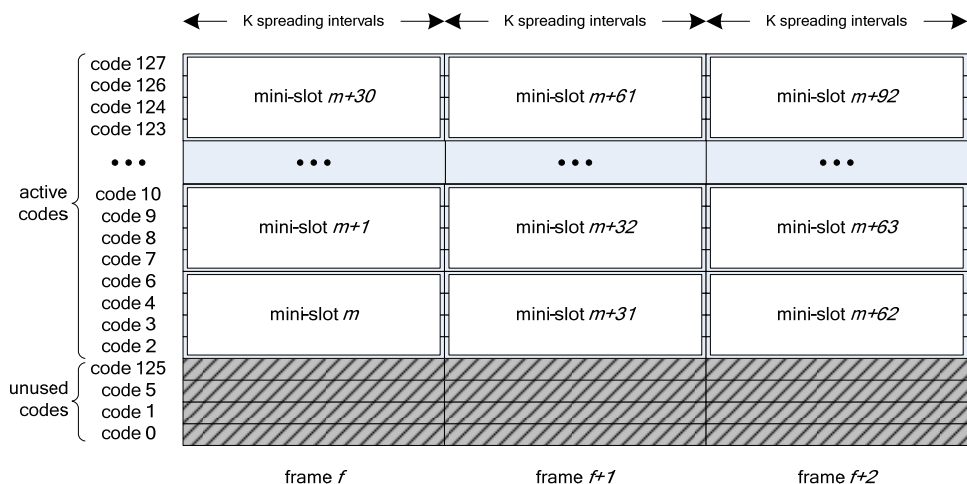


Figure 6-16 - Mini-slot Mapping with Four Codes per Mini-slot, 124 Active Codes, Codes 0, 1, 5, and 125 Unused; Selectable Active Codes Mode 2

There is no implication that physical layer processing is performed on a per-mini-slot basis. As in a TDMA channel, the physical layer is concerned only with the burst start time (mini-slot number), and the burst length.

6.2.12.3 Transmission Time

Ideally, all the mini-slots contained in one S-CDMA frame are received simultaneously. These mini-slots may be transmitted from a single CM or may be transmitted from multiple CMs, as defined by the bandwidth allocation MAP message and the mini-slot mapping configuration settings (from the UCD). Note that a single CM may have more than one allocation active in a single S-CDMA frame.

6.2.12.4 Latency Considerations

S-CDMA frame timing is derived directly from (is phase locked to) the 10.24 MHz CMTS Master Clock. Based on the allowable modulation rates, and the fact that there are 128 modulation intervals (chip periods) in a spreading interval, the S-CDMA frame time MUST always be a multiple of 25 μ s.

The number of spreading intervals per frame and the modulation rate exactly define the S-CDMA frame duration. As a specific example, a burst profile defined with 10 spreading intervals per frame, with a modulation rate of 2.56 MHz, would result in a frame duration of 500 μ s.

The amount of additional upstream latency added by the use of S-CDMA mode is approximately one S-CDMA frame, with the exact value described in Section 6.2.18.

6.2.12.5 Spreader-off Bursts for Maintenance on S-CDMA channel

Spreader-off bursts are defined as bursts on an S-CDMA channel which has attributes that specify the spreader be turned off. For a spreader-off burst, both the S-CDMA framer and S-CDMA spreader are bypassed. The Initial Maintenance burst type MUST be specified (via UCD) to use spreader-off bursts. The Station Maintenance burst type MAY be specified (via UCD) to use spreader-off or spreader-on bursts. The CM MUST support both spreader-on and spreader-off modes for Station Maintenance bursts. All remaining IUC burst types MUST be specified (via UCD) to use spreader-on bursts. The S-CDMA channel will be programmed (via UCD) for C_{ms} codes per mini-slot, p number of active codes, K spreading intervals per S-CDMA frame, and a resultant s mini-slots per frame, where $s = p/C_{ms}$.

Then each S-CDMA frame, where a transmission with the spreader off is to occur, will contain exactly s mini-slots, where each mini-slot consists of $C_{ms} * K$ symbols, as shown in Figure 6-17.

In the case where the number of active codes (p) is less than 128, the frame will still contain exactly s mini-slots, where each mini-slot consists of $C_{ms} * K$ symbols. The first mini-slot of a frame will start with the first symbol of the frame. If a burst spans multiple frames, the burst will start relative to the first frame and continue without interruption into the next frame.

Spreader-off bursts for Station Maintenance Regions (IUC 4) MUST be padded with zero data symbols from the end of the R-S encoded data until the end of the burst as defined by the burst boundaries of Section 6.2.5.1.1. Spreader-off bursts for Initial Maintenance Regions (IUC3) MUST be padded with zero data symbols from the end of the R-S encoded data until the end of the burst as defined by the burst boundaries of Section 6.2.5.1.1. Differential encoding and R-S byte interleaving MUST NOT be used with spreader-off bursts on S-CDMA channels.

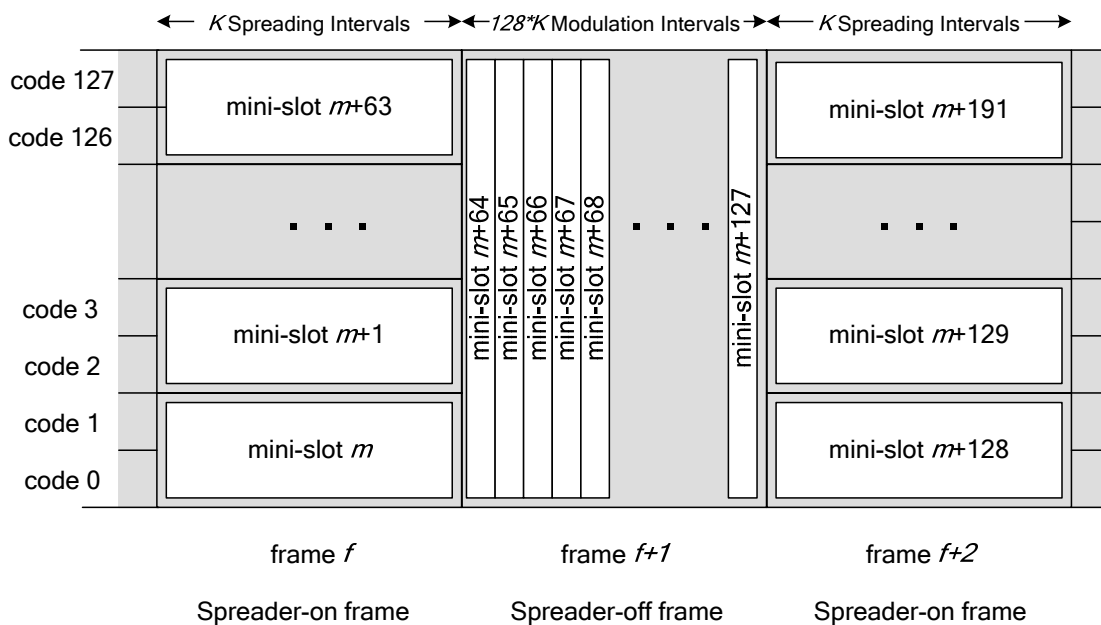


Figure 6-17 - S-CDMA Spreader-on and Spreader-off Intervals

The CMTS scheduler MUST ensure that the spreader-off interval is aligned to the start of an S-CDMA frame, occurs completely within one or more S-CDMA frames, and MUST ensure that no spreader-on bursts are scheduled during these same frames. The CMTS scheduler MUST grant at most one spreader-off burst per CM per frame. It is the responsibility of the CMTS to allocate mini-slots to the NULL SID, as required to prevent interference between bursts (*i.e.*, before and after spreader-off bursts when the CM might not be sufficiently synchronized). Specifically, the CMTS MUST issue a NULL grant (to the NULL SID) of 1 mini-slot immediately before each spreader-off burst, which corresponds to either Station Maintenance, or Unicast Initial Maintenance. The CMTS MUST also issue a NULL grant (to the NULL SID) of 1 mini-slot or guarantee a quiet mini-slot (dead time) immediately after these bursts and before a spreader-on interval starts.

During spreader-off bursts on S-CDMA channels when less than 128 active codes are in use, the spreader-off frame will contain quiet mini-slots (dead time) equal to the number of inactive codes.

6.2.12.6 Limiting the Number of Codes Assigned to a CM

In certain situations, it may be useful for a CMTS to limit the number of codes that a single CM is required to simultaneously transmit. By doing so, the CM can divide its transmit power across a smaller number of codes than it would otherwise, which results in a higher power per code. This can be especially useful when a population of CMs is subject to an unusually high upstream attenuation, such that the CMs are transmitting at the maximum total transmit power. When the value of Maximum Scheduled Codes is set less than the Number of Active Codes, the CMTS MUST ensure that each compliant CM will not, via scheduled grants or multicast IEs with IUC=1, exceed its assigned Maximum Scheduled Codes transmission limit in any S-CDMA frame. To accomplish this, the CMTS must avoid scenarios that would potentially cause the CM to attempt to transmit on more codes than its Maximum Scheduled Code limit would allow. For instance, the CMTS must manage the number of codes assigned to contention IEs with IUC=1 in all frames. In frames where IEs with IUC=1 could not be inserted by a CMTS because of a CM's Maximum Scheduled Codes, the CMTS MAY provide multicast IEs with IUC=2 for contention request opportunities. CMs with Maximum Scheduled Codes enabled MUST be configurable via SNMP to control usage of IEs with IUC=2 [OSSI3.0]. By default, CMs with Maximum Scheduled Codes Enabled MUST NOT use IEs with IUC=2. Maximum Scheduled Codes MUST be equivalent to an integer number of mini-slots.

A CM MUST NOT concatenate packets beyond the size permitted by the S-CDMA Maximum Scheduled Codes, if Maximum Scheduled Codes specified in the RNG-RSP is not 0. This is in order to reduce fragmentation overhead, which can become significant as the number of codes reduces. A CM receiving a Maximum Scheduled Codes value MUST be capable of fragmenting any MAC frame, including frames transmitted prior to completing the registration process. To support 1.0 style configuration files, a CM and CMTS using the Maximum Scheduled Codes value SHOULD support fragmentation in 1.0 mode.

If a UGS flow is requested to provide an Unsolicited Grant Size greater than the value permitted by the Maximum Scheduled Codes value, the CMTS MUST reject the request for the UGS flow or change the CM's Maximum Scheduled Codes value such that it would permit the UGS grants.

6.2.13 S-CDMA Framing

The S-CDMA framer maps mini-slots to spreading codes and spreading intervals by arranging them as symbols within an S-CDMA frame. It also performs an interleaving function, to provide protection against impulse noise. The S-CDMA framer's function of mapping mini-slots to spreading codes and spreading intervals is illustrated in Section 6.2.12, "S-CDMA Framing and Interleaver." As previously described, an S-CDMA frame is defined by the number of spreading intervals per frame, codes per mini-slot, and number of active codes. The framer uses this information to map the mini-slots of a transmission into frames. The framer maps complete grants so that any interleaving which is performed is not constrained by individual mini-slot boundaries. The framer MUST align transmissions to begin and end on mini-slot boundaries. Within a transmission, the framer numbers the symbols or bits and allocates them to codes and spreading intervals independent of the mini-slot mapping. When using TCM encoding, the TCM encoded symbols from the TCM encoder are split into two subsymbols consisting of the coded subsymbol which is the two bits and the parity generated from the convolution encoder, and the uncoded subsymbol consisting of the rest of the bits. When TCM is off, the randomizer output is treated as a continuous bit stream ignoring byte boundaries, as specified in Section 6.2.14, "Symbol Mapping."

6.2.13.1 Subframe Definition

The S-CDMA framer performs interleaving independently of mini-slots. Interleaving is constrained by subframe boundaries, where a subframe is a rectangular subset of an S-CDMA frame over which interleaving is performed. A subframe is normally an integer number of Reed-Solomon code words to enhance protection from impulse noise.

Given an S-CDMA frame which is N_a active code by K spreading intervals, a subframe is defined to be a group of R contiguous rows, where R is an integer in the range from 1 to N_a . A subframe is defined to exist entirely within a single frame and does not span multiple frames. Each subframe contains $R*K$ locations and each location holds one symbol used for mapping and spreading. Each transmission MUST start with a new subframe. The last subframe of a frame MUST be shortened to fit entirely within a single S-CDMA frame. The last subframe of a transmission

MUST be shortened to fit within the granted mini-slots. In both of these cases the subframe will be only R' rows instead of R rows where $R' \leq R$. Figure 6-18 shows a subframe consisting of R rows and K spreading intervals within an S-CDMA frame.




Figure 6-18 - Subframe structure

The parameters that define a subframe and the numbering within a subframe are *codes per subframe* and *interleaver step size*. These two parameters are specified as part of the burst attributes and can vary between burst profiles. These parameters determine the size of the subframe, and also how the subframe is filled with symbols. The valid range for *codes per subframe* is from one to the number of active codes in use. The parameter *interleaver step size* is used while putting TCM coded subsymbols and preamble symbols into the frame. Both of these types of symbols fill in subframes first along a row, and the *interleaver step size* parameter indicates the spreading interval increment to be used while filling in the symbols.

6.2.13.2 Framer Operation

The symbols entering the framer MUST be placed into the framer according to the following sets of rules. There are two sets of rules which apply to different types of input symbols. Preamble symbols and coded TCM subsymbols follow one set of rules, while non-TCM encoded symbols and uncoded TCM subsymbols follow the second set of rules. The rules are specified in the following sections.

6.2.13.2.1 Rules for Preamble and Coded TCM Symbols

The CM preamble (whether TCM is on or off) and the coded TCM subsymbols MUST fill in the frame according to the following rules:

1. The first symbol or subsymbol MUST be placed in the first spreading interval of the first row of the granted mini-slot. In Figure 6-18 this would be row 1, spreading interval 0, assuming that this is the start of the first mini-slot of the grant.
2. Subsequent symbols MUST be placed at the next available spreading interval Interleaver Step Size away from the previous. For instance if the previous symbol was placed at spreading interval X , the next symbol is placed at $X + \text{Interleaver Step Size}$.

3. If the addition of the Interleaver Step Size results in the next location being beyond the end of the frame, the next location MUST be located modulo the frame length. For instance if $J + \text{Interleaver Step Size} = K+1$, then the next location would be spreading interval 1.
4. If the next location is already occupied, then the spreading interval MUST be incremented by 1 until the next unoccupied spreading interval is located. For instance if the desired location is spreading interval X and spreading interval X is occupied, but not X+1, then X+1 would be used.
5. After filling all of the spreading intervals of a single row, the operation is repeated, starting with the next row and step 1 above.
6. After placing all of the preamble and data symbols into the frame, the remaining symbols in the burst, as defined by the burst boundaries of Section 6.2.5.1.1, MUST be filled with zero data symbols which will be mapped to non-zero power.
7. Any locations that have only a TCM uncoded subsymbol MUST be filled with zero bits in the coded subsymbol portion before mapping and spreading.

6.2.13.2.2 Rules for Uncoded Symbols and the Uncoded TCM Subsymbols

Symbols without TCM encoding and uncoded TCM subsymbols MUST fill subframes according to the following rules.

1. The first symbol MUST be placed in the first available code of the first available spreading interval of the subframe after the preamble has been placed into the frame. The symbols are filled from row 1 through row R and, after filling a spreading interval, the next spreading interval is filled from row 1 through row R.
2. Uncoded symbols and the uncoded portion of TCM symbols MUST NOT be placed in the same frame location (spreading interval, code) as a preamble symbol. For instance if there is a preamble symbol in row X, spreading interval Y; and row (X+1), spreading interval Y is unused, the symbol should be placed into row (X+1), spreading interval Y.
3. Subsequent symbols MUST be placed in the next available row of the first available spreading interval of the current subframe. This causes the subframe to be filled column-wise bottom to top and then from left to right. For instance if row 1 through row R of spreading interval X is already occupied, the next symbol would be placed into the first available row of spreading interval X+1.
4. After completely filling a subframe, the next subframe MUST begin as specified in step 1 above.
5. The number of rows contained in the last subframe of a frame MUST be reduced to fit entirely within the frame if there is not adequate space for a full subframe.
6. The number of rows contained in the last subframe of a grant of mini-slots MUST be reduced to fit entirely within the granted mini-slots if there is not adequate space for a full subframe within the grant.
7. After placing all of the data symbols into the frame, the remaining symbols in the burst, as defined by the burst boundaries of Section 6.2.5.1.1, MUST be filled with zero data symbols which will be mapped to non-zero power.
8. Any locations that have only a TCM coded subsymbol MUST be filled with zero bits in the uncoded subsymbol portion before mapping and spreading.

6.2.13.2.3 Subframe Example

Figure 6-19 below shows an example which follows the above specified rules. Each box in the figure represents a symbol which can contain either a preamble symbol, an uncoded symbol when not using TCM, or an uncoded and coded subsymbol when using TCM. In this example there are 9 spreading intervals in the frame, 3 rows for the subframe, an Interleaver Step Size of 3, and the preamble is 4 symbols. Based on these parameters, the subframe would be filled as shown. If the data is TCM encoded, the Cs would represent locations of the coded subsymbols and the Us represent the locations of the uncoded subsymbols. If the TCM is not used, then the symbols would be placed according to the Us only.




Figure 6-19 - Symbol Numbering With and Without TCM

6.2.13.2.4 Frame Transmission

Once a frame is completed and ready for transmission, the symbols MUST be mapped and spread in spreading interval order. This means that spreading interval 0, as shown in Figure 6-18, MUST be the first spreading interval on the wire. For TCM encoded data, the coded and uncoded subsymbols from each location in the frame MUST be combined to create complete symbols before mapping and spreading. This corresponds to creating a new symbol where the coded portion of the symbol is C^i and the uncoded portion is U^j . The preamble symbols remain intact.

6.2.14 Symbol Mapping

The modulation mode is configurable via MAC messages. Differential encoded QPSK and 16 QAM are available for TDMA channels. QPSK, 8 QAM, 16 QAM, 32 QAM, and 64 QAM are available for TDMA and S-CDMA channels. TCM encoded QPSK, 8 QAM, 16 QAM, 32 QAM, 64 QAM and 128 QAM are available for S-CDMA channels. The symbols transmitted in each mode and the mapping of the input bits to the I and Q constellation MUST be as defined in Table 6-3. In the table, x^1 represents the LSB of each of the symbol maps and x^2, x^3, x^4, x^5, x^6 and x^7 represents the MSB for QPSK, 8 QAM, 16 QAM, 32 QAM, 64 QAM, and 128 QAM respectively. The MSB MUST be the first bit in the serial data into the symbol mapper. The MSB MUST be mapped to the MSB of the symbol map. The number of data bytes may not map into an integer number of symbols. In this case, the last symbol MUST be padded with zero bits in the LSB locations after all data bits are processed.

All constellations are defined on a common integer grid in Figure 6-20. This defines each QAM symbol with 5-bit values on each (I and Q) axis. The relative symbol amplitudes defined by the grid MUST be maintained across all constellations. Different constellations may be used, for example, in different burst profiles, in preamble and data symbols within the same burst, and in modulating different spreading codes within a frame.

In Figure 6-20, E_{av} denotes the average constellation energy for equally likely symbols. For each constellation the integer values of E_{av} and differences in dB compared to 64 QAM, G_{const} , are given. The QPSK0 constellation is employed for low-power preamble and all QPSK data symbols. Use of QPSK1 is restricted to high-power preamble symbols.

Table 6-3 - I/Q Mapping

QAM Mode	Input bit Definitions
QPSK	x^2x^1
8 QAM	$x^3x^2x^1$
16 QAM	$x^4x^3x^2x^1$
32 QAM	$x^5x^4x^3x^2x^1$
64 QAM	$x^6x^5x^4x^3x^2x^1$
128 QAM	$x^7x^6x^5x^4x^3x^2x^1$

The upstream symbol constellations MUST be as shown in Figure 6-20.

The upstream QPSK Gray-coded and differential symbol mapping MUST be as shown in Figure 6-21.

The upstream 8 QAM symbol mapping MUST be as shown in Figure 6-22.

The upstream 16 QAM Gray-coded symbol mapping MUST be as shown in Figure 6-23.

The upstream 16 QAM differential symbol mapping MUST be as shown in Figure 6-23.

The upstream 32 QAM symbol mapping MUST be as shown in Figure 6-24.

The upstream 64 QAM Gray-coded symbol mapping MUST be as shown in Figure 6-25.

The TCM symbol mapping used for S-CDMA are shown in Figure 6-26 through Figure 6-28.

The upstream QPSK TCM symbol mapping MUST be as shown in Figure 6-26.

The upstream 8 QAM TCM symbol mapping MUST be as shown in Figure 6-26.

The upstream 16 QAM TCM symbol mapping MUST be as shown in Figure 6-27.

The upstream 32 QAM TCM symbol mapping MUST be as shown in Figure 6-27.

The upstream 64 QAM TCM symbol mapping MUST be as shown in Figure 6-28.

The upstream 128 QAM TCM symbol mapping MUST be as shown in Figure 6-28.


If differential quadrant encoding is enabled, then the currently-transmitted symbol quadrant is derived from the previously transmitted symbol quadrant and the current input bits via Table 6-4. If differential quadrant encoding is enabled, the upstream PMD sublayer MUST apply these differential encoding rules to all transmitted symbols (including those that carry preamble bits). Differential quadrant encoding is only available for QPSK and 16 QAM on TDMA channels. In Table 6-4, I(1)Q(1) refers to x^2x^1 and x^4x^3 from Table 6-3 for QPSK and 16 QAM cases respectively.

Table 6-4 - Definition of Differential Quadrant Coding


Current Input Bits I(1) Q(1)	Quadrant Phase Change	MSBs of Previously Transmitted Symbol	MSBs for Currently Transmitted Symbol
00	0	11	11
00	0	01	01
00	0	00	00
00	0	10	10
01	90	11	01
01	90	01	00
01	90	00	10
01	90	10	11
11	180	11	00
11	180	01	10
11	180	00	11
11	180	10	01
10	270	11	10
10	270	01	11
10	270	00	01
10	270	10	00

QPSK0: $E_{av} = 128$ ($G_{const} = -1.18$ dB rel to 64 QAM)


QPSK1: $E_{av} = 288$ ($G_{const} = +2.34$ dB)




8 QAM-DS: $E_{av} = 160$ ($G_{const} = -0.21$ dB)




16 QAM-SQ: $E_{av} = 160$ ($G_{const} = -0.21$ dB)



32 QAM-DS: $E_{av} = 168$ ($G_{const} = 0$ dB)



64 QAM-SQ: $E_{av} = 168$ ($G_{const} = 0$ dB)



128 QAM-DS: $E_{av} = 170$ ($G_{const} = +0.05$ dB)




Figure 6-20 - Symbol Constellations




Figure 6-21 - QPSK Gray and Differential Symbol Mapping




Figure 6-22 – 8 QAM Symbol Mapping

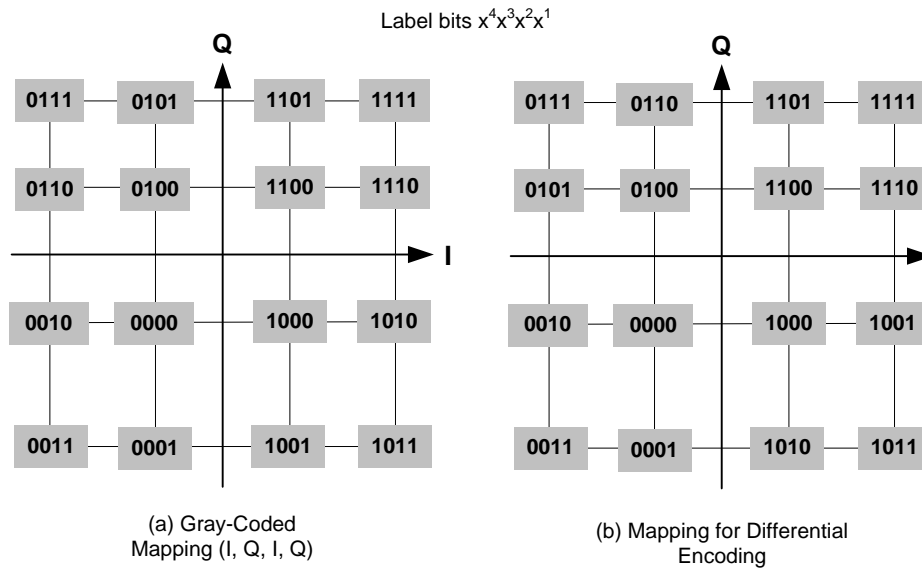


Figure 6-23 – 16 QAM Symbol Mapping




Figure 6-24 – 32 QAM Symbol Mapping

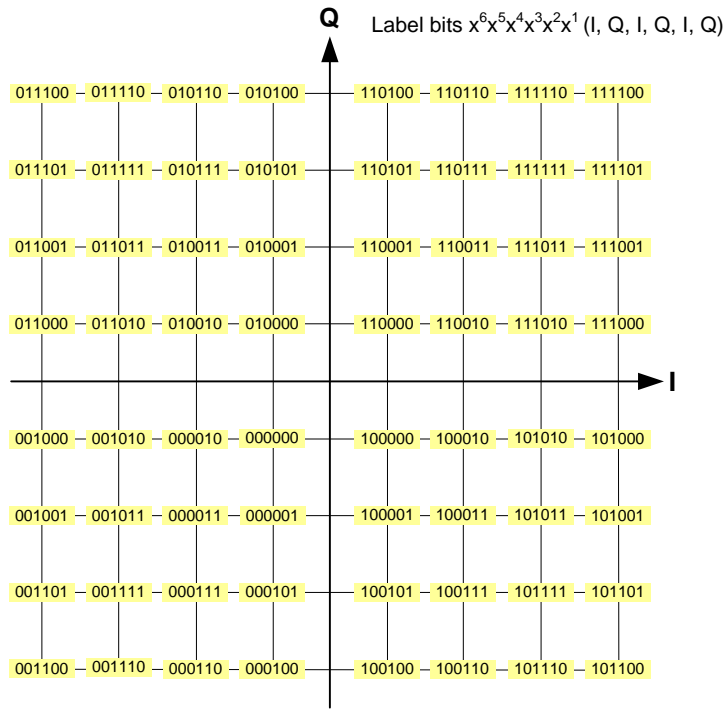


Figure 6-25 – 64 QAM Symbol Mapping






Figure 6-26 - QPSK and 8 QAM TCM Symbol Mapping

Octal labels o^2o^1 (= bin. labels $x^4x^3x^2x^1$)



Subset B0


Octal labels o^2o^1 (= bin. labels $x^5x^4x^3x^2x^1$)



Subset B1


Figure 6-27 – 16 QAM and 32 QAM TCM Symbol Mapping

64 QAM: octal labels o^3o^1
(= bin. labels $x^4x^3x^2x^1$)



Subset B0

128 QAM: octal labels $o^3o^2o^1$
(= bin. labels $x^7x^6x^5x^4x^3x^2x^1$)



Subset B1

Figure 6-28 – 64 QAM and 128 QAM TCM Symbol Mapping

6.2.15 S-CDMA Spreader

The basis of signal transmission with S-CDMA is direct-sequence spread-spectrum modulation. S-CDMA employs a family of orthogonal digital code words, called spreading codes, to simultaneously transmit up to 128 modulation symbols. In each spreading interval, a vector P_k is transmitted such that:

$$P_k = S_k * C,$$

where S_k is a vector, $[s_{k,127}, s_{k,126}, \dots, s_{k,0}]$, of modulation symbols on the integer grid of Section 6.2.14 to be transmitted in spreading interval k , and C is a matrix:

$$C = \begin{bmatrix} C_{127,127} & C_{127,126} & \dots & C_{127,0} \\ C_{126,127} & C_{126,126} & \dots & C_{126,0} \\ \dots & \dots & \dots & \dots \\ C_{0,127} & C_{0,126} & \dots & C_{0,0} \end{bmatrix} = \begin{bmatrix} x_1 & x_{127} & \dots & x_2 & -1 \\ x_2 & x_1 & \dots & x_3 & -1 \\ \dots & \dots & \dots & \dots & -1 \\ x_{127} & x_{126} & \dots & x_1 & -1 \\ 1 & 1 & 1 & 1 & 1 \end{bmatrix}$$

where the rows of C are the 128 spreading codes such that code $i = [c_{i,127}, c_{i,126}, \dots, c_{i,0}]$. The matrix C may be written in shorthand form in terms of its rows as follows:

$$C = \begin{bmatrix} \text{code 127} \\ \text{code 126} \\ \dots \\ \text{code 2} \\ \text{code 1} \\ \text{code 0} \end{bmatrix}$$

The result of the spreading operation is the transmission vector P_k which has 128 elements, $(p_{k,127}, p_{k,126}, \dots, p_{k,0})$, where each element is transmitted at the signaling rate, with element $p_{k,0}$ transmitted first in time. The first element S_0 into the spreader is defined as follows. As a point of reference, for 128 allocated codes, and considering the first column of the framer ($k = 0$), S_0 is the first symbol in time to enter the framer, occupies the lower left element of the framer, and is the first element into the spreader.

The set of orthogonal codes used for the spreading operation is quasi-cyclic and consists of values which are either +1 or -1. Code 0 consists of 128 elements, each of which has value +1. For each of the other spreading codes, code i , the element $c_{i,0}$ is -1 and the remaining elements are obtained by a cyclic shift of a sequence x as is shown in the above matrix in this section.

The sequence x_i is defined such that the elements corresponding to the following set of indices are equal to -1:

{2 3 4 5 6 7 9 10 11 13 16 17 18 19 20 21 25 26 28 30 31 33 34 35 37 39 40 41 49 51 52 55 56 59 60 61 65 66 67 69 72 73 74 77 78 79 81 84 90 92 94 97 100 101 103 106 109 110 111 114 117 119 121};

The remaining elements of code 1 have a value of +1.

Each code i is obtained by cyclic shift to the left (in the direction of increasing indices) of code $(i-1)$, where the element $c_{i,0}$ has a value of -1 and does not take part in the cyclic shift.

Although each code is defined to have equal power, the spread symbols may have slightly unequal power since the symbols at the input to the spreader have varying values of E_{av} according to the integer symbol grid of Section 6.2.14.

If a CM has not been assigned to use a particular code i at a spreading time interval k , then in its computation of its transmission vector P_k , it will set $s_{k,i}$ to numerical zero. The assignment of codes to the CM is performed by the framer as it assigns a burst of symbols a particular order in the two-dimensional space of codes and time. This symbol sequencing is described in detail in Section 6.2.13.

The I and Q components of the symbols are spread using the same spreading code.

In the matrix multiple of the equation above, and subsequent CM processing prior to the D/A, there is an essential clipping operation wherein, as an example, filtered (pulse-shaped) elements of P_k in excess of some vendor-specific absolute value are clipped (retaining complex angle) to this absolute value. This non-linear operation, deviating from the equation above and the subsequent linear processing prior to the D/A, is essential for meeting spurious emission and MER requirements safely and efficiently, while operating at the highest CM average transmit power levels (see Table 6-12, "User Unique Burst Parameters").

6.2.15.1 Code Hopping

Code hopping refers to a systematic re-ordering of the rows of the original spreading matrix C , so that at each spreading interval k , a new code matrix C_k is produced. A pseudo-random number generator determines a cyclic shift of a subset of the rows of the original matrix C . Two code hopping modes are defined: Code Hopping Mode 1 is used exclusively with Selectable Active Codes Mode 1; Code Hopping Mode 2 is used exclusively with Selectable Active Codes Mode 2.

In S-CDMA mode, the CM MUST support Code Hopping Mode 1 and SHOULD support Code Hopping Mode 2. The CMTS MAY support code hopping.

6.2.15.1.1 Code Hopping Mode 1

In Code Hopping Mode 1, when the number of active codes $N_a < 128$, the cyclic codes are hopped (cyclically shifted), while code 0, the all 1s code, remains fixed at the bottom of the matrix. The unused codes are located in the $(128 - N_a)$ bottom rows of the hopped matrix C_k . Thus, for $N_a < 127$, the set of unused codes changes at each spreading interval. When $N_a = 128$, all codes are hopped, including code 0. The hopped spreading matrix is defined by:

$$C_k = \begin{bmatrix} C_{f(k,127),127} & C_{f(k,127),126} & \cdots & C_{f(k,127),0} \\ C_{f(k,126),127} & C_{f(k,126),126} & \cdots & C_{f(k,126),0} \\ \cdots & \cdots & \cdots & \cdots \\ C_{f(k,0),127} & C_{f(k,0),126} & \cdots & C_{f(k,0),0} \end{bmatrix}$$

where,

$$f(k,i) = \begin{cases} \text{modulo}(128 - lfsr_out(k) + i, 128), & N_a = 128, 0 \leq i \leq 127 \\ \text{modulo}(126 - lfsr_out(k) + i, 127) + 1, & N_a < 128, 1 \leq i \leq 127 \end{cases}$$

The matrix elements $c_{i,j}$ are the elements of the original code matrix C

k indicates the spreading interval;

i indicates the row in the matrix;

$lfsr_out$ is a pseudo-random number described in Section 6.2.15.1.3.

6.2.15.1.2 Code Hopping Mode 2

In Code Hopping Mode 2, the active codes are hopped, while the unused codes remain fixed at the bottom of the matrix. A horizontal line in the following equation shows the division between active and unused codes. The hopped spreading matrix is defined by:

$$C_k = \left[\begin{array}{cccc} C_{f(k,127),127} & C_{f(k,127),126} & \dots & C_{f(k,127),0} \\ C_{f(k,126),127} & C_{f(k,126),126} & \dots & C_{f(k,126),0} \\ \dots & \dots & \dots & \dots \\ C_{f(k,128-N_a),127} & C_{f(k,128-N_a),126} & \dots & C_{f(k,128-N_a),0} \\ \hline C_{u(128-N_a-1),127} & C_{u(128-N_a-1),126} & \dots & C_{u(128-N_a-1),0} \\ \dots & \dots & \dots & \dots \\ C_{u(0),127} & C_{u(0),126} & \dots & C_{u(0),0} \end{array} \right]$$

where,

$$f(k,i) = \text{active_code_list} [\text{modulo}(2*N_a - 128 - \text{hop_number}(k) + i, N_a)]$$

$$128 - N_a \leq i \leq 127$$

The matrix elements $c_{i,j}$ are the elements of the original code matrix C :

k indicates the spreading interval;

i indicates the row in the matrix;

N_a is the number of active codes;

active_code_list is the list of active codes, in increasing order, with each element in the range 0 to 127;

u is the list of unused codes, in increasing order, with each element in the range 0 to 127;

hop_number is a pseudo-random number described in 6.2.15.1.3.

The following is a numerical example of Selectable Active Codes Mode 2 and Code Hopping Mode 2. Figure 6-16 shows $N_a = 124$ active codes with unused code list $u = \{0, 1, 5, 125\}$. Thus, $\text{active_code_list} = \{2, 3, 4, 6, 7, 8, 9, 10, \dots, 123, 124, 126, 127\}$, as annotated at the left of the figure; $\text{active_code_list}(0) = 2$, $\text{active_code_list}(123) = 127$, $u(0) = 0$, and $u(3) = 125$. Assume that at a given spreading interval k , the pseudo-random hop generator returns a value of $\text{hop_number} = 0$. (This results in the same matrix that would occur with code hopping turned off). The spreading matrix C_k consists of the unused codes at the bottom and active codes in increasing order on top:

$$C_k = \begin{bmatrix} \text{code 127} \\ \text{code 126} \\ \text{code 124} \\ \text{code 123} \\ \dots \\ \text{code 8} \\ \text{code 7} \\ \text{code 6} \\ \text{code 4} \\ \text{code 3} \\ \text{code 2} \\ \hline \text{code 125} \\ \text{code 5} \\ \text{code 1} \\ \text{code 0} \end{bmatrix}$$

At another spreading interval, if $\text{hop_number} = 3$, the active-code rows are cyclically shifted vertically by 3 rows, while the unused codes remain fixed, and the spreading matrix C_k becomes:

$$C_k = \begin{bmatrix} \text{code 123} \\ \text{code 122} \\ \text{code 121} \\ \text{code 120} \\ \dots \\ \text{code 9} \\ \text{code 8} \\ \text{code 7} \\ \text{code 6} \\ \text{code 4} \\ \text{code 3} \\ \text{code 2} \\ \text{code 127} \\ \text{code 126} \\ \text{code 124} \\ \hline \text{code 125} \\ \text{code 5} \\ \text{code 1} \\ \text{code 0} \end{bmatrix}$$

6.2.15.1.3 Code hopping generator

The pseudo-random number generator, which determines the spreading matrix reordering during code hopping, utilizes the linear-feedback shift register (LFSR), shown in Figure 6-29. In order to align the CM's code-hopping pseudo-random sequence with that of the CMTS, the pseudo-random generator must output the following value at the first spreading interval of each frame:

$$\text{lfsr_out}(\text{frame_number} * \text{spreading_interval_per_frame})$$

where $\text{lfsr_out}(k)$ is the value of lfsr_out , after k shift register updates, following the code hopping seed load into the LFSR. The bits are defined as $\text{lfsr_out}_{7:1} = s_{7:1}$, where s is the shift register contents, and lfsr_out bit 7 is considered the MSB. In Code Hopping Mode 1, a shift register update equates to one shift of the LFSR. In Code Hopping Mode 2, a shift register update equates to 15 shifts of the LFSR.

The description of the frame counter and the procedures for its synchronization are contained in Section 6.2.12.2, "Mini-slot Numbering." At this reset, a 15-bit initialization value (seed) is loaded into the shift register and is used at the first spreading interval. The 15-bit seed value is configured in response to the Upstream Channel Descriptor message from the CMTS.

At each subsequent spreading interval k , the LFSR is updated. The code hopping mechanism (LFSR and spreading interval index k) is advanced, every spreading interval (128 modulation intervals), in both spreader-on and spreader-off frames.

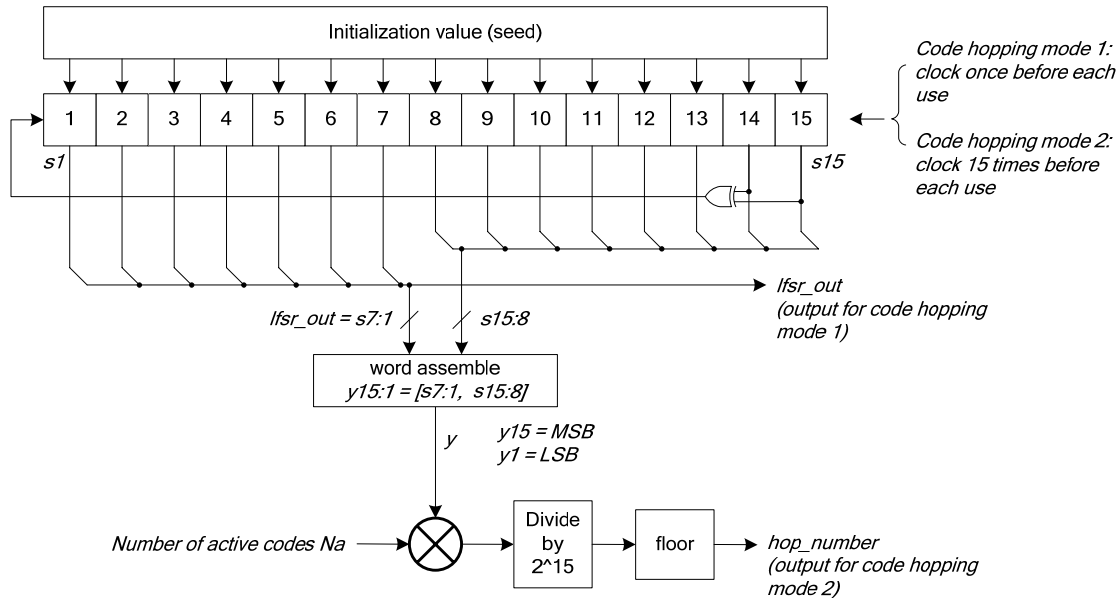


Figure 6-29 - Code Hopping Random Number Generator

6.2.15.1.3.1 Code Hopping Mode 1 pseudo-random number generation

In Code Hopping Mode 1, the LFSR is updated by clocking it once before each use on each spreading interval. The output `lfsr_out` is used to compute the spreading matrix indices, as given by the equation for Code Hopping Mode 1 above.

6.2.15.1.3.2 Code Hopping Mode 2 pseudo-random number generation

In Code Hopping Mode 2, the LFSR is updated by clocking it 15 times before use on each spreading interval, thereby flushing it to reduce the correlation between hops. The LFSR contents are then scaled to produce an approximately uniform distributed pseudo-random hop number in the range 0 to $N_a - 1$. The scaling consists of the following operation:

$$\text{hop_number} = \text{floor} \left(\frac{N_a y}{2^{15}} \right)$$

In this equation, N_a is the number of active codes, and y is the result of assembling the output of the LFSR into a 15-bit word as follows: Bits 7:1 of the LFSR are mapped to the most significant 7 bits of y . The remaining bits 15:8 of the LFSR are used as the least significant 8 bits of y . That is, $y_{15:1} = [s_{7:1}, s_{15:8}]$, where y_{15} is considered the MSB. The output `hop_number` is used to compute the spreading matrix indices, as given by the equation, for Code Hopping Mode 2 above.

Full precision MUST be maintained in the multiplication of $N_a * y$. The product MUST then be truncated, discarding the 15 LSBs, as shown in the figure. Or, equivalent processing may be used that produces the same `hop_number` output under all input conditions. The use of the full 15 bits of the LFSR provides a dithering effect which prevents "clumping" in the hopping distribution.

6.2.16 Transmit Pre-Equalizer

A transmit pre-equalizer of a linear equalizer structure, as shown in Figure 6-30, MUST be configured for each upstream channel used by the CM in response to the Ranging Response (RNG-RSP) message transmitted by the CMTS.

There are two modes of operation for the pre-equalizer of a CM: DOCSIS 1.1 mode, and DOCSIS 2.0 pre-equalization mode: In DOCSIS 1.1 mode, the CM MUST support a (T)-spaced equalizer structure, with 8 taps. The pre-equalizer MAY have 1, 2, or 4 samples per symbol, with a tap length longer than 8 symbols: In DOCSIS 1.1 pre-equalization mode, for backwards compatibility, the CMTS MAY support fractionally spaced equalizer format (T/2 and T/4). In DOCSIS 2.0 pre-equalization mode, the pre-equalizer MUST support a symbol (T)-spaced equalizer structure with 24 taps.

In DOCSIS 1.x-only logical channels (type 1 logical channels [MULPI]), the CM and the CMTS MUST use DOCSIS 1.1 pre-equalization mode.

In DOCSIS 2.0-only or DOCSIS 3.0-only logical channels (type 3 or type 4 logical channels [MULPI]), the CM and the CMTS MUST use DOCSIS 2.0 pre-equalization mode.

In DOCSIS 1.x/2.0 mixed logical channels (type 2 logical channels [MULPI]), the CM and the CMTS MUST use DOCSIS 1.1 pre-equalization mode from initial ranging until DOCSIS 2.0 mode is enabled or until Multiple Transmit Channel mode is enabled in the registration process (if either is enabled). The CM and CMTS MUST use DOCSIS 2.0 pre-equalization mode after DOCSIS 2.0 is enabled for the CM. In the case that a CM is placed into Multiple Transmit Channel mode, the CM and CMTS use DOCSIS 2.0 pre-equalization for all bursts on type 2 logical channels whether the burst descriptor in the UCD for the transmitted burst is a Type 4 or Type 5 burst descriptor.

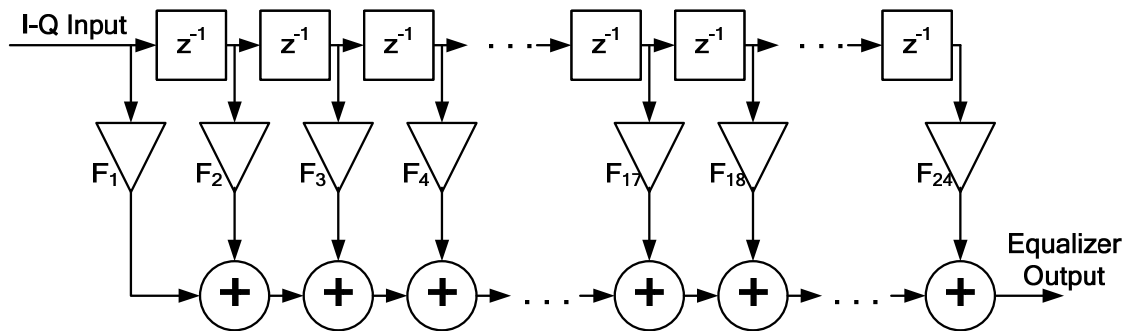


Figure 6-30 - Transmit Pre-Equalizer Structure

The RNG-RSP MAC message carries the CM transmit equalization information, and may instruct the CM to either convolve the equalizer coefficients, or (in DOCSIS 2.0 pre-equalization mode only) load them directly. When the CM is instructed to convolve the transmit equalizer coefficients, it MUST convolve the coefficients sent by the CMTS in the RNG-RSP with the existing coefficients, to get the new coefficients. After convolving, the CM MUST truncate the convolution result, such that 24 taps (8 taps in DOCSIS 1.1 pre-equalization mode) remain after the truncation, with the main tap located at the tap designated by the last RNG-RSP received by the CM. The operation of the convolution is formulized by the following equation:

$$F_n^{m+1} = \sum_{k=\max(1-L^{m+1}, n+L^m-L^{m+1}-24)}^{\min(24-L^{m+1}, n+L^m-L^{m+1}-1)} F_{n-k+L^m-L^{m+1}}^m \cdot \hat{F}_{k+L^{m+1}}^{m+1}, \quad n = 1 \dots 24$$

where:

F_n^m are the coefficients prior to the convolution;

F_n^{m+1} are the coefficients after the convolution;

\hat{F}_n are the coefficients sent from the CMTS;

L^m is the main tap location prior to the convolution;

L^{m+1} is the main tap location after the convolution as dictated by the CMTS.

In DOCSIS 2.0 pre-equalization mode, the CMTS MAY instruct the CM to load the transmit pre-equalizer coefficients. When the CM is instructed to load the transmit equalizer coefficients, it MUST load the coefficients sent by the CMTS into the pre-equalizer coefficients after proper normalization, if necessary.

In DOCSIS 1.x-only logical channels, in response to an initial ranging request and periodic ranging requests prior to CM registration, when the CMTS sends the pre-equalizer coefficients, the CMTS MUST compute and send them with an equalizer length of 8 and in T-spaced format, where T is the modulation interval. After registration, the CMTS MAY use a fractionally spaced equalizer format (T/2- or T/4-spaced), with a longer tap length, to match the CM pre-equalizer capabilities that the CMTS learned from the REG-REQ message modem capabilities field.

In DOCSIS 2.0-only or DOCSIS 3.0-only logical channels, the CMTS MUST compute and send the pre-equalizer coefficients with an equalizer length of 24, and in T-spaced format, at all times.

In DOCSIS 1.x/2.0 mixed logical channels, in response to an initial ranging request and periodic ranging requests prior to CM registration, when the CMTS sends the pre-equalizer coefficients, the CMTS MUST compute and send them with an equalizer length of 8 and in T-spaced format. After registration, if the DOCSIS 1.1 pre-equalization mode is enabled, the CMTS MAY use a fractionally spaced equalizer format (T/2- or T/4-spaced), with a longer tap length, to match the CM pre-equalizer capabilities that the CMTS learned from the REG-REQ message modem capabilities field. If DOCSIS 2.0 pre-equalization mode or Multiple Transmit Channel mode is enabled for the CM, the CMTS MUST use a T-spaced equalizer structure with 24 taps. If the first update of the pre-equalizer after the activation of DOCSIS 2.0 pre-equalization mode uses "convolve" mode, the CM MUST zero-pad the existing 8-tap filter to a 24-tap filter, and convolve, according to the rules above.

Prior to making an initial ranging request, and whenever the upstream channel frequency or upstream channel modulation rate changes, the CM MUST initialize the coefficients of the pre-equalizer to a default setting in which all coefficients are zero, except the real coefficient of the first tap (i.e., F_1). Whenever the main location is changed, the CM, not the CMTS, MUST compensate for the delay (ranging offset), due to a shift from the previous main tap location, to a new main tap location of the equalizer coefficients, sent by the CMTS (in both "convolve" and "load" operations). The pre-equalizer coefficients are then updated through the subsequent ranging process (unicast initial ranging and periodic ranging).

In DOCSIS 1.1 pre-equalization mode, the CMTS MUST NOT move the main tap location during periodic ranging.

In DOCSIS 1.1 pre-equalization mode, the CMTS MUST NOT instruct the CM to load the transmit equalizer coefficients.

In DOCSIS 2.0 pre-equalization mode, the CMTS MAY move the main tap location during unicast initial ranging or periodic ranging.

Equalizer coefficients may be included in every RNG-RSP message, but typically they only occur when the CMTS determines that the channel response has significantly changed. The frequency of equalizer coefficient updates in the RNG-RSP message is determined by the CMTS.

The CM MUST normalize the transmit equalizer coefficients in order to guarantee proper operation (such as not to overflow or clip). The CM MUST NOT change its target transmit power due to gain or loss of the new coefficients in both "convolve" and "load" operations. The target power is defined in Section 6.2.19, "Transmit Power Requirements".

In DOCSIS 1.1 mode, if the CM equalizer structure implements the same number of coefficients as assigned in the RNG-RSP message, then the CM MUST NOT change the location of the main tap in the RNG-RSP message. If the CM equalizer structure implements a different number of coefficients than defined in the RNG-RSP message, the CM MAY shift the location of the main tap value. The CM MUST adjust its ranging offset, in addition to any adjustment in the RNG-RSP message, by an amount that compensates for the movement of the main tap location.

6.2.17 Spectral Shaping

The CM upstream transmitter MUST approximate a Nyquist square-root raised-cosine pulse-shaping filter with roll-off factor $\alpha = 0.25$. For operation with a DOCSIS 3.0 CMTS (unless the CMTS is operating as a pre-3.0 DOCSIS CMTS), the CM -30 dB transmitted bandwidth MUST NOT exceed the Channel Width values in Table 6-5. For backward compatibility with a CMTS operating as a pre-3.0-DOCSIS CMTS, the CM -30 dB transmitted bandwidth MUST NOT exceed the Channel Width values in Table 6-6. The Channel Width values are given analytically by:

$$\text{ChannelWidth} = \text{ModulationRate} * (1 + \alpha).$$

Table 6-5 - Maximum Channel Width for Operation with a DOCSIS 3.0 CMTS

Modulation Rate (kHz)	Channel Width (kHz)
1280	1600
2560	3200
5120	6400

Table 6-6 - Maximum Channel Width for Operation with a pre-3.0-DOCSIS CMTS

Modulation Rate (kHz)	Channel Width (kHz)	Applicability
160	200	CM MAY support
320	400	CM MAY support
640	800	CM MAY support
1280	1600	CM MUST support
2560	3200	CM MUST support
5120	6400	CM MUST support

6.2.17.1 Upstream Frequency Agility and Range

The CM MUST support operation over the frequency range of 5-42 MHz.

Additionally, the CM MAY support selectable operation over a) the frequency range of 5-42 MHz, or b) the extended frequency range of 5-85 MHz, edge to edge.

Offset frequency commands MUST be supported per Table 6-12.

6.2.17.2 Spectrum Format

The upstream modulator MUST provide operation with the format $s(t) = I(t)*\cos(\omega t) - Q(t)*\sin(\omega t)$, where t denotes time and ω denotes angular frequency.

6.2.18 Relative Processing Delays

The CM MAP processing delay is the time provided between arrival of the last bit of a MAP message at a CM and the effectiveness of this MAP. During this time, the CM should process the MAP message and fill its interleavers (or its framer, in S-CDMA mode) with encoded data. The CMTS MUST transmit the MAP message early enough to allow the CM MAP processing delay, as specified below.

The CM MAP processing delay, D_p , is given by the equations:

$$D_p = 600 + \frac{M}{5.12} \mu\text{sec}$$

$$M = \begin{cases} I_r N_r, & I_r \neq 0 \\ B_r, & I_r = 0 \end{cases}$$

where M is the number of elements in the CM interleavers (in the case of TDMA), or framer (in the case of S-CDMA). In DOCSIS 1.x mode, $M = 0$. Note that in the above equations, the values for B_r and $I_r * N_r$ are taken to be the maximum from all of the specified burst types in a particular UCD.

In S-CDMA mode, $M = 128(K+1)$, where K is the number of spreading intervals per frame. This is the time required for processing an S-CDMA frame, plus an extra spreading interval. For example, in the case of $K = 32$, which corresponds to the maximum framer size, the CM MAP processing time is 1425 μs , assuming a modulation rate of 5.12 MHz.

Note: The CM MAP processing delay does not include downstream FEC de-interleaving delay.

Note: The "effectiveness of the MAP" relates to the beginning of the burst frame at the RF output of the CM. In the S-CDMA mode, "effectiveness of the MAP" relates to the beginning (at the RF output of the CM) of the first spreading interval of the S-CDMA frame which contains the burst.

6.2.19 Transmit Power Requirements

The following requirements apply with Multiple Transmit Channel mode enabled. Requirements with Multiple Transmit Channel mode disabled are addressed in Section 6.2.19.4 below.

The CM MUST 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 [MULPI], Section 6.4.5. Such adjustments by the CM MUST be within the ranges of tolerances described below. A CM MUST confirm that the transmit power per channel limits are met after a RNG-RSP is received or after a UCD change, for each of the CM's active channels that is referenced. An active channel for a CM is defined as any channel for which the CM has been given a grant, which it will use for channel initialization, or ranging, or any channel for which the CM is "ranged." The set of "active channels" is also called the Transmit Channel Set. (There will be some mechanism by which a CMTS can command a CM to deactivate an active channel, for the purpose of either reducing the number of active channels on that CM or for the purpose of adding a different active channel as a replacement upstream channel for that CM.) 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) transmitted in the data symbols of a burst, assuming equal likely QAM symbols, measured at the F-connector of the CM. Total transmit power is defined as the sum of the transmit power per channel of each channel transmitting a burst at a given time. Maximum and minimum transmit power per channel level requirements refer to the CM's target transmit power per channel levels, defined as the CM's estimate of its actual transmit power per channel. The actual transmitted power per channel MUST be within 2 dB of the target power. The target transmit power per channel MUST be variable over the range specified in Table 6-12.

With Multiple Transmit Channel mode enabled, let $P_{load} = P_{hi} - P_r$, for each channel, using the definitions for P_{hi} and P_r in the following subsections of 6.2.19. The channel corresponding to the minimum value of P_{load} is called the highest loaded channel, and its value is denoted as P_{load_1} , even if there is only one channel in the Transmit Channel Set. Note that the highest loaded channel is not necessarily the channel with the highest transmit power. Similarly, the channel with the second lowest value of P_{load} is denoted as the second highest loaded channel, and its loading value is denoted as P_{load_2} ; the channel with the third lowest value of P_{load} is the third highest loaded channel, and its loading value is denoted as P_{load_3} ; and the channel with the fourth lowest loading value is denoted as the fourth highest loaded channel, and its loading value is P_{load_4} . $P_{load_min_set}$ defines the upper end of the Dynamic Range Window for the CM with respect to P_{hi} for each channel. $P_{load_min_set}$ will limit the maximum power possible for each active channel to a value less than P_{hi} when $P_{load_min_set}$ is greater than zero. $P_{load_min_set}$ is a value commanded to the CM from the CMTS. (In the case of fewer than 4 channels in the Transmit Channel Set, P_{load_n} will only be valid for the n active upstream channels.) $P_{load_min_set}$, P_{load_n} , P_{hi_n} , P_r_n , etc., are defined only when Multiple Transmit Channel mode is enabled.

CMTS commands adjusting P_r_n MUST be consistent with the $P_{load_min_set}$ previously assigned to the CM and with the following limits:

$$P_{load_min_set} \leq P_{hi_n} - P_r_n \leq P_{load_min_set} + 12 \text{ dB}$$

A value is computed, P_{low_multi} , which sets the lower end of the transmit power Dynamic Range Window given the upper end of the range (which is determined by $P_{load_min_set}$).

$$P_{low_multi} = P_{hi} - P_{load_min_set} - 12 \text{ dB}.$$

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

When the CMTS sends a new value of $P_{load_min_set}$ 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 $P_{load_min_set}$ 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_{r_n} may not be implemented immediately upon reception at the CM if the n^{th} channel is bursting. Commanded changes to P_{r_n} may occur simultaneously with the command to change $P_{load_min_set}$. The CMTS SHOULD NOT issue a change in $P_{load_min_set}$ after commanding a change in P_{r_n} until after also providing a sufficient reconfiguration time on the n^{th} channel. The CMTS SHOULD NOT issue a change in $P_{load_min_set}$ after commanding a prior change in $P_{load_min_set}$ until after also providing a global reconfiguration time for the first command. Also, the CMTS SHOULD NOT issue a change in P_{r_n} a) until after providing a global reconfiguration time following a command for a new value of $P_{load_min_set}$; and b) until after providing a sufficient reconfiguration time on the n^{th} channel after issuing a previous change in P_{r_n} . In other words, the CMTS is to avoid sending consecutive changes in P_{r_n} and/or $P_{load_min_set}$ to the CM without a sufficient reconfiguration time for instituting the first command. When a concurrent new value of $P_{load_min_set}$ and change in P_{r_n} are commanded, the CM MAY wait to apply the change in P_{r_n} at the next global reconfiguration time (i.e., concurrent with the institution of the new value of $P_{load_min_set}$) rather than applying the change at the first sufficient reconfiguration time of the n^{th} channel; the value of $P_{load_min_set}$ which applies to the new P_{r_n} is the concurrently commanded $P_{load_min_set}$ value, and thus if the change to P_{r_n} falls outside the Dynamic Range Window of the old $P_{load_min_set}$, then the CM MUST wait for the global reconfiguration time to apply the change in P_{r_n} .

The CM MUST ignore the command to increase the per channel transmit power if such a command would cause P_{load} for that channel to drop below $P_{load_min_set}$. 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 $P_{load_min_set}$ to a smaller value than normal, and fluctuation of the power per channel in the highest loaded channel is expected to wander. The CM also MUST ignore any commanded change of per channel transmit power which would result in P_r falling below the Dynamic Range Window, determined by $P_{load_min_set}$. The CM also MUST ignore a command changing $P_{load_min_set}$ such that existing values of P_r would fall outside the new Dynamic Range Window.

The spurious performance requirements of Sections 6.2.22.1, 6.2.22.1.1, and 6.2.22.1.2 apply when the CM is operating within certain ranges of values for P_{load_n} , for $n = 1$ to the active number of upstream channels, and for certain ranges of the number of spreading codes being transmitted as a fraction of the total number of active codes on a channel for S-CDMA channels, as detailed in those sections.

Transmit power per channel, for each channel, as reported by the CM in the MIB, is referenced to the 64 QAM constellation. When transmitting with other constellations, a slightly different transmit power will result, depending on the constellation gain in Table 6-7, Table 6-8 and Table 6-9 below. As an example, if the reported power in a channel is 30 dBmV, 64 QAM will be transmitted with a target power of 30 dBmV in that channel, while QPSK will be transmitted with 28.82 dBmV in that channel.

Table 6-7 - Constellation Gains and Per Channel Power Limits with One Channel in the Transmit Channel Set

Constellation	Constellation Gain G_{const} Relative to 64 QAM (dB)	P_{min} (dBmV)			P_{max} (dBmV) TDMA	P_{max} (dBmV) S-CDMA	$P_{min} - G_{const}$ (dBmV)	$P_{max} - G_{const}$ (dBmV) TDMA	$P_{max} - G_{const}$ (dBmV) S-CDMA
		L	M	H					
QPSK	-1.18	17	20	23	61	56	18.18	62.18	57.18
8 QAM	-0.21	17	20	23	58	56	17.21	58.21	56.21
16 QAM	-0.21	17	20	23	58	56	17.21	58.21	56.21
32 QAM	0.00	17	20	23	57	56	17.00	57.00	56.00
64 QAM	0.00	17	20	23	57	56	17.00	57.00	56.00
128 QAM	0.05	17	20	23	N/A	56	16.95	N/A	55.95

(P_{min} is a function of Modulation Rate, with L = 1280 kHz, M = 2560 kHz, and H = 5120 kHz.)

Table 6-8 - Constellation Gains and Per Channel Power Limits with Two Channels in the Transmit Channel Set

Constellation	Constellation Gain G_{const} Relative to 64 QAM (dB)	P_{min} (dBmV)			P_{max} (dBmV) TDMA	P_{max} (dBmV) S-CDMA	$P_{min} - G_{const}$ (dBmV)	$P_{max} - G_{const}$ (dBmV) TDMA	$P_{max} - G_{const}$ (dBmV) S-CDMA
		L	M	H					
QPSK	-1.18	17	20	23	58	53	18.18	59.18	54.18
8 QAM	-0.21	17	20	23	55	53	17.21	55.21	53.21
16 QAM	-0.21	17	20	23	55	53	17.21	55.21	53.21
32 QAM	0.00	17	20	23	54	53	17.00	54.00	53.00
64 QAM	0.00	17	20	23	54	53	17.00	54.00	53.00
128 QAM	0.05	17	20	23	N/A	53	16.95	N/A	52.95

(P_{min} is a function of Modulation Rate, with L = 1280 kHz, M = 2560 kHz, and H = 5120 kHz.)

Table 6-9 - Constellation Gains and Per Channel Power Limits with Three or Four Channels in the Transmit Channel Set

Constellation	Constellation Gain G_{const} Relative to 64 QAM (dB)	P_{min} (dBmV)			P_{max} (dBmV) TDMA	P_{max} (dBmV) S-CDMA	$P_{min} - G_{const}$ (dBmV)	$P_{max} - G_{const}$ (dBmV) TDMA	$P_{max} - G_{const}$ (dBmV) S-CDMA
		L	M	H					
QPSK	-1.18	17	20	23	55	53	18.18	56.18	54.18
8 QAM	-0.21	17	20	23	52	53	17.21	52.21	53.21
16 QAM	-0.21	17	20	23	52	53	17.21	52.21	53.21
32 QAM	0.00	17	20	23	51	53	17.00	51.00	53.00
64 QAM	0.00	17	20	23	51	53	17.00	51.00	53.00
128 QAM	0.05	17	20	23	N/A	53	16.95	N/A	52.95

(P_{min} is a function of Modulation Rate, with L = 1280 kHz, M = 2560 kHz, and H = 5120 kHz.)

The actual transmitted power per channel, within a burst, MUST 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 for S-CDMA, spreading and varying number of allocated codes.

The CM MUST support the transmit power calculations defined in Section 6.2.19.1 and Section 6.2.19.2.

6.2.19.1 TDMA Transmit Power Calculations

In TDMA mode, the CM determines its target transmit power per channel P_t as follows, for each channel that is active. Define for each active channel, for example, upstream channel X:

P_r = Reported power level (dBmV) of CM in MIB (refers to 64 QAM constellation) for channel X

ΔP = Power level adjustment (dB), for example, as commanded in ranging response message

G_{const} = Constellation gain (dB) relative to 64 QAM constellation (see above table)

P_{min} = Minimum target transmit power per channel permitted for the CM per Section 6.2.19 (see above table)

P_{max} = Maximum target transmit power permitted for the CM per Section 6.2.19 (see above table)

P_{hi} = $\min(P_{\text{max}} - G_{\text{const}})$ over all burst profiles used by the CM in channel X (see above table)

P_{low} = $\max(P_{\text{min}} - G_{\text{const}})$ over all burst profiles used by the CM in channel X (see above table)

P_t = Target transmit power per channel level (dBmV) of CM in channel X (actual transmitted power per channel in channel X as estimated by CM)

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

1. $P_r = P_r + \Delta P$ //Add power level adjustment (for each channel) to reported power level for each channel.
2. $P_r = \min[P_r, P_{\text{hi}}]$ //Clip at max power limit per channel.
3. $P_r = \max[P_r, P_{\text{low}}]$ //Clip at min power limit per channel.
4. Test, IF [$P_r = < P_{\text{low_multi}}$], ignore command and restore previous P_r //Power per channel from this command would violate the set Dynamic Range Window.
5. Test, IF [$P_r = > P_{\text{hi}} - P_{\text{low_min_set}}$], ignore command and restore previous P_r //Power per channel from this command would violate the set Dynamic Range Window.

The CM then transmits, in channel X with target power per channel $P_t = P_r + G_{\text{const}}$, i.e., the reported power, plus the constellation gain.

Usually the reported power level is a relatively constant quantity, while the transmitted power level in channel X varies dynamically as different burst profiles, with different constellation gains, are transmitted. A CM's target transmit power per channel MUST never be below P_{min} or above P_{max} . This implies that in some cases, the extreme transmit power levels (e.g., 61 dBmV for QPSK and 17 dBmV) may not be permitted if burst profiles with multiple constellations are active. Also, if only QPSK is used, the reported power per channel may be greater than 61 dBmV, although the target transmit power per channel will not exceed 61 dBmV.

For example, if in channel X, only QPSK and 64 QAM burst profiles are active, $P_{\text{hi}} = 54$ dBmV and $P_{\text{low}} = 18.2$ dBmV for 1280 kHz modulation rate. P_{low} is dependent on modulation rate. The maximum permitted QPSK transmitted power in channel X is 54 dBmV - 1.2 dB = 52.8 dBmV, the minimum QPSK power in channel X is 18.2 dBmV - 1.2 dB = 17 dBmV (with 1280 kHz modulation rate), the maximum 64 QAM power in channel X is 54 dBmV, and the minimum 64 QAM power in channel X is 18.2 dBmV (with 1280 kHz modulation rate).

6.2.19.2 S-CDMA Transmit Power Calculations

In S-CDMA mode, the power calculations depend on whether the Maximum Scheduled Codes feature is enabled.

6.2.19.2.1 S-CDMA Transmit Power Calculations when Maximum Scheduled Codes Is Not Enabled

In S-CDMA mode when Maximum Scheduled Codes is not enabled, the CM determines its target transmit power per channel P_t , as follows, for each channel which is active. Define for each active channel, for example, upstream channel X:

P_r = reported power level (dBmV) of CM in MIB (refers to 64 QAM constellation and all active codes transmitted) for channel X

P_{hi} = $\min[P_{max} - G_{const}]$ over all burst profiles used by the CM in channel X (see Table 6-7, Table 6-8, and Table 6-9)

P_{low} = $\max[P_{min} - G_{const}] + 10 \log(\text{number_active_codes} / \text{number_of_codes_per_mini_slot})$ where the maximum is over all burst profiles used by the CM in channel X (see Table 6-7, Table 6-8, and Table 6-9)

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

1. $P_r = P_r + \Delta P$ //Add power level adjustment (for each channel) to reported power level for each channel
2. $P_r = \min[P_r, P_{hi}]$ //Clip at max power limit per channel
3. $P_r = \max[P_r, P_{low}]$ //Clip at min power limit per channel
4. Test, IF $[P_r < P_{lo_multi}]$, ignore command and restore previous P_r //Power per channel from this command would violate the set Dynamic Range Window.
5. Test, IF $P_r > P_{hi} - P_{load_min_set}$, ignore command and restore previous P_r //Power per channel from this command would violate the set Dynamic Range Window.

In a spreader-on frame, the CM then transmits each code i with target power:

$$P_{t,i} = P_r + G_{const,i} - 10 \log(\text{number_active_codes})$$

(i.e., the reported power for channel X, plus the constellation gain $G_{const,i}$ of that code, less a factor taking into account the number of active codes. The total transmit power in channel X, P_t , in a frame is the sum of the individual transmit powers $P_{t,i}$ of each code in channel X, where the sum is performed using absolute power quantities [non-dB domain]).

In a spreader-off frame, the CM target transmit power in channel X is $P_t = P_r + G_{const}$.

The transmitted power level in channel X varies dynamically as the number of allocated codes varies, and as different burst profiles, with different constellation gains, are transmitted. A CM's target transmit power per channel MUST never be below P_{min} or above P_{max} , including over all numbers of allocated codes and all burst profiles. This implies that in some cases the extreme transmit power levels per channel (e.g., 17 and 56 dBmV) may not be permitted. Also, if only QPSK is used, the reported power in a channel may be greater than 56 dBmV, although the target transmit power per channel will not exceed 56 dBmV.

If, for example, the Transmit Channel Set contains only one channel, and this channel, channel X, has QPSK and 64 QAM burst profiles, the number of active codes is 128 and the number of codes per mini-slot is 2, then $P_{hi} = 56$ dBmV and $P_{low} = 36.24$ dBmV for 1280 kHz modulation rate. P_{low} is dependent on modulation rate. The maximum

permitted QPSK transmitted power in channel X is $56 - 1.18 = 54.82$ dBmV when all active codes are transmitted. The minimum QPSK power in channel X is 36.24 dBmV $- 1.18$ dB $- 10\log(128)$ dB $+ 10\log(2)$ dB = 17 dBmV (with 1280 kHz modulation rate) when one mini-slot is transmitted. The last term in the sum is the result of summing the individual powers over two codes. Similarly, the maximum 64 QAM power in channel X is 56 dBmV, when all active codes are transmitted and the minimum 64 QAM power in channel X is 36.24 dBmV $- 10\log(128)$ dB $+ 10\log(2)$ dB = 18.18 dBmV (with 1280 kHz modulation rate) when one mini-slot is transmitted. The minimum QPSK power permitted in channel X while transmitting, for example, 2 mini-slots is 20 dBmV (with 1280 kHz modulation rate), and the minimum 64 QAM power permitted in channel X, while transmitting 2 mini-slots, is 21.2 dBmV (with 1280 kHz modulation rate).

The CM needs to implement some form of clipping on the transmitted waveform at the higher output powers in order to prevent peak to average ratio (PAR) issues.

The power received at the CMTS in a spreader-on frame will sometimes be less than the nominal power of a spreader-off frame because of such factors as: 1) broadcast opportunities not used by any CM; 2) unicast grants not used by one or more CMs; or 3) mini-slots assigned to the NULL SID.

6.2.19.2.2 S-CDMA Transmit Power Calculations When Maximum Scheduled Codes is Enabled.

In S-CDMA mode on channels on which Maximum Scheduled Codes is enabled, the CM determines its target transmit power per channel P_t as follows, for each channel which is active. Define for each upstream channel, for example, upstream channel X:

P_r = reported power level (dBmV) of CM in MIB (operational transmit power of the spreader-off ranging burst referenced to 64 QAM modulation) for channel X.

$P_{hi_S} = \min[P_{max} - G_{const}]$ over all spreader-on burst profiles used by the CM in channel X (see Table 6-7, Table 6-8, and Table 6-9)

$P_{low_S} = \max[17 - G_{const}] + 10 \log(\text{number_active_codes} / \text{number_of_codes_per_mini-slot})$ where the maximum is over all burst profiles used by the CM in channel X (see Table 6-7, Table 6-8, and Table 6-9)

P_{max_T} = Maximum target transmit power permitted for the CM in channel X in TDMA mode (see Table 6-7, Table 6-8, and Table 6-9) for the constellation used in ranging.

$P_{hi_T} = \min[P_{max_T} - G_{const}]$ over all spreader-off burst profiles used by the CM in channel X (see Table 6-7, Table 6-8, and Table 6-9).

$P_{on} = P_r$ clipped at the maximum spreader-on limit.

P_{sf} = CM Power Shortfall.

P_{hr} = S-CDMA Power Headroom in dB. Equivalent to the value of the Ranging Response Message TLV-11 divided by 4.

ΔP = power level adjustment in dB sent from CMTS to CM for channel X.

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

1. $P_r = P_r + \Delta P$ //Add power level adjustment (for each channel) to reported power level for each channel.
2. $P_r = \min[P_r, P_{hi_T}]$ //Clip at max TDMA power limit per channel.
3. $P_r = \max[P_r, P_{low_S}]$ //Clip at min S-CDMA power limit per channel.
4. Test, IF[$P_r < P_{low_multi}$], ignore command and restore previous P_r //Power per channel from this command would violate the set Dynamic Range Window.
5. Test, IF[$P_r > P_{hi_T - P_{load_min_set}}$], ignore command and restore previous P_r //Power per channel from this command would violate the set Dynamic Range Window.
6. $P_{on} = \min[P_r, P_{hi_S - P_{load_min_set}}]$ //Clip at max S-CDMA power limit per channel for the multiple channel dynamic range limitation, as set by $P_{load_min_set}$.

In spreader-off frames, the CM transmits in channel off with target power:

$$P_t = P_r + G_{const}$$

Based on the spreader-off transmit power in channel X, the CM updates its power shortfall in channel X according to the following steps:

with only one channel in the Transmit Channel Set

$$P_{sf} = P_t - (56 - P_{load_min_set}) // \text{Difference between spreader-off and max spreader-on target powers in channel X;}$$

$$P_{sf} = \max[P_{sf}, 0] // \text{Set } P_{sf} \text{ to 0 if } P_t \text{ is less than } 56 \text{ dBmV} - P_{load_min_set} \text{ in channel X;}$$

and with more than one channel in the Transmit Channel Set

$$P_{sf} = P_t - (53 - P_{load_min_set}) // \text{Difference between spreader-off and max spreader-on target powers in channel X;}$$

$$P_{sf} = \max[P_{sf}, 0] // \text{Set } P_{sf} \text{ to 0 if } P_t \text{ is less than } 53 \text{ dBmV} - P_{load_min_set} \text{ in channel X.}$$

In spreader-on frames, the CM transmits each code i with target power

$$P_{t,i} = P_{on} + G_{const,i} - 10 \log(\text{number_active_codes}) + P_{hr}$$

i.e., the clipped reported power in channel X plus the constellation gain $G_{const,i}$ of that code, less a factor taking into account the number of active codes, plus the Power Headroom P_{hr} . P_{hr} is the power (in dB) added to account for CMs that have maximum scheduled code limits and can transmit additional power per code. The total transmit power in channel X, P_t , in a frame is the sum of the individual transmit powers $P_{t,i}$ of each code in channel X, where the sum is performed over all N_{alloc} allocated codes using absolute power quantities (non-dB domain).

$$P_t = 10 \log \sum_{i=1}^{N_{alloc}} 10^{P_{t,i}/10}$$

If, for example, with more than one channel in the Transmit Channel Set, and the burst profile in channel X contains QPSK for IUCs 1, 2, 3, and 4 and 64 QAM for IUCs 9 and 10, the number of active codes is 128, and the number of codes per mini-slot is 2, then $P_{hi_S} = 53$ dBmV, $P_{low_S} = 36.24$ dBmV, and $P_{hi_T} = 58$ dBmV. Assume $P_{load_min_set} = 0$

dB. Assume the CM ranges channel X at spreader-off target transmit power of 57 dBmV. The CM reports for channel X $P_{sf} = 57 \text{ dBmV} - 53 \text{ dBmV} = 4 \text{ dB}$. The CMTS uses P_{sf} to set for channel X (using its vendor-specific algorithm) $\text{max_scheduled_codes} = 32$ and $P_{hr} = 6 \text{ dB}$. (The S-CDMA power headroom may differ from the power shortfall, at the discretion of the CMTS.) The CM sets its transmitted power per code in channel X to:

$$\begin{aligned} P_{t,i} &= P_{on} + G_{const,i} - 10 \log(\text{number_active_codes}) + P_{hr} \\ &= 53 \text{ dBmV} + 0 \text{ dB} - 21 \text{ dB} + 6 \text{ dB} \quad // \text{ For a code with 64 QAM modulation in channel X} \\ &= 38 \text{ dBmV} \end{aligned}$$

A parameter that may be used to illustrate the effect of increased power per code in channel X is the Effective Transmit Power, P_{eff} , the power that would result in channel X hypothetically if all N_a active codes were transmitted. It is computed as:

$$\begin{aligned} P_{eff} &= 10 \log \sum_{i=1}^{N_{act}} 10^{P_{t,i}/10} \\ &= P_{on} + P_{hr} + 10 \log \frac{1}{N_{act}} \sum_{i=1}^{N_{act}} 10^{G_{const,i}/10} \end{aligned}$$

where the last term is the average constellation gain.

For a reference case with all codes transmitted in channel X using 64 QAM modulation ($G_{const} = 0 \text{ dB}$), the effective transmit power reduces to:

$$P_{eff} = P_{on} + P_{hr}$$

Continuing the above example, the result is:

$$\begin{aligned} P_{eff} &= 53 \text{ dBmV} + 6 \text{ dB} \\ &= 59 \text{ dBmV} \end{aligned}$$

Limiting the number of codes has given the CM an enhanced effective power of 59 dBmV, which is 6 dB above the normal maximum power per channel of 53 dBmV, and 2 dB above the ranging power of 57 dBmV. In this example, the CMTS used its discretion to ask for 2 dB more enhancement in channel X than was needed ($P_{hr} = 6 \text{ dB}$ vs $P_{sf} = 4 \text{ dB}$), perhaps due to some known impairment in the channel.

The *effective_SNR* is an SNR estimate for a given code corresponding to the effective transmit power in channel X. It is defined as the measured SNR at the last station maintenance, minus the CM power shortfall, plus the power headroom, plus the difference in constellation gain between the ranging burst and the code under consideration. Its equation is:

$$\text{effective_SNR} = \text{measured_SNR} - P_{sf} + P_{hr} + (G_{const,i} - G_{const,ranging})$$

where $G_{const,ranging}$ is the constellation gain of the ranging burst in channel X that resulted in the SNR measurement.

In the MIB, *effective_SNR* corresponds to a reference case with 64 QAM modulation ($G_{const,i} = 0 \text{ dB}$):

$$\text{effective_SNR} = \text{measured_SNR} - P_{sf} + P_{hr} - G_{const,ranging}$$

Continuing the example, if the measured SNR in channel X in the last station maintenance was 17 dB, using QPSK modulation ($G_{\text{const, ranging}} = -1.2$ dB), then the effective SNR referenced to 64 QAM modulation is:

$$\text{effective_SNR} = 17\text{dB} - 4\text{dB} + 6\text{dB} + 1.2 \text{ dB} = 20.2 \text{ dB}$$

6.2.19.3 Transmit Power Step Size

The step resolution in transmit power for each channel MUST be 0.5 dB or less. The step size accuracy MUST be within 0.4 dB. For example, the actual power increase in a channel X resulting from a command to increase the power level by 1 dB in a CM's next transmitted burst in channel X MUST be between 0.6 dB and 1.4 dB.

All active channels with no power change commanded MUST NOT change by more than 0.1 dB during a burst even with a power change command to any subset of the other active channels.

After a change in the Dynamic Range Window (i.e., $P_{\text{load_min_set}}$), the transmit power MUST NOT change by more than 0.4 dB in any of the active channels when they resume transmission (referenced against any commanded change in their transmit power).

A relaxation in step size accuracy to 1.4 dB is allowed for one gain change (and one setting of Dynamic Range Window) for each 12 dB, when changing the power control range in either direction (from low-end to high-end power and vice versa). The locations of these gain changes (increasing and decreasing to provide hysteresis), MUST be at least 2 dB apart, thus enabling the use of large step attenuators in the coverage of the full power control range.

6.2.19.4 Transmit Power Requirements for Multiple Transmit Channel mode Not Enabled

With Multiple Transmit Channel mode disabled, the CM MUST operate as in [DOCSIS RFI 2.0], with the exceptions that the minimum upstream power P_{min} MUST be as given in Table 6-10 below, and the maximum upstream power P_{max} MUST be as given in Table 6-7 above. Note that in Section 6.2.18.2.2 of the [DOCSIS RFI 2.0] the value of 53 in the equation for $P_{\text{hi_s}}$ MUST be replaced with 56 since this value is actually intended to be P_{max} from Table 6-7 above.

Table 6-10 – Minimum Transmit Power P_{min} for Multiple Transmit Channel mode Not Enabled

Modulation Rate (kHz)	P_{min} (dBmV)	Applicability
160	17	CM MAY support
320	17	CM MAY support
640	17	CM MAY support
1280	17	CM MUST support
2560	20	CM MUST support
5120	23	CM MUST support

6.2.20 Burst Profiles

The transmission characteristics are separated into three portions: a) Channel Parameters; b) Burst Profile Attributes; and c) User Unique Parameters. The Channel Parameters include: a) the modulation rate (six rates from 160 kHz to 5120 kHz in octave steps); b) the center frequency (Hz); c) the 1536-bit Preamble Superstring; and d) the S-CDMA channel parameters. The Channel Parameters are shared by all users on a given channel. The Burst Profile Attributes are listed in Table 6-11, and these parameters are the shared attributes corresponding to a burst type.

The CM MUST generate each burst at the appropriate time, as conveyed in the mini-slot grants provided by the CMTS MAPs.

The CM MUST support all burst profiles commanded by the CMTS, via the Burst Descriptors in the UCD, and subsequently assigned for transmission in a MAP.

Table 6-11 - Burst Profile Attributes

Burst Profile Attributes	Configuration Settings
Modulation	QPSK, 8 QAM, 16 QAM, 32 QAM, 64 QAM, 128 QAM(TCM Only)
Differential Encoding	On/Off
TCM Encoding	On/Off
Preamble Length	0-1536 bits (Note Section 6.2.9)
Preamble Value offset	0 to 1534
R-S FEC Error Correction (T)	0 to 16 (0 implies no R-S FEC. The number of codeword parity bytes is $2*T$)
R-S FEC Codeword Information Bytes (k)	Fixed: 16 to 253 (assuming R-S FEC on) Shortened: 16 to 253 (assuming R-S FEC on)
Scrambler Seed	15 bits
Maximum Burst Length (mini-slots) ¹	0 to 255^2
Guard Time	4 to 255 modulation intervals There is no guard time in S-CDMA
Last Codeword Length	Fixed, shortened
Scrambler On/Off	On/Off
Byte Interleaver Depth (I_r) ³	0 to floor ($2048/N_r$) ⁴
Byte Interleaver Block Size (B_r) ⁵	$2*N_r$ to 2048
Preamble Type	QPSK0/QPSK1
S-CDMA Spreader ⁶	On/Off
S-CDMA Codes per Subframe ⁶	1 to 128
S-CDMA Interleaver Step ⁶	1 to (spreading intervals per frame - 1)
S-CDMA Selection Mode for Active Codes and Code Hopping	Mode 1 or Mode 2
S-CDMA Selection String for Active Codes	128-bit string indicating which codes are active in Selectable Active Codes Mode 2
<p>¹. A burst length of 0 mini-slots in the Channel Profile means that the burst length is variable on that channel for that burst type. The burst length, while not fixed, is granted explicitly by the CMTS to the CM in the MAP.</p> <p>². N_r is the R-S codeword size $k+2T$ as defined in Section 6.2.5.1.</p> <p>³. If depth=1, no interleaving; if depth=0, dynamic mode.</p> <p>⁴. N_r is the R-S code word size $K+2T$ as defined in Section 6.2.5.1.</p> <p>⁵. Used only in dynamic mode</p> <p>⁶. Used only for S-CDMA channels.</p>	

The User Unique Parameters may vary for each user even when using the same burst type on the same channel as another user (for example, Power Level), and are listed in Table 6-12:

Table 6-12 - User Unique Burst Parameters

User Unique Parameter	Adjustment Command	Resulting Parameter Value
Power Level ¹	8-bit two's complement, resolution = 0.25 dB	TDMA: +17 to +57 dBmV (32 QAM, 64 QAM) +17 to +58 dBmV (8 QAM, 16 QAM) +17 to +61 dBmV (QPSK) S-CDMA: +17 to +56 dBmV (all modulations) Resolution = 0.25 dB or better
Offset Frequency	Range = ±32 kHz, resolution = 1 Hz	Frequency Range per Section 6.2.17.1 Upstream Frequency Agility and Range
Ranging Offset	Integer part: 32-bit two's complement, resolution = (1 / 10.24 MHz) = 6.25 μs/64 = 97.65625 ns Fractional part: unsigned 8-bit fractional extension, resolution = 6.25 μs/(64*256) = 0.3814697265625 ns	Range: sufficient for maximum cable plant length per Section 1.2.1 Resolution: TDMA: 6.25μs/64. S-CDMA: 6.25 μs/(64*256)
Burst Length (mini-slots) if variable on this channel (changes burst-to-burst)	N/A	1 to 255 mini-slots
Transmit Equalizer Coefficients (See Section 6.2.16, Transmit Pre-Equalizer)	DOCSIS 3.0 & 2.0 modes: 24 complex coefficients, 4 bytes per coefficient (2 real and 2 imaginary), load and convolve modes DOCSIS 1.1 mode: up to 64 complex coefficients, 4 bytes per coefficient (2 real and 2 imaginary), convolve mode only	DOCSIS 3.0 & 2.0 modes: 24 complex coefficients DOCSIS 1.1 mode: up to 64 complex coefficients
¹ . The minimum power level limit depends on modulation rate. The maximum power level limit depends on modulation order, whether or not Multiple Transmit Channel mode is enabled and the number of upstream channels in the Transmit Channel Set if MTC Mode is enabled.		

The CM MUST implement the Offset Frequency Adjustment to effect a change in upstream carrier frequency within ±10 Hz of the commanded change.

6.2.20.1 Ranging Offset

Ranging Offset is the time difference between the CM upstream frame time base and the CMTS upstream frame time base. It is an advancement equal to roughly the round-trip delay between the CM and the CMTS, and is needed to synchronize upstream transmissions in the TDMA and S-CDMA schemes. The CMTS MUST provide the CM with feedback adjustments of this offset, based on reception of one or more successfully received bursts (i.e.,

satisfactory result from each technique employed: error correction and/or CRC). The CMTS sends these Timing Adjust commands to the CM in the Ranging Response MAC message, where a negative value implies the Ranging Offset is to be decreased, resulting in later times of transmission at the CM.

For TDMA channels the CM MUST implement the Timing Adjust command with resolution of at most 1 symbol duration (of the symbol rate in use for a given burst), and (other than a fixed bias) with accuracy within $\pm 0.25 \mu\text{s}$ plus $\pm 1/2$ symbol owing to resolution. As an example, for the maximum symbol rate of 5.12 Msym/s, the corresponding symbol period would be 195 ns, the corresponding maximum resolution for the Timing Adjust MUST be 195 ns. For the maximum symbol rate of 5.12 Msym/s, the corresponding minimum accuracy MUST be ± 348 ns. The accuracy of CM burst timing of $\pm 0.25 \mu\text{s}$ plus $\pm 1/2$ symbol is relative to the mini-slot boundaries derivable at the CM, based on an ideal processing of the timestamp signals received from the CMTS.

The resolution of the integer part of the Timing Adjust parameter, which is used for TDMA channels, is $(1 / 10.24 \text{ MHz}) = 6.25 \mu\text{s}/64 \approx 97.66$ ns. For S-CDMA channels, the CMTS provides an additional fractional field in the Timing Adjust command, with resolution of $1/16384$ of the frame tick increment = $6.25 \mu\text{s}/(64*256) \approx 0.3814$ ns. For S-CDMA channels, the CM MUST implement the Timing Adjust to within ± 0.01 of the nominal chip period. As an example, for the maximum chip rate of 5.12 MHz, the corresponding maximum resolution for implementation of the timing correction would be $(\pm 0.01)*195$ ns or roughly ± 2 ns.

For CMTS/CM systems using multiple downstream channels, the CMTS MUST designate a single channel as the primary channel, with the timing information such as timestamps or timestamp snapshots.

For CMTS/CM systems using multiple upstream channels, the CM must range each upstream channel separately, using the resolutions described above.

6.2.20.2 TDMA Reconfiguration Times

The CM MUST be capable of switching burst profiles with no reconfiguration time required between bursts, except for changes in the following parameters: 1) Output Power; 2) Symbol Rate; 3) Offset Frequency; 4) Channel Frequency; 5) Ranging Offset; and 6) Pre-Equalizer Coefficients. In the case of Multiple Transmit Channel mode enabled, all of these parameters are adjusted independently on each upstream channel and reconfiguration time on the non-adjusted channels is not required by the CM.

In the case of Multiple Transmit Channel mode enabled, global reconfiguration time (simultaneous on all of the active channels) could be required by the CM when the Dynamic Range Window is adjusted (new value of $P_{\text{load_min_set}}$), or when any active channel accumulates a transmit power change of more than 3 dB since the last global reconfiguration time (see Section 6.2.22.1). The CM MUST meet the requirements, on each active channel, for "Dynamic Range Window change" (below) whenever $P_{\text{load_min_set}}$ is changed or re-commanded.

For **Dynamic Range Window** change or re-command: If $P_{\text{load_min_set}}$ is changed, or re-commanded, the CM MUST be able to implement the change between bursts, as long as the CMTS allocates at least 96 symbols, plus $10 \mu\text{s}$, between the last symbol center of one burst (on any active channel) and the first symbol center of the following burst (on any active channel -- same channel or different channel from the previous burst). (Global "quiet" across all active channels requires the intersection of ungranted burst intervals across all active TDMA channels to be at least 96 symbols plus $10 \mu\text{s}$ duration.) The Output Power of the CM MUST be settled to within ± 0.1 dB of its final output power level within $10 \mu\text{s}$ from the beginning of the change in $P_{\text{load_min_set}}$, and from the beginning of any Output Power change on any channel that accompanies the change or re-command of $P_{\text{load_min_set}}$. The Dynamic Range Window, i.e., $P_{\text{load_min_set}}$, MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS. The Dynamic Range Window MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted.

For **Output Power** changes: If Output Power is to be changed by 1 dB or less, the CM MUST be able to implement the change between bursts as long as the CMTS allocates at least 96 symbols, plus $5 \mu\text{s}$ between the last symbol center of one burst and the first symbol center of the following burst. If Output Power is to be changed by more than

1 dB, the CM MUST be able to implement the change between bursts, as long as the CMTS allocates at least 96 symbols, plus 10 μ s, between the last symbol center of one burst and the first symbol center of the following burst. The maximum reconfiguration time of 96 symbols should compensate for the ramp-down time of one burst and the ramp-up time of the next burst, as well as the overall transmitter delay time, including the pipeline delay and pre-equalizer delay. The Output Power of the CM MUST be settled to within ± 0.1 dB of its final output power level: a) within 5 μ s from the beginning of a change of 1 dB or less; and b) within 10 μ s from the beginning of a change of greater than 1 dB. Output Power MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS. Output Power MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted.

For **Symbol Rate** changes: The CM MUST be able to transmit consecutive bursts as long as the CMTS allows the required time between bursts for UCD parameter changes. Symbol Rate MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS. Symbol Rate MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted.

For **Offset Frequency** changes: The CM MUST be able to transmit consecutive bursts as long as the CMTS allocates at least 96 symbols in between the last symbol center of one burst and the first symbol center of the following burst. The maximum reconfiguration time of 96 symbols should compensate for the ramp-down time of one burst and the ramp-up time of the next burst, as well as the overall transmitter delay time, including the pipeline delay and optional pre-equalizer delay. Offset frequency MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS. Offset frequency MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted.

For **Channel Frequency** changes: The CM MUST be able to implement the change between bursts as long as the CMTS allocates at least 96 symbols plus 100 ms between the last symbol center of one burst and the first symbol of the following burst. The Channel Frequency of the CM MUST be settled within the phase noise and accuracy requirements of Section 6.2.22.5 and Section 6.2.22.6 within 100 ms from the beginning of the change. Channel Frequency MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS. Channel Frequency MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted.

For **Ranging Offset** and/or **Pre-Equalizer Coefficient** changes: The CM MUST be able to transmit consecutive bursts as long as the CMTS allocates at least 96 symbols in between the last symbol center of one burst and the first symbol center of the following burst. The maximum reconfiguration time of 96 symbols should compensate for the ramp-down time of one burst and the ramp-up time of the next burst, as well as the overall transmitter delay time, including the pipeline delay and pre-equalizer delay. Ranging Offset MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS. Ranging Offset MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted.

For **Modulation Type** changes: The CM MUST be able to transmit consecutive bursts with no reconfiguration time between them (except for the minimum guard time). The modulation MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted, EXCLUDING the effect of the transmit equalizer (if present in the CM). (This is to be verified with the transmit equalizer providing no filtering; delay only. Note that if the CMTS has decision feedback in its equalizer, it may need to provide more than the 96 symbol gap between bursts of different modulation type which the same CM may use; this is a CMTS decision.)

6.2.20.3 S-CDMA Reconfiguration Times

In S-CDMA mode, for changes in Output Power per mini-slot, Offset Frequency, Pre-equalizer coefficients, and/or Ranging Offset, the CM MUST be able to transmit consecutive bursts as long as the CMTS allocates the time duration of at least one frame in between the bursts. For all other burst profile parameter changes, no

reconfiguration is required beyond what is provided by the MAC for such changes. When Multiple Transmit Channel mode is enabled, all of these parameters are adjusted independently on each upstream channel and reconfiguration time on the non-adjusted channels is not required by the CM.

In the case of Multiple Transmit Channel mode enabled, global reconfiguration time could be required by the CM when the Dynamic Range Window is adjusted (new value of $P_{load_min_set}$), or when any active channel accumulates a transmit power change of more than 3 dB since the last global reconfiguration time (see Section 6.2.22.1). (Global "quiet" across all active channels requires the intersection of ungranted burst intervals across all active S-CDMA channels to be at least one frame.) The CM MUST be able to transmit consecutive bursts, even with a change or re-command of $P_{load_min_set}$, as long as the CMTS allocates at least one frame in between bursts, across all channels in the Transmit Channel Set. (From the end of a burst on one channel to the beginning of the next burst on any channel, there must be at least one frame duration to provide a "global reconfiguration time" for S-CDMA channels.)

6.2.20.4 CM Timing Offsets When Changing Modulation Rate

When making a modulation rate change, the CM MUST employ the timing offsets shown in Table 6-13. The offsets in the table correspond to the contribution of DOCSIS 1.0 and 1.1 legacy upstream receivers to changes in latency when making modulation rate changes. The timing offset to apply is the difference between the entry in the table corresponding to the new modulation rate and the entry corresponding to the original modulation rate. The offsets are referenced to the center of the first symbol in the burst, which is the reference point for burst timing, as stated in Section 6.2.21. Specification of these offsets is needed so that CMs apply uniform adjustments to their ranging offsets, and CMTSs can appropriately handle CMs that apply these offsets when making modulation rate changes.

Table 6-13 - Timing Offset for Modulation Rate Changes

Modulation Rate (kHz)	Timing Offset (in units of 1/64 time ticks referenced to 5.12 MHz)	Applicability
5120	0	CM MUST support
2560	0	CM MUST support
1280	24	CM MUST support
640	72	CM MAY support
320	168	CM MAY support
160	360	CM MAY support

For example, suppose a CM is on an upstream channel, operating at a modulation rate of 1.28 MHz. Now, suppose the UCD message from the CMTS changes the modulation rate of the channel to 0.32 MHz. The CM applies an additional timing offset of $168 - 24 = 144$ to its ranging offset to compensate for this modulation rate change. The value of 144 is positive, and thus, the CM will add to its ranging offset, so that it effectively transmits earlier by 144 units of 1/64 time ticks.

Furthermore, in changing modulation rates, if a CM has its own contribution to a change in latency, the CM MUST also compensate for this CM-specific latency difference. This is in addition to the offset applied from the values in the table above, which result from legacy CMTS upstream receiver contributions to changes in latency. The requirements for CM burst timing accuracy found earlier in this section for TDMA mode, referenced to the modulation rate that is the lower of the original, and the new modulation rate, apply after the modulation rate change, with the required timing offsets above considered. Specifically, the CM MUST implement the timing adjustments with accuracy within $\pm 0.25 \mu\text{s}$, plus $\pm 1/2$ symbol, in both TDMA and S-CDMA modes.

A CMTS that does not apply the same internal physical delay offsets as the legacy DOCSIS upstream CMTS receiver implementation is capable of receiving a CM burst after a modulation rate change in any of the following ways, but is not limited to only these ways:

1. The CMTS may implement the internal physical delay offset, as specified in the above table.
2. The CMTS may implement an internal timing compensation based on the expected offset in the above table.
3. The CMTS may increase the guard time.
4. The CMTS may send an unsolicited RNG-RSP to each CM to adjust the delay offset. As discussed in [MULPI], Section 6.4.4, the CM is expected to be capable of adjusting its timing offset at any time with the accuracy specified within this section.

6.2.21 Burst Timing Convention

Figure 6-31 illustrates the nominal burst timing for TDMA channels.




Figure 6-31 - Nominal TDMA Burst Timing

Figure 6-32 indicates worst-case burst timing for a TDMA channel. In this example, burst N arrives 1.5 symbols late, and burst N+1 arrives 1.5 symbols early, but separation of 5 symbols is maintained; the 8-symbol guard band is shown below:




Figure 6-32 - Worst-Case TDMA Burst Timing

At a symbol rate of R_s , symbols occur at a rate of one each $T_s = 1/R_s$ s. Ramp-Up and Ramp-Down are the spread of a symbol in the time domain beyond T_s duration, owing to the symbol-shaping filter, and any residual effect from the transmit equalizer. If only one symbol were transmitted, its duration would be longer than T_s due to the shaping filter impulse response being longer than T_s . The spread of the first and last symbols of a burst transmission effectively extends the duration of the burst to longer than $N * T_s$, where N is the number of symbols in the burst.

For S-CDMA channels, the bursts from all CMs are synchronized. This means that the ramp-down of one burst may occur at the same time as the ramp-up of the subsequent burst. The CM MUST meet the ranging and synchronization requirements of S-CDMA to assure that the ramp-down and ramp-up of bursts are aligned.

6.2.22 Fidelity Requirements

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

6.2.22.1 Spurious Emissions

The noise and spurious power MUST NOT exceed the levels given in Table 6-14, Table 6-15, and Table 6-17. The noise and spurious power SHOULD NOT exceed the levels given in Table 6-16 and Table 6-18. When Multiple Transmit Channel mode is enabled, the noise and spurious power requirements for the individual channels are summed (absolute power, NOT in dB) to determine the composite noise floor for the multi-channel transmission condition. Also 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 Transmit Channel Set, and for the ratio of *number_allocated_codes / number_active_codes* from 1 down to $1/4^{\text{th}}$, for S-CDMA channels.

For S-CDMA mode, when a modem is transmitting fewer than $1/4^{\text{th}}$ of the active spreading codes on a channel, the spurious emissions requirement limit is the power value (in dBmV), corresponding to the specifications for the

power level associated with $1/4^{\text{th}}$ of the active spreading codes (i.e., 6 dB below the channel power corresponding to all active codes for that channel allocated to the CM).

When Multiple Transmit Channel mode is enabled and there are two or more channels in the Transmit Channel Set, the spurious performance requirements MUST be met only when the following loading ranges are satisfied:

Two or more channels in the TCS: applicable when $P_{\text{load}_1} - P_{\text{load}_{\text{min_set}}} \leq 3$ dB, AND

Two channels in the TCS: applicable when $P_{\text{load}_2} - P_{\text{load}_1} \leq 12$ dB.

Three channels in the TCS: applicable when $P_{\text{load}_3} - P_{\text{load}_1} \leq 8$ dB.

Four channels in the TCS: applicable when:

$$P_{\text{load}_2} - P_{\text{load}_1} \leq 4 \text{ dB.}$$

AND $P_{\text{load}_3} - P_{\text{load}_2} \leq 4$ dB.

AND $P_{\text{load}_4} - P_{\text{load}_3} \leq 4$ dB.

When a modem is transmitting at power levels falling outside these loading ranges (which is allowed by the 12 dB dynamic range), the spurious emissions requirement limits are the power values (in dBmV rather than dBc) corresponding to the specifications associated with transmissions at the second and higher loaded channels if the loading of those channels were decreased (transmit power increased) to just meet the applicability ranges.

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. "Global reconfiguration time" is defined as the inactive time interval provided between active transmissions, which simultaneously satisfies the requirement in Section 6.2.20.2 for all TDMA channels in the TCS and the requirement in Section 6.2.20.3 for all S-CDMA channels in the TCS. With Multiple Transmit Channel mode enabled, the CMTS SHOULD provide global reconfiguration time to a CM before the CM has been commanded to change any upstream channel transmit power by ± 3 dB cumulative since its last global reconfiguration time.

As one example, with three S-CDMA channels in the TCS, at transmit powers per channel of 53 dBmV, 41 dBmV, and 41 dBmV, the absolute spurious emissions requirements, when two or more channels in the TCS are transmitting, correspond to the absolute spurious emissions requirements with transmit power per channel of 53 dBmV, 45 dBmV, and 45 dBmV. The 41 dBmV power levels are raised for the purposes of determining the spurious emissions limits, up to $53 \text{ dBmV} - 8 \text{ dB} = 45 \text{ dBmV}$.

As an example, with four S-CDMA channels in the TCS, at transmit powers per channel of 53 dBmV, 41 dBmV, 41 dBmV, and 41 dBmV, the absolute spurious emissions requirements, when two or more channels in the TCS are transmitting, correspond to the absolute spurious emissions requirements with transmit power per channel of 53 dBmV, 49 dBmV, 45 dBmV, and 41 dBmV. The 41 dBmV power levels are raised for the purposes of determining the spurious emissions limits, up to $53 \text{ dBmV} - 4 \text{ dB} = 49 \text{ dBmV}$; and the artificial $49 \text{ dBmV} - 4 \text{ dB} = 45 \text{ dBmV}$; and the artificial $45 \text{ dBmV} - 4 \text{ dB} = 41 \text{ dBmV}$.

In Table 6-14, 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 is equal to the modulation rate (e.g., 1280 to 5120 kHz). All requirements expressed in dBc are relative to the actual transmit power that the CM emits in one channel.

The measurement bandwidth is 160 kHz for the Between Bursts (none of the channels in the TCS are bursting) specs of Table 6-14, except where called out as 4 MHz or 250 kHz.

The Transmitting Burst specs apply during the mini-slots granted to the CM (when the CM uses all or a portion of the grant), and for 32 modulation intervals before and after the granted mini-slots. The Between Bursts specs apply except during a used grant of mini-slots on any active channel for the CM, and the 32 modulation intervals before and after the used grant.

In TDMA mode, a mini-slot may be as short as 32 modulation intervals, or 6.25 μ s at the 5.12 Msym/s rate, or as short as 25 μ s at the 1.28 Msym/s rate.

Table 6-14 - Spurious Emissions

Parameter	Transmitting Burst	Between Bursts
Inband	-40 dBc	-72 dBc
Adjacent Band	See Table 6-15 and Table 6-16	-72 dBc
Within the upstream operating range 5-42 MHz or 5-85 MHz (excluding assigned channel, adjacent channels)	See Table 6-17 and Table 6-18	-72 dBc
For the case where the upstream operating range is 5-42 MHz: CM Integrated Spurious Emissions Limits (all in 4 MHz, includes discretely) ¹ 42 to 54 MHz 54 to 60 MHz 60 to 88 MHz 88 to 870 MHz	max(-40 dBc, -26 dBmV) -35 dBmV -40 dBmV -45 dBmV	-26 dBmV -40 dBmV -40 dBmV max(-45 dBmV, -40 dB ref d/s ²)
For the case where the upstream operating range is 5-42 MHz: CM Discrete Spurious Emissions Limits ¹ 42 to 54 MHz 54 to 88 MHz 88 to 870 MHz	-max(-50 dBc, -36 dBmV) -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 discretely) ³ 108 to 136 MHz 136 to 870 MHz	-40 dBmV -45 dBmV	-40 dBmV max(-45 dBmV, -40 dB ref d/s ⁴)
For the case where the upstream operating range is 5-85 MHz: CM Discrete Spurious Emissions Limits ¹ 108 to 870 MHz	-50 dBmV	-50 dBmV

1. These spec limits exclude a single discrete spur related to the tuned received channel; this single discrete spur MUST be no greater than -40 dBmV.
2. "dB ref d/s" is relative to the received downstream signal level. Some spurious outputs are proportional to the receive signal level.
3. These spec limits exclude a single discrete spur related to the tuned received channel; this spur MUST be no greater than -40 dBmV.
4. "dB ref d/s" is relative to the received downstream signal level.

6.2.22.1.1 Adjacent Channel Spurious Emissions

Spurious emissions from a transmitted carrier may occur in an adjacent channel which could be occupied by a carrier of the same or different modulation rate. Table 6-15 lists the required adjacent channel spurious emission levels for all combinations of transmitted carrier modulation rates and adjacent channel modulation rates. The measurement is performed in an adjacent channel interval of appropriate bandwidth and distance from the transmitted carrier, based on the modulation rates of the transmitted carrier, and the carrier in the adjacent channel.

Table 6-15 - Adjacent Channel Spurious Emissions Requirements Relative to the Per Channel Transmitted Burst Power Level for Each Channel

Specification in the interval	Measurement interval and distance from carrier edge	Adjacent channel carrier modulation rate
-50 dBc	20 kHz to 180 kHz	160 kHz
-50 dBc	40 kHz to 360 kHz	320 kHz
-50 dBc	80 kHz to 720 kHz	640 kHz
-50 dBc	160 kHz to 1440 kHz	1280 kHz
-47 dBc	320 kHz to 2880 kHz	2560 kHz
-44 dBc	640 kHz to 5760 kHz	5120 kHz

In addition, the CM SHOULD meet the requirements in Table 6-16, under the same conditions as listed for Table 6-15 above. With a lower power channel bursting as a lone bursting channel (at that point in time), 2 dB relaxation will apply to the absolute spurious emissions requirement of Table 6-16.

Table 6-16 - Adjacent Channel Spurious Emissions Goals Relative to the Per Channel Transmitted Burst Power Level for Each Channel

Specification in the interval	Measurement interval and distance from carrier edge	Adjacent channel carrier modulation rate
-53 dBc	20 kHz to 180 kHz	160 kHz
-53 dBc	40 kHz to 360 kHz	320 kHz
-52 dBc	80 kHz to 720 kHz	640 kHz
-51 dBc	160 kHz to 1440 kHz	1280 kHz
-50 dBc	320 kHz to 2880 kHz	2560 kHz
-48 dBc	640 kHz to 5760 kHz	5120 kHz

6.2.22.1.2 Spurious Emissions in the Upstream Frequency Range

Table 6-17 lists the possible modulation rates that could be transmitted in an interval, the required spurious level in that interval, and the initial measurement interval at which to start measuring the spurious emissions. Measurements

should start at the initial distance and be repeated at increasing distance from the carrier until the upstream band edge is reached. The lower band edge of the upstream is 5 MHz; the upper band edge of the upstream is 42 MHz or if upstream frequency range selection is available in the CM, 85 MHz may be selected.

Table 6-17 - Spurious Emissions Requirements in the Upstream Frequency Range Relative to the Per Channel Transmitted Burst Power Level for Each Channel

Possible modulation rate in this interval	Specification in the interval	Initial measurement interval and distance from carrier edge
160 kHz	-54 dBc	220 kHz to 380 kHz
320 kHz	-52 dBc	240 kHz to 560 kHz
640 kHz	-50 dBc	280 kHz to 920 kHz
1280 kHz	-50 dBc	360 kHz to 1640 kHz
2560 kHz	-47 dBc	520 kHz to 3080 kHz
5120 kHz	-44 dBc	840 kHz to 5960 kHz

In addition, the CM SHOULD meet the requirements in Table 6-18, under the same conditions as listed for Table 6-17 above. With a lower power channel bursting as a lone bursting channel (at that point in time), 2 dB relaxation will apply to the absolute spurious emissions requirement of Table 6-18.

Table 6-18 - Spurious Emissions Goals in the Upstream Frequency Range Relative to the Per Channel Transmitted Burst Power Level for Each Channel

Possible modulation rate in this interval	Specification in the interval	Initial measurement interval and distance from carrier edge
160 kHz	-60 dBc	220 kHz to 380 kHz
320 kHz	-58 dBc	240 kHz to 560 kHz
640 kHz	-56 dBc	280 kHz to 920 kHz
1280 kHz	-54 dBc	360 kHz to 1640 kHz
2560 kHz	-52 dBc	520 kHz to 3080 kHz
5120 kHz	-50 dBc	840 kHz to 5960 kHz

6.2.22.2 Spurious Emissions During Burst On/Off Transients

Each transmitter MUST control spurious emissions, prior to and during ramp-up, during and following ramp-down, and before and after a burst.

On/off spurious emissions, such as the change in voltage at the upstream transmitter output, due to enabling or disabling transmission, MUST be no more than 100 mV, and such a step MUST be dissipated no faster than 2 μ s 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 maximum change in voltage MUST 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 7 mV at 31 dBmV per channel and below. This requirement does not apply to CM power-on and power-off transients.

6.2.22.3 Modulation Error Ratio (MER)

MER measures the cluster variance caused by the transmit waveform. It includes the effects of ISI, spurious, phase noise, and all other transmitter degradations.

6.2.22.3.1 Definitions

Symbol MER: MER_{symp} is defined as follows for TDMA or S-CDMA symbols. The transmitted RF waveform (after appropriate down conversion) is applied to the ideal receive symbol matched filter and is sampled once per symbol. For TDMA, the matched filter is a square-root raised cosine filter with $\alpha = 0.25$. For S-CDMA, the matched filter is a square-root raised cosine filter with $\alpha = 0.25$, convolved with the time-reversed spreading code sequence. (In this convolution, the spreading code sequence is expressed as a weighted impulse train spaced at the chip period.) No external noise (AWGN) is added to the signal. The carrier frequency offset, carrier phase offset, symbol timing and gain may be adjusted during each burst to maximize MER_{symp} . Equalization of the received waveform is not permitted. For cases where the CM transmit equalizer is ON, the transmit equalizer coefficients may be adjusted to maximize MER_{symp} . MER_{symp} is defined at the F connector of the CM, except that when an echo channel is inserted, MER_{symp} is defined at the output of the echo channel. MER_{symp} is computed by the formula:

$$MER_{\text{symp}} \text{ (dB)} = 10 \cdot \log_{10} \left\{ \frac{E_{\text{av}}}{\frac{1}{N} \sum_{j=1}^N |e_j|^2} \right\}$$

where:

- E_{av} is the average constellation energy for equally likely symbols (see Section 6.2.14 and Figure 6-20)
- N is the number of symbols averaged
- e_j is the error vector from the j th received symbol to the ideal transmitted QAM symbol on the grid of Figure 6-20

For S-CDMA, MER_{symp} is averaged over all active codes.

MER of composite chips: MER_{chip} is specified for composite S-CDMA chips to ensure that high SNR is maintained, especially for a small number of allocated codes, to prevent noise funneling effects when many modems transmit simultaneously. A composite S-CDMA chip is defined as the output of the spreader during one chip interval, that is, an element of the transmission vector defined in Section 6.2.15, "S-CDMA Spreader".

MER_{chip} is defined as follows: The transmitted RF waveform (after appropriate down conversion) is applied to the ideal receive chip matched filter and is sampled once per chip. The matched filter is a square-root raised cosine filter with $\alpha = 0.25$. No external noise (AWGN) is added to the signal. The carrier frequency offset, carrier phase offset, timing and gain may be adjusted during each burst to maximize MER_{chip} . Equalization of the received waveform is not permitted. For cases where the CM transmit equalizer is ON, the transmit equalizer coefficients may be adjusted to maximize MER_{chip} . MER_{chip} is defined at the F connector of the CM. MER_{chip} is computed by the formula where:

$$MER_{chip}(dB) = 10 \bullet \log_{10} \left\{ \frac{\sum_{j=1}^N |p_j|^2}{\sum_{j=1}^N |p_j - r_j|^2} \right\}$$

where:

- p_j is the j th ideal transmitted composite chip
- r_j is the j th received composite chip
- N is the number of composite chips observed

6.2.22.3.2 Requirements

Unless otherwise stated, the MER MUST meet or exceed the following limits over the full transmit power range of Table 6-12 for each modulation, modulation rate, and over the full carrier frequency range, and for S-CDMA, over any valid number of active and allocated codes. The 5-85 MHz carrier frequency range refers more precisely to $[5 \text{ MHz} + \text{modulation rate} * 1.25 / 2]$ to $[85 \text{ MHz} - \text{modulation rate} * 1.25 / 2]$. At the break points between regions, the higher MER specification applies.

Case 1: Flat channel, transmit equalization OFF

Case 1a: for modulation rates 2.56 MHz and below for 5 MHz to 42 MHz mode of operation:

- $MER_{symp} \geq 30 \text{ dB}$ over 15 to 30 MHz carrier frequency
- $MER_{symp} \geq 27 \text{ dB}$ over 10 MHz to 15 MHz and 30 MHz to 35 MHz carrier frequency
- $MER_{symp} \geq 26 \text{ dB}$ over 5 MHz to 10 MHz and 35 MHz to 42 MHz carrier frequency

Case 1b: for modulation rate 5.12 MHz for 5 MHz to 42 MHz mode of operation:

- $MER_{symp} \geq 27 \text{ dB}$ over 15 to 30 MHz carrier frequency
- $MER_{symp} \geq 24 \text{ dB}$ over 10 MHz to 15 MHz and 30 MHz to 35 MHz carrier frequency
- $MER_{symp} \geq 23 \text{ dB}$ over 5 MHz to 10 MHz and 35 MHz to 42 MHz carrier frequency

Case 1c: for modulation rates 2.56 MHz and below for 5 MHz to 85 MHz mode of operation:

- $MER_{symp} \geq 30 \text{ dB}$ over 15 to 61 MHz carrier frequency
- $MER_{symp} \geq 27 \text{ dB}$ over 10 MHz to 15 MHz and 61 MHz to 71 MHz carrier frequency
- $MER_{symp} \geq 26 \text{ dB}$ over 5 MHz to 10 MHz and 71 MHz to 85 MHz carrier frequency

Case 1d: for modulation rate 5.12 MHz for 5 MHz to 85 MHz mode of operation:

- $MER_{symp} \geq 27 \text{ dB}$ over 15 to 61 MHz carrier frequency
- $MER_{symp} \geq 24 \text{ dB}$ over 10 MHz to 15 MHz and 61 MHz to 71 MHz carrier frequency
- $MER_{symp} \geq 23 \text{ dB}$ over 5 MHz to 10 MHz and 71 MHz to 85 MHz carrier frequency

Case 2: Flat channel, transmit equalization ON

Case 2a: for TDMA/QPSK, $MER_{\text{symp}} \geq 30$ dB.

Case 2b: for S-CDMA and all TDMA modulations except QPSK, $MER_{\text{symp}} \geq 35$ dB.

Case 2c: for S-CDMA, $MER_{\text{chip}} \geq 33$ dB.

Case 3: Echo channel, transmit equalization ON. (Echo channel MUST not be limited by dynamic range and MUST be a true representation described in Table 5-2)

Case 3a: In the presence of a single echo selected from the channel micro-reflections defined in Table 5-2, the measured MER_{symp} MUST be ≥ 30 dB for TDMA/QPSK, and ≥ 33 dB for S-CDMA and all TDMA modulations, except QPSK.

Case 3b: In the presence of two or three of the echoes defined in Table 5-2 (at most one of each specified magnitude and delay), the measured MER_{symp} MUST be ≥ 29 dB.

Since the table does not bound echo delay for the -30 dBc case, for testing purposes, it is assumed that the time span of the echo at this magnitude is less than or equal to 1.5 μ s.

The CMTS MUST provide a test mode in which it:

- accepts equalizer coefficients via an external interface, (e.g., Ethernet)
- sends the coefficients to the CM's pre-equalizer via ranging response message (both load and convolve modes)
- does not adjust the CM's frequency, timing or power.

6.2.22.4 Filter Distortion

The following requirements assume that any pre-equalization is disabled.

6.2.22.4.1 Amplitude

The spectral mask MUST be the ideal square-root raised-cosine spectrum with $\alpha = 0.25$, within the ranges given in Table 6-19.

Table 6-19 – Single Channel Filter Amplitude Distortion

Frequency	Spectral Mask Limits	
	Lower	Upper
$f_c - 5R_s/8$	-	-30dB
$f_c - R_s/2$	-3.5dB	-2.5dB
$f_c - 3R_s/8$ to $f_c - R_s/4$	-0.5dB	+0.3dB
$f_c - R_s/4$ to $f_c + R_s/4$	-0.3dB	+0.3dB
$f_c + R_s/4$ to $f_c + 3R_s/8$	-0.5dB	+0.3dB
$f_c + R_s/2$	-3.5dB	-2.5dB
$f_c + 5R_s/8$	-	-30dB

Where f_c is the center frequency, R_s is the modulation rate, and the spectral density is measured with a resolution bandwidth of 10 kHz or less.

6.2.22.5 Carrier Phase Noise

The upstream transmitter total integrated phase noise (including discrete spurious noise) MUST be less than or equal to -46 dBc, summed over the spectral regions spanning 200 Hz to 400 kHz above and below the carrier.

The upstream transmitter total integrated phase noise (including discrete spurious noise) MUST be less than or equal to -44 dBc, summed over the spectral regions spanning 8 kHz to 3.2 MHz above and below the carrier.

The CM MUST provide a test mode in which:

- A continuous (non-burst), unmodulated (CW) upstream signal is transmitted at the commanded carrier frequency, modulation rate and level. This is equivalent to replacing the chip sequence at the spreader output with the constant sequence (1, 1, 1, 1, 1, 1,...) at nominal amplitude, equal on both I and Q.
- The CM tracks the downstream symbol clock and uses it to generate the upstream symbol clock as in normal synchronous operation.

6.2.22.6 Channel Frequency Accuracy

The CM MUST implement the assigned channel frequency within 50 parts per million, over a temperature range of 0 to 40 degrees C, and up to five years from date of manufacture.

6.2.22.7 Modulation Rate Accuracy

The MDD message [MULPI] provides a TLV field that indicates whether the CMTS is providing a synchronous downstream symbol clock, i.e., whether the downstream symbol clock is locked to the 10.24 MHz Master Clock. If the MDD indicates that the downstream symbol clock is synchronous, the CM MUST lock its timing to the downstream symbol clock. If the MDD indicates that the downstream symbol clock is asynchronous, the CM MUST NOT lock its timing to the downstream symbol clock.

In synchronous operation (i.e., when the CM is locked to the downstream symbol clock), the upstream modulator MUST lock the upstream modulation rate of all upstream channels (S-CDMA and/or TDMA) to the downstream symbol rate, subject to the symbol timing jitter requirements of Section 6.2.22.8.

In asynchronous operation (i.e., when the CM timing is not locked to the downstream symbol clock), the upstream modulator MUST provide an absolute modulation rate accuracy of symbol rates, 50 parts per million, over a temperature range of 0 to 40 degrees C, up to five years from date of manufacture.

6.2.22.8 Modulation Timing Jitter

6.2.22.8.1 Symbol Timing Jitter for TDMA Upstream Channels in Asynchronous Operation

For TDMA upstream channels in asynchronous operation, peak-to-peak symbol jitter, referenced to the previous symbol zero-crossing of the transmitted waveform, MUST be less than 0.02 of the nominal symbol duration over a 2 s period. In other words, the difference between the maximum and the minimum symbol duration during the 2 s period shall be less than 0.02 of the nominal symbol duration for each of the five upstream symbol rates.

For TDMA upstream channels in asynchronous operation, the peak-to-peak cumulative phase error, referenced to the first symbol time, and with any fixed symbol frequency offset factored out, MUST be less than 0.04 of the nominal symbol duration over a 0.1 s period. In other words, the difference between the maximum and the minimum cumulative phase error during the 0.1 s period shall be less than 0.04 of the nominal symbol duration for each of the five upstream symbol rates. Factoring out a fixed symbol frequency offset is to be done by using the computed mean symbol duration during the 0.1 s.

6.2.22.8.2 Modulation Timing Jitter for Synchronous Operation

All jitter specifications assume a downstream input to the CM per 6.3.5, 6.3.3 and [DRFI].

For S-CDMA and TDMA upstream channels in synchronous operation, upstream modulation clock timing error (with the mean error subtracted out) relative to the CMTS master clock **MUST** be less than 0.005 RMS of the modulation interval over a 35 s measurement duration. This applies to: 1) the worst-case jitter and frequency drift specified for the CMTS Master clock and the CMTS downstream symbol clock in the requirements above; and 2) for any round-trip propagation delay up to the maximum allowed.

The CM upstream modulation clock **SHOULD** track the jitter components below 10 Hz in the input downstream symbol clock with an error transfer function below -25 dB. The CM upstream modulation clock **SHOULD** attenuate the jitter components in the input downstream symbol clock above 200 Hz.

The CM **MUST** provide a test mode in which:

- A continuous (non-burst) upstream signal is transmitted at the commanded carrier frequency, modulation rate and level.
- The modulation sequence at the transmit equalizer input is replaced with an alternating binary sequence (1, -1, 1, -1, 1, -1,...) at nominal amplitude, equal on both I and Q.
- The CM tracks the downstream symbol clock and uses it to generate the upstream modulation clock as in normal synchronous operation.

6.2.23 Upstream Demodulator Input Power Characteristics

The instantaneous input signal level, including ingress and noise to the upstream demodulator, **MUST NOT** exceed 29 dBmV in the 5-85 MHz frequency range of operation. The intended received power in each carrier **MUST** be within the values shown in Table 6-20.

The demodulator **MUST** operate within its defined performance specifications with received bursts within ± 6 dB of the nominal commanded received power. Thus the maximum power in a received burst **MUST** be less than or equal 29 dBmV.

Table 6-20 - Upstream Channel Demodulator Input Power Characteristics

Modulation Rate (kHz)	Maximum Range (dBmV)	Applicability
160	-13 to +17	CMTS MAY support
320	-13 to +17	CMTS MAY support
640	-13 to +17	CMTS MAY support
1280	-13 to +17	CMTS MUST support
2560	-10 to +20	CMTS MUST support
5120	-7 to +23	CMTS MUST support

6.2.24 Upstream Electrical Output from the CM

The CM MUST output an RF modulated signal with the characteristics delineated in Table 6-21.

Table 6-21 –Electrical Output from CM

Parameter	Value
Frequency	5 to 42 MHz edge to edge The following option MAY be provided: Mode Selection Option: Mode 1: 5 to 42 MHz edge to edge Mode 2: 5 to 85 MHz edge to edge
Level range per channel (Multiple Transmit Channel mode disabled, or only Multiple Transmit Channel mode enabled with one channel in the TCS)	TDMA: P_{\min} to +57 dBmV (32 QAM, 64 QAM) P_{\min} to +58 dBmV (8 QAM, 16 QAM) P_{\min} to +61 dBmV (QPSK) S-CDMA: P_{\min} to +56 dBmV (all modulations) where P_{\min} = +17 dBmV, 1280 kHz modulation rate P_{\min} = +20 dBmV, 2560 kHz modulation rate P_{\min} = +23 dBmV, 5120 kHz modulation rate
Level range per channel (two channels in the TCS)	TDMA: P_{\min} to +54 dBmV (32 QAM, 64 QAM) P_{\min} to +55 dBmV (8 QAM, 16 QAM) P_{\min} to +58 dBmV (QPSK) S-CDMA: P_{\min} to +53 dBmV (all modulations) where P_{\min} = +17 dBmV, 1280 kHz modulation rate P_{\min} = +20 dBmV, 2560 kHz modulation rate P_{\min} = +23 dBmV, 5120 kHz modulation rate
Level range per channel (three or four channels in the TCS)	TDMA: P_{\min} to +51 dBmV (32 QAM, 64 QAM) P_{\min} to +52 dBmV (8 QAM, 16 QAM) P_{\min} to +55 dBmV (QPSK) S-CDMA: P_{\min} to +53 dBmV (all modulations) where P_{\min} = +17 dBmV, 1280 kHz modulation rate P_{\min} = +20 dBmV, 2560 kHz modulation rate P_{\min} = +23 dBmV, 5120 kHz modulation rate
Modulation Type	QPSK, 8 QAM, 16 QAM, 32 QAM, 64 QAM, and 128 QAM

Parameter	Value
Modulation Rate (nominal)	TDMA: 1280, 2560, and 5120 kHz S-CDMA: 1280, 2560, and 5120 kHz Optional pre-3.0-DOCSIS operation, TDMA: 160, 320, and 640 kHz
Bandwidth	TDMA: 1600, 3200, and 6400 kHz S-CDMA: 1600, 3200, and 6400 kHz Optional pre-3.0-DOCSIS operation, TDMA: 200, 400, and 800 kHz
Output impedance	75 ohms
Output Return Loss	>6 dB (across the selected upstream frequency range 5-42 MHz or 5-85 MHz). > 6 dB (108 to 870 MHz)
Connector	F connector per [ISO/IEC-61169-24] or [SCTE 02] (common with the input)

6.2.25 Upstream CM Transmitter Capabilities

The CM advertises its capabilities to the CMTS. These capabilities include:

- Maximum number of active channels, which MUST be 4 or greater.
- Maximum number of 3.2 MHz channels, which MUST be no greater than the maximum number of active channels supported.
- Maximum number of 6.4 MHz channels, which MUST be no greater than the maximum number of 3.2 MHz channels supported, and MUST be 4 or greater.
- Selective Active Codes Mode 2 capability – yes/no.
- Code Hopping Mode 2 capability – yes/no.
- Supports 5 - 85 MHz upstream range – yes/no.

The CM MUST support all the capability it reports to the CMTS.

The CMTS MUST NOT command the CM to operate outside its advertised capabilities.

6.2.25.1 Description of CM Upstream Transmit Channel Set Capability

The CM communicates to the CMTS its capabilities in regards to the number of active upstream channels it can support. The method of conveyance must clearly and unambiguously describe the capabilities of the CM to the CMTS. The parameters for this messaging are defined in this section, and Appendix IV contains example calculations for this reporting from the CM to the CMTS of the number of active upstream channels of the various bandwidths supported by the CM.

Let X = number of 6.4 MHz channels, Y = number of 3.2 MHz channels, and Z = number of 1.6 MHz channels that are active at one time (the Transmit Channel Set). Let X_{\max} be the maximum number of 6.4 MHz channels that the CM can support, Y_{\max} be the maximum number of 3.2 MHz channels that the CM can support, and Z_{\max} be the maximum number of 1.6 MHz channels that the CM can support. Z_{\max} is equivalent to the maximum number of upstream transmitters that the CM can support.

The CM reports X_{\max} , Y_{\max} , and Z_{\max} during registration using modem capabilities encodings specified in [MULPI], Annex subsection, "Modem Capabilities Encoding." The CM MUST provide that $X_{\max} \leq Y_{\max} \leq Z_{\max}$.

Communications of X_{\max} , Y_{\max} , and Z_{\max} by a CM provides enough information to describe all the combinations possible for X, Y, and Z.

From these values the CMTS may calculate $B_{\max} = \max(6.4 * X_{\max}, 3.2 * Y_{\max}, 1.6 * Z_{\max})$.

Then, all values for (X, Y, Z) are allowed, which satisfy the four following inequalities:

$$6.4 * X + 3.2 * Y + 1.6 * Z \leq B_{\max}$$

$$X \leq X_{\max}$$

$$Y \leq Y_{\max}$$

$$X + Y + Z \leq Z_{\max}$$

6.3 Downstream

6.3.1 Downstream Protocol and Interleaving Support

The downstream PMD sublayer MUST conform to ITU-T Recommendations J.83, [ITU-T J.83-B], except for Section B.6.2. Interleaver depths are defined in Section 6.3.3 of the Downstream Radio Frequency Specification [DRFI]. The downstream RF output electrical requirements including downstream frequency plan, interleaver depths, spectrum format, clock, and symbol requirements are defined in [DRFI]. The CM MUST support the interleaver depths defined in Table 6-1 of [DRFI]. The CM MAY support the interleaver depths defined in Table 6-2 of [DRFI].

6.3.2 Downstream Electrical Input to CM

The CM MUST be able to accept any number of signals between one to MDBC simultaneously, located in frequency intervals that comply with the CMs TB list and Demod List. A CM MUST be capable of being reconfigured to receive different channels. At a minimum, the CM MUST be able to receive at least four downstream channels, located independently within an arbitrary 60 MHz window in the downstream frequency band. The CM MUST support bonding of any number of downstream channels up to its maximum. The CM MUST be able to locate and accept RF modulated signals located within channels defined in [CEA-542-B] for Harmonic Related Carrier (HRC), Incremental Related Carrier (IRC), and Standard (STD) North American frequency plans. Operation below a center frequency of 111 MHz is not required. The signals will have the characteristics defined in Table 6-22.

Table 6-22 - Electrical Input to CM

Parameter	Value
Center Frequency	111 to 867 MHz \pm 30 kHz
Level Range (one DOCSIS channel)	-15 dBmV to +15 dBmV
Modulation Type	64 QAM and 256 QAM
Symbol Rate (nominal)	5.056941 Msym/s (64 QAM) and 5.360537 Msym/s (256 QAM)
Bandwidth	6 MHz (alpha = 0.18 Square Root Raised Cosine shaping for 64 QAM and alpha = 0.12 Square Root Raised Cosine shaping for 256 QAM)
Total Input Power (40 MHz and above)	< 33 dBmV
Maximum power of any carrier input to CM	20 dBmV
Input (load) Impedance	75 ohms
Input Return Loss	> 6 dB (across the selected upstream frequency range 5-42 MHz or 5-85 MHz). > 6 dB (108 to 870 MHz)
Connector	F connector per [ISO/IEC-61169-24] or[SCTE 02] (common with the output)

6.3.3 CM BER Performance

The bit-error-rate performance of a CM MUST be as described in this section. The requirements apply to each individual downstream received channel with the interleaver set for the I = 128, J = 1 mode of interleaving.

6.3.3.1 64 QAM

6.3.3.1.1 64 QAM CM BER Performance

Implementation loss of the CM MUST be such that the CM achieves a post-FEC BER less than or equal to 10^{-8} when operating at a carrier to noise ratio (E_s/N_0) of 23.5 dB or greater. If it is not possible to measure post-FEC BER directly, Codeword Error Rate, R_C (as defined below) may be used. In this case, the CM MUST achieve a Codeword Error Rate of less than or equal to 9×10^{-7} when operating at a carrier to noise ratio (E_s/N_0) of 23.5 dB or greater.

Computation of Codeword Error Rate, R_C :

$$R_C = \frac{(E_U - E_{U0})}{(E_U - E_{U0}) + (E_C - E_{C0}) + (C - C_0)}$$

Where:

E_U is the value of the count of code words with uncorrectable errors;

E_C is the value of the count of code words with correctable errors and;

C is the value of the count of code words without errors.

Sample the values at the beginning of the test interval (indicated by the subscripted values, E_{U0} , E_{C0} , and C_0) and the end of the test interval (indicated by the values, E_U , E_C , and C).

6.3.3.1.2 Image Rejection Performance

Performance as described in Section 6.3.3.1.1 MUST be met with analog or digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

6.3.3.1.3 64 QAM Adjacent Channel Performance

Performance as described in Section 6.3.3.1.1 MUST be met with a digital signal at 0 dBc in the adjacent channels.

Performance as described in Section 6.3.3.1.1 MUST be met with an analog signal at +10 dBc in the adjacent channels.

Performance as described in Section 6.3.3.1.1, with an additional 0.2 dB allowance, MUST be met with a digital signal at +10 dBc in the adjacent channels.

6.3.3.2 256 QAM

6.3.3.2.1 256 QAM CM BER Performance

Implementation loss of the CM MUST be such that the CM achieves a post-FEC BER less than or equal to 10^{-8} when operating at a carrier to noise ratio (E_s/N_0) as shown below. If it is not possible to measure post-FEC BER directly, Codeword Error Rate, R_C (as defined in Section 6.3.3.1.1) may be used. In this case, the CM MUST achieve a Codeword Error Rate of less than or equal to 9×10^{-7} when operating at a carrier to noise ratio (E_s/N_0), as shown in the following:

Input Receive Signal Level	E_s/N_0
-6 dBmV to +15 dBmV	30 dB or greater
Less than -6 dBmV down to -15 dBmV	33 dB or greater

6.3.3.2.2 256 QAM Image Rejection Performance

Performance as described in Section 6.3.3.2.1 MUST be met with an analog or a digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

6.3.3.2.3 256 QAM Adjacent Channel Performance

Performance as described in Section 6.3.3.2.1 MUST be met with an analog or a digital signal at 0 dBc in the adjacent channels.

Performance as described in Section 6.3.3.2.1, with an additional 0.5 dB allowance, MUST be met with an analog signal at +10 dBc in the adjacent channels.

Performance as described in Section 6.3.3.2.1, with an additional 1.0 dB allowance, MUST be met with a digital signal at +10 dBc in the adjacent channels.

6.3.4 Downstream Multiple Receiver Capabilities

This section describes the mechanism by which the CM communicates to the CMTS the constraints on receive channel assignment imposed by the CM's capabilities. The CMTS needs to be aware of the CM's capabilities when it assigns or changes the parameters of a downstream channel received by a CM. If the CM were reconfigured

without regard to its capabilities, a disruption of downstream data and/or DOCSIS master clock synchronization could result.

The CM reports its capabilities to the CMTS by sending its Receive Channel Profiles (RCP). In response, the CMTS configures the CM by sending a Receive Channel Configuration (RCC).

The CM MUST support all the capability it reports to the CMTS.

The CMTS MUST NOT command the CM to operate outside its advertised capabilities.

The [MULPI] specification defines the detailed encodings by which Receive Channel Profiles and Receive Channel Configurations describe Receive Modules, Receive Channels, and the interconnections between them.

6.3.4.1 Receive Module Parameters

A Receive Module (RM) is a grouping of channels, together with parameters describing the constraints on those channels. Examples of RMs are:

- Tuner, with constraints on capture bandwidth and frequency range.
- Demodulator, with constraints on contiguous positioning of channels, frequency range, modulation order, and interleaver depth.

A Receive Channel Profile communicated from CM to CMTS defines the following attributes of each Receive Module:

- **Adjacent Channels:** The number of contiguous channels processed by the Receive Module.
- **Channel Block Range:** The minimum center frequency of the lowest-frequency channel in the block and the maximum center frequency of the highest-frequency channel in the block.
- **Common Physical Layer Parameters:** A list of physical layer parameters which are shared by all Receive Channels connected to the Receive Module. Examples may include modulation type and interleaver settings.
- **Connection Capability:** A list of the higher level (closer to RF port) Receive Modules to which this Receive Module can connect.

A Receive Channel Configuration (RCC), communicated from CMTS to CM, assigns one or more of the following attributes of a Receive Module:

- **First Channel Center Frequency:** Center frequency of the lowest-frequency channel of a block of adjacent channels.
- **Connection Assignment:** Specifies to which single higher-level Receive Module to connect.

6.3.4.2 Receive Channel Parameters

A Receive Channel (RC) is the packet stream output to the MAC layer originating from a single QAM channel. A Receive Channel Profile, communicated from CM to CMTS, defines the following attributes of each Receive Channel:

- **Connection Capability:** A list of the Receive Modules to which the Receive Channel can connect.
- **Connected Offset:** The offset of the Receive Channel within a block of adjacent channels.
- **Primary Downstream Channel Capability:** A flag indicating whether the Receive Channel is capable of providing the DOCSIS master clock reference to the CM.

A Receive Channel Configuration communicated from CMTS to CM assigns the following attributes to a Receive Channel:

- Center Frequency Assignment: The RF center frequency of the Receive Channel.
- Primary Downstream Channel Indicator: A flag indicating that the CMTS assigns this Receive Channel the responsibility to provide master clock reference timing to the CM.
- Connection Assignment: Specifies to which Receive Module the Receive Channel is to connect.

6.3.4.3 Standard Receive Channel Profile

In order to limit the complexity in the CMTS of configuring arbitrarily complicated Receive Channel Profiles, DOCSIS defines a set of "Standard" Receive Channel Profiles that describe a minimal set of constraints.

A CM reports to the CMTS at least one standard RCP as well as its "Manufacturer" RCP that gives more details of its capabilities and constraints. If the CMTS configures the CM with a Receive Channel Configuration based on a Standard RCP, some CM capabilities represented in the Manufacturer RCP but outside the Standard RCP may not be available.

The Standard Receive Channel Profile named "6-DOCSIS-01" is defined for 6 MHz operation as depicted in Figure 6-33. The RF input from the cable plant is connected to a Receive Module with a block of 10 adjacent channels. This Receive Module represents the constraints of a tuner with a 60 MHz capture bandwidth that can be placed at any point in the full DOCSIS frequency range. The demodulation of four channels anywhere within the capture bandwidth is represented by the four Receive Channels, each corresponding to the packet stream from a single QAM channel output to the MAC layer.

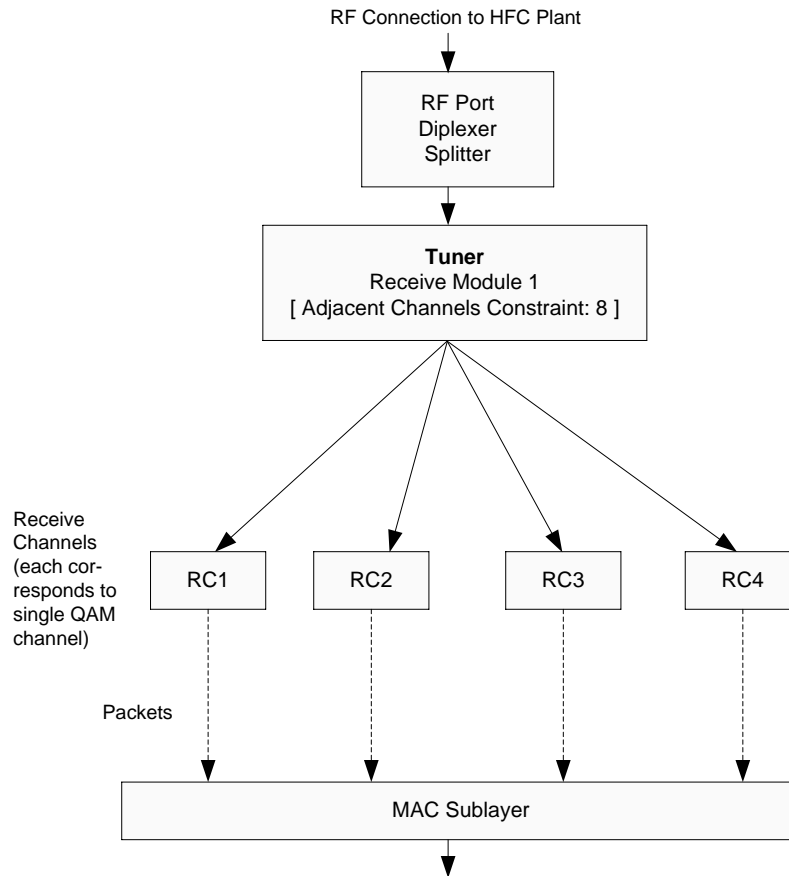


Figure 6-33 - 6 MHz DOCSIS Standard Receive Channel Profile 01

For this Standard Receive Channel Profile, the CMTS sends the following parameters to the CM in the Receive Channel Configuration:

- Center frequency of the first channel in the 10-channel Receive Module.
- Center frequency of each Receive Channel within the capture bandwidth of the Receive Module.

6.3.4.4 Example Manufacturer Receive Channel Profile

A more complex example showing a Manufacturer Receive Channel Profile as reported by the CM to the CMTS is depicted in Figure 6-34. The RF input from the cable plant is distributed to three Receive Modules, corresponding to tuner modules with capture bandwidth constraints. A network connects the tuners to a group of four Receive Modules, corresponding to demodulators with contiguous channel constraints. Each demodulator connects to four Receive Channels, each representing a single QAM channel output to the MAC layer.

In response, the CMTS configures each Receive Module and Receive Channel as follows:

- Frequency range of each analog tuner.
- Center frequency of first channel of each demodulator.
- Center frequency of each Receive Channel.
- Interconnection between Receive Modules.
- Interconnection between lower-level (demodulator) Receive Modules and Receive Channels.

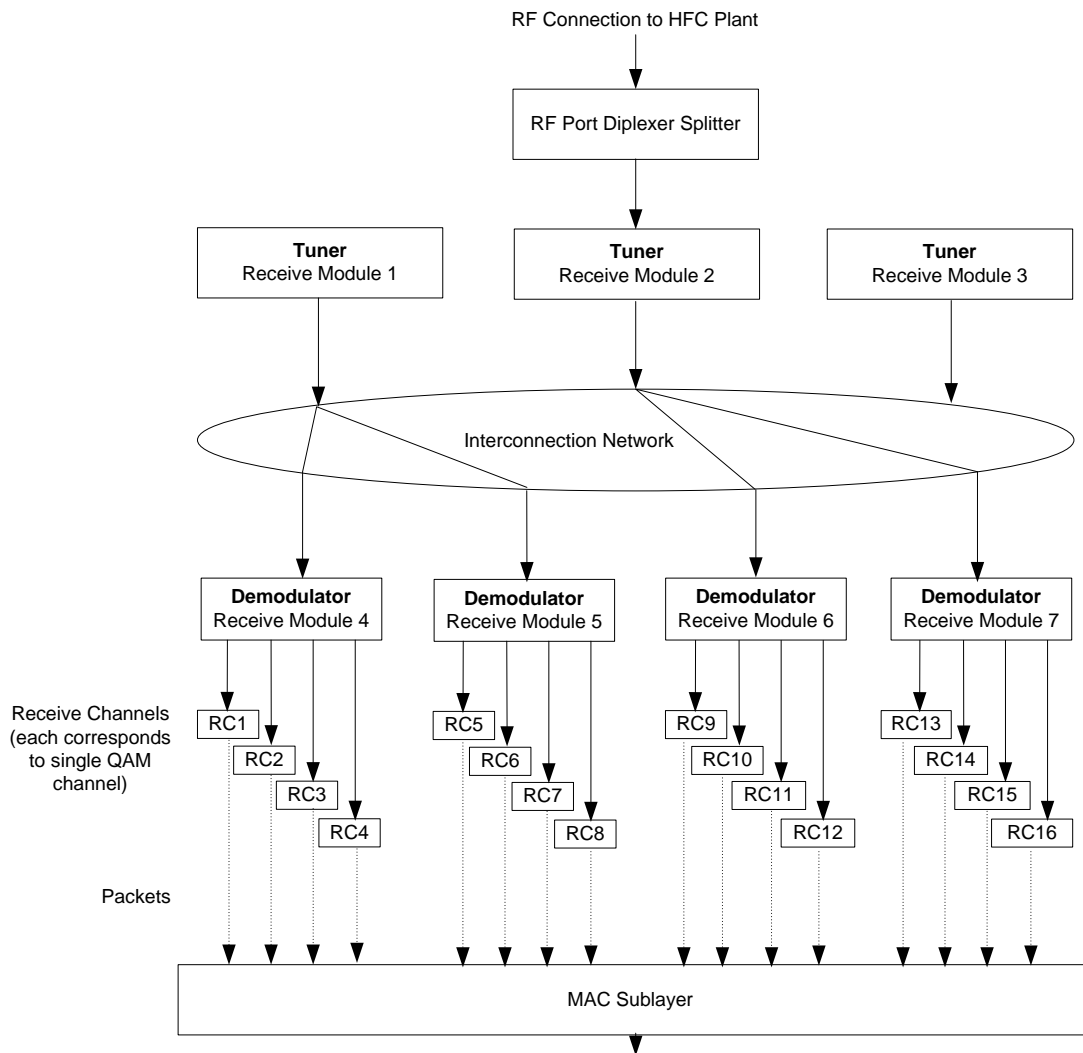


Figure 6-34 - Example Manufacturer Receive Channel Profile

6.3.4.5 Reconfiguration of CM Receive Channels

Before the CMTS sends a DBC-REQ (dynamic bonding change request) asking the CM to change parameters of a receive channel, the CMTS can determine the constraints and connectivity of the CM based on the Receive Channel Profile and Receive Channel Configuration. From this information, the CMTS can determine whether the requested change can be implemented by the CM, and whether the change will interrupt other downstream channels. In this way, the CMTS can minimize and/or schedule the disruption of downstream traffic and any perturbations to DOCSIS master clock timing resulting from receive channel parameter changes.

6.3.5 Non-Synchronous DS Channel Support

The "primary downstream channel" of a CM is defined as the downstream channel from which it derives CMTS master clock timing for upstream transmission. All other concurrently received channels are called "secondary downstream channels."

A DOCSIS 3.0 CM MUST support secondary downstream channels that are asynchronous to the primary downstream channel, i.e., with a different long-term average QAM symbol rate, but that conform to all DOCSIS 3.0 specifications.

Annex A Timing Requirements for Supporting Business Services over DOCSIS

The CMTS transmits its 10.24 MHz reference clock timing to the CM via the downstream QAM symbol clock and via timing information embedded in downstream SYNC messages. The CM utilizes this downstream clock information to regenerate a local 10.24 MHz reference clock. The CM reference clock, together with ranging commands, is used to accurately time upstream TDMA or S-CDMA burst transmissions. This same clock transfer mechanism can be leveraged to transfer TDM timing from the CMTS to the CM as follows. The CMTS locks its reference to a TDM source. The CM recovers the reference and regenerates a TDM output clock.

A.1 CMTS

DOCSIS 3.0 CMTSs have the following timing and synchronization requirements. To fully support T1 and E1 Services, the CMTS should, in addition, comply with the requirements as outlined in [TEI]. Specifically, the CMTS should be able to send TDM timing information down the DOCSIS RF downstream. The definitions and abbreviations used here are defined in [TEI]:

- The CMTS MUST be able to lock the downstream DOCSIS symbol clock to an external Stratum 1 clock source (PRS or DTI).
- The CMTS MUST be able to lock the downstream DOCSIS SYNC messages to an external Stratum 1 clock source (PRS or DTI).

There are further timing performance requirements for the CMTS symbol clock and SYNC messages, as specified in [TEI] Section, "TE-CMTS Timing Requirements".

A.2 CM

DOCSIS 3.0 CMs have the following timing and synchronization requirements. To fully support T1 and E1 Services, the CM should comply with the requirements as outlined in [TEI]. Specifically, the CM should be able to receive TDM timing information from the DOCSIS RF downstream. The definitions and abbreviations used here are defined in [TEI]:

- The CM SHOULD have a TDM clock output (in DOCSIS 3.0 operation, and in pre-3.0 DOCSIS S-CDMA operation, synchronized to the downstream Symbol Clock) from which a T1 or E1 clock can be regenerated.
- The CM SHOULD have a TDM clock output (in pre-3.0 DOCSIS TDMA operation, synchronized to the SYNC message timing information) from which a T1 or E1 clock can be regenerated.

The same clock source that is used for upstream timing SHOULD be used to generate the TDM clock output. The 10.24 MHz (or related) reference clock is recovered by the CM by locking to the downstream symbol clock (normal operation, highest performance) or downstream SYNC messages (legacy operation, reduced performance). The recovered reference clock is then used to generate a T1 or E1 clock via a fractional (M/N) PLL or equivalent method.

The CM TDM clock output should not have frequency components that cannot be attenuated by practical filters to regenerate a compliant T1 or E1 clock. Further guidelines for regenerating T1 and E1 clocks from downstream TDM timing information are outlined in [TEI] Section, "Network Clock Recovery." The compliance and timing performance requirements for regenerated T1 and E1 clocks are specified in [TEI] Section, "TSP – TDM Service Processor".

Annex B Additions and Modifications for European Specification

This annex applies to the second technology option referred to in Section 1.1. For the first option, refer to Sections 5 and 6.

This annex defines the interface for the physical layer used in conjunction with the European multi-program television distribution. It describes the electrical characteristics and signal processing operations for what is generally called EuroDOCSIS Cable Modem (CM) and Cable Modem Termination System (CMTS). This is an optional annex and in no way affects certification of equipment adhering to the North American technology option described in the sections referenced above.

Requirements referring to the absolute level of electrical power are represented in units of dBmV throughout this specification. This is in contrast to other European standards applicable in cabled distribution systems for television and sound signals (e.g., [EN 50083-7] or [EN 50083-10]) where such requirements are expressed traditionally in units of dB μ V. The values in dBmV used in this specification can be transformed into values in dB μ V by adding 60 dB.

The numbering of the paragraphs in this annex has been maintained such that the suffix after the letter for the annex refers to the part of the specification where the described changes apply to. In cases where the requirements for both technology options are identical, a reference is provided to the main text.

B.1 Scope

B.1.1 Introduction and Purpose

See Section 1.1.

B.1.2 Background

B.1.2.1 Broadband Access Network

A coaxial-based broadband access network is assumed. This may take the form of either an all-coax or hybrid fiber/coax (HFC) network. The generic term "cable network" is used here to cover all cases.

A cable network uses a shared-medium, tree-and-branch architecture, with analog transmission. The key functional characteristics assumed in this document are the following:

- Two-way transmission
- A maximum optical/electrical spacing between the CMTS and the most distant CM of 160 km (route meters) in each direction
- A maximum differential optical/electrical spacing between the CMTS and the closest and most distant modems of 160 km (route meters) in each direction

At a propagation velocity in fiber of approximately 5 ns/m, 160 km of fiber in each direction results in a round-trip delay of approximately 1.6 ms.

B.1.2.2 Network and System Architecture**B.1.2.2.1 The DOCSIS Network**

See Section 1.2.2.1.

B.1.2.3 Service Goals

See Section 1.2.3.

B.1.2.4 Statement of Compatibility

This annex specifies an interface, commonly referred to as EuroDOCSIS 3.0, which is the third generation of the interface, (earlier versions are commonly referred to as EuroDOCSIS 1.x and 2.0). EuroDOCSIS 3.0 MUST be backward- and forward-compatible with equipment built to the previous specifications. EuroDOCSIS 3.0-compliant CMs MUST interoperate seamlessly with EuroDOCSIS 2.0 and EuroDOCSIS 1.x CMTSSs, albeit in the 2.0 and 1.x modes, as the case may be. EuroDOCSIS 3.0-compliant CMTSSs MUST seamlessly support EuroDOCSIS 2.0 and EuroDOCSIS 1.x CMs.

B.1.2.5 Reference Architecture

See Section 1.2.5.

B.1.2.6 DOCSIS 3.0 Documents

A list of the specifications in the DOCSIS 3.0 series is provided in Table B–1. If additions and modifications are required to define the European technology option, they are included in a normative annex to each of the specifications. Otherwise, the listed specifications equally apply to EuroDOCSIS 3.0. For further information, please refer to <http://www.cablemodem.com>.

Table B–1 - DOCSIS 3.0 Series of Specifications

Designation	Title
CM-SP-PHYv3.0	Physical Layer Specification
CM-SP-MULPIv3.0	Media Access Control and Upper Layer Protocols Interface Specification
CM-SP-OSSIV3.0	Operations Support System Interface Specification
CM-SP-SECv3.0	Security Specification

This annex is defining the interface for the physical layer of the European technology option.

B.1.3 Requirements

See Section 1.3.

B.1.4 Conventions

See Section 1.4.

B.1.5 Organization of Document

See Section 1.5.

B.2 References

B.2.1 Normative References

See Section 2.1.

B.2.2 Informative References

See Section 2.2.

B.2.3 Reference Acquisition

See Section 2.3.

B.3 Terms and Definitions

See Section 3.

B.4 Abbreviations and Acronyms

See Section 4.

B.5 Functional Assumptions

This section describes the characteristics of a cable television plant, assumed to be for the purpose of operating a data-over-cable system. It is not a description of CMTS or CM parameters. The data-over-cable system **MUST** be interoperable within the environment described in this section.

Whenever a reference in this section to frequency plans or to compatibility with other services conflicts with any legal requirement for the area of operation, the latter shall take precedence. Any reference to analog TV signals in a particular frequency band does not imply that such signals are physically present.

B.5.1 Equipment Assumptions

B.5.1.1 Frequency Plan

In the downstream direction, the cable system is assumed to have a pass band with a lower edge of down to 47 MHz. Typically, the lower edge of the downstream pass band is at 87.5 MHz. The upper edge is implementation-dependent but is typically in the range of 300 to 862 MHz. Within that pass band, PAL/SECAM analog television signals in 7/8 MHz channels and FM radio signals are assumed to be present, as well as other narrowband and wideband digital signals. 8 MHz channels are used for data communication.

In the upstream direction, the cable system is assumed to have a pass band of 5-65 MHz. PAL/SECAM analog television signals in 7/8 MHz channels may be present, as well as other signals.

B.5.1.2 Compatibility with Other Services

See Section 5.1.2.

B.5.1.3 Fault Isolation Impact on Other Users

See Section 5.1.3.

B.5.1.4 Cable System Terminal Devices

Compliance with EMC requirements is not covered by this specification. The protection requirements with respect to electromagnetic compatibility are contained in harmonized standards published in the Official Journal of the European Union.

Any reference in the present document to the transmission of television in the forward channel that is not consistent with [EN 300 429] is outside the normative scope as only [EN 300 429] is used for digital multi-program TV distribution by cable in European applications.

Requirements for safety are outside the scope of the present document. Safety standards for European applications are published by CENELEC. Examples of such CENELEC product safety standards are [EN 60950-1] and [EN 50083-1]. For CENELEC safety categories of interfaces, see [EG 201 212].

B.5.2 RF Channel Assumptions

See Section 5.3.

B.5.2.1 Transmission Downstream

The RF channel transmission characteristics of the cable network in the downstream direction assumed for the purposes of minimal operating capability are described in Table B–2. These numbers assume total average power of a digital signal in an 8 MHz channel bandwidth for carrier levels unless indicated otherwise. For impairment levels, the numbers in Table B–2 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 B–2 assume nominal analog video carrier level (peak envelope power) in a 7/8 MHz channel bandwidth. Transmission is from head-end combiner to the CM input at the customer location. All conditions are present concurrently.

Table B–2 - Assumed Downstream RF Channel Transmission Characteristics

Parameter	Value
Frequency range	Cable system downstream operating range is from 47 MHz to 862 MHz. However, the operating range for data communication is from 108 to 862 MHz and the values in this table only apply to this frequency range. The use of frequencies between 108 and 136 MHz may be forbidden due to national regulation with regard to interference with aeronautical navigation frequencies.
RF channel spacing (design bandwidth)	7/8 MHz, 8 MHz channels are used for data communication
Transit delay from head-end to most distant customer	≤ 0.800 ms (typically much less)
Carrier-to-noise ratio in an 8 MHz band (analog video level)	Not less than 44 dB ¹
Carrier-to-interference ratio for total power (discrete and broadband ingress signals)	Not less than 52 dB within the design bandwidth
Composite triple beat distortion for analog modulated carriers	Not greater than -57 dBc within the design bandwidth ²
Composite second-order distortion for analog modulated carriers	Not greater than -57 dBc within the design bandwidth ³
Cross-modulation level	Under consideration

Parameter	Value
Amplitude ripple	2.5 dB in 8 MHz
Group delay ripple in the spectrum occupied by the CMTS	100 ns over frequency range 0.5 – 4.43 MHz
Micro-reflections bound for dominant echo	-10 dBc @ ≤ 0.5 μs -15 dBc @ ≤ 1.0 μs -20 dBc @ ≤ 1.5 μs -31.5 dBc @ > 1.5 μs
Carrier hum modulation	Not greater than -46 dBc (0.5%)
Burst noise	Not longer than 25 μs at a 10 Hz average rate
Seasonal and diurnal signal level variation	8 dB
Signal level slope, 85 – 862 MHz	Maximum slope of 12 dB in either the positive or negative direction
Maximum analog video carrier level at the system outlet, inclusive of above signal level variation	17 dBmV ³
Lowest analog video carrier level at the system outlet, inclusive of above signal level variation	0 dBmV ⁴
¹ This presumes that the average digital carrier is operated at analog peak carrier level. When the digital carrier is operated below the analog peak carrier level, this C/N may be less ² For SECAM systems the value is no greater than -52 dBc within the design bandwidth ³ For SECAM systems the value is 14 dBmV ⁴ For SECAM systems the value is -3 dBmV	

B.5.2.2 Transmission Upstream

The RF channel transmission characteristics of the cable network in the upstream direction assumed for the purposes of minimal operating capability are described in Table B–3. Transmission is from head-end combiner to the CM input at the customer location. All conditions are present concurrently.

Table B–3 - Assumed Upstream RF Channel Transmission Characteristics

Parameter	Value
Frequency range.	5 to 65 MHz edge to edge
Transit delay from head-end to most distant customer	≤ 0.800 ms (typically much less)
Carrier-to-noise ratio in active channel	Not less than 22 dB
Carrier-to-ingress power (the sum of discrete and broadband ingress signals) ratio in active channel	Not less than 22 dB ¹
Carrier-to-interference (the sum of noise, distortion, common-path distortion and cross-modulation) ratio in active channel	Not less than 22 dB ²
Carrier hum modulation	Not greater than -23 dBc (7.0%)
Burst noise	Not longer than 10 μs at a 1 KHz average rate for most cases ^{2,3}

Parameter	Value
Amplitude ripple across upstream operating frequency range (maximum)	2.5 dB in 2 MHz
Group delay ripple across upstream operating frequency range (maximum)	300 ns in 2 MHz
Micro-reflections (maximum) – single echo	-10 dBc @ $\leq 0.5 \mu\text{s}$ -20 dBc @ $\leq 1.0 \mu\text{s}$ -31.5 dBc @ $> 1.0 \mu\text{s}$
Seasonal and diurnal signal level variation	Not greater than 12 dB min to max
¹ Ingress avoidance or tolerance techniques may be used to ensure operation in the presence of time-varying discrete ingress signals that could be as high as 0 dBc. ² Amplitude and frequency characteristics sufficiently strong to partially or wholly mask the data carrier. ³ Impulse noise levels more prevalent at lower frequencies (<15 MHz).	

B.5.2.2.1 Availability

See Section 5.2.2.1.

B.5.3 Transmission Levels

See Section 5.3.

B.5.4 Frequency Inversion

See Section 5.4.

B.6 Physical Media Dependent Sublayer Specification

B.6.1 Scope

This section defines the electrical characteristics and signal processing operations for a cable modem (CM) and Cable Modem Termination System (CMTS). It is the intent of this specification 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 this specification to imply any specific implementation.

This section applies to the second technology option referred to in Section 1. In cases where the requirements for both technology options are identical, a reference is provided to the main text.

B.6.2 Upstream

B.6.2.1 Overview

The upstream Physical Media Dependent (PMD) sublayer uses an FDMA/TDMA (herein called TDMA mode) or FDMA/TDMA/S-CDMA (herein called S-CDMA mode) burst type format, which provides six modulation rates and multiple modulation formats. The use of TDMA or S-CDMA mode is configured by the CMTS via MAC messaging.

FDMA (frequency division multiple access) indicates that multiple RF channels are assigned in the upstream band. A CM transmits on one or more RF channels and may be reconfigured to change channels.

A CM MUST support at least 4 active upstream channels (which are referred to as the Transmit Channel Set for that CM).

The CM reports its maximum number of upstream channels capability, and certain other of its capability characteristics, to the CMTS (Section B.6.2.25).

The CM MUST be able to operate each channel in the Transmit Channel Set, simultaneously, anywhere in the upstream band, subject to restrictions on transmit power across the channels and on reconfiguration of certain transmit properties (see Sections B.6.2.19 and B.6.2.20, and subsections of each). The CMTS MUST be capable of assigning and receiving each RF channel anywhere in the upstream band. The CMTS MUST set the number of assigned channels and assigned channels' center frequency, and all other channel attributes. The CMTS MAY change the number of assigned channels and the channel attributes. Each RF channel has its own set of UCD parameters as defined in [MULPI] Section 6.4.3.

TDMA (time division multiple access) indicates that upstream transmissions have a burst nature. A given RF channel is shared by multiple CMs via the dynamic assignment of time slots. S-CDMA (synchronous code division multiple access) indicates that multiple CMs can transmit simultaneously on the same RF channel and during the same TDMA time slot, while being separated by different orthogonal codes.

In this document, the following naming conventions are used. For TDMA, the term "modulation rate" refers to the RF channel symbol rate (160 to 5120 ksym/s). For S-CDMA, the term "modulation rate" refers to the "chip rate," which is the rate (1280 to 5120 kHz) of the individual elements (chips) of the S-CDMA spreading code. Modulation rates are represented in units of "Hz" denoting the number of symbols per second in TDMA mode or the number of chips per second in S-CDMA mode. The "modulation interval" is the symbol period (TDMA mode) or chip period (S-CDMA mode) and is the reciprocal of the modulation rate. At the output of the spreader, a group of 128 chips which comprise a single S-CDMA spreading code, and are the result of spreading a single information (QAM constellation) symbol is referred to as a "spread symbol." The period of a spread symbol (128 chips) is called a "spreading interval." A "burst" is a physical RF transmission that contains a single preamble plus data, and (in the absence of preceding and following bursts) exhibits RF energy ramp-up and ramp-down.

In some cases logical zeros or logical ones are used to pad data blocks; this indicates data with zero-valued or one-valued binary bits, which result in non-zero transmitted RF energy. In other cases a numerical zero is used; this denotes, for example, symbols which result in zero transmitted RF energy (after ramp-up and ramp-down are taken into account).

The modulation format includes pulse shaping for spectral efficiency, is carrier-frequency agile, and has selectable output power level.

Each burst supports a flexible modulation order, modulation rate, preamble, randomization of the payload, and programmable FEC encoding.

All of the upstream transmission parameters associated with burst transmission outputs from the CM are configurable by the CMTS via MAC messaging. Many of the parameters are programmable on a burst-by-burst basis.

The PMD sublayer can support a near-continuous mode of transmission, wherein ramp-down of one burst MAY overlap the ramp-up of the following burst, so that the transmitted envelope is never zero. In TDMA mode, the system timing of the TDMA transmissions from the various CMs MUST provide that the center of the last symbol of one burst and the center of the first symbol of the preamble of an immediately following burst are separated by at least the duration of five symbols. The guard band MUST be greater than or equal to the duration of five symbols plus the maximum timing error. Timing error is contributed by both the CM and CMTS. CM timing performance is specified in Section B.6.2.20.1. Maximum timing error and guard band may vary with CMTSs from different vendors. The term guard time is similar to the guard band, except that it is measured from the end of the last symbol of one burst to the beginning of the first symbol of the preamble of an immediately following burst. Thus, the guard time is equal to the guard band – 1.

The PMD sublayer also supports a synchronous mode of transmission when using S-CDMA, wherein ramp-down of one burst MAY completely overlap the ramp-up of the following burst, so that the transmitted envelope is never zero. There is no guard time for transmission on S-CDMA channels. The system timing of the S-CDMA transmissions from the various CMs MUST provide adequate timing accuracy so that different CMs do not appreciably interfere with each other. S-CDMA utilizes precise synchronization so that multiple CMs can transmit simultaneously.

The upstream modulator is part of the cable modem which interfaces with the cable network. The modulator contains the electrical-level modulation function and the digital signal-processing function; the latter provides the FEC, preamble prepend, symbol mapping, and other processing steps.

At the Demodulator, similar to the Modulator, there are two basic functional components: the demodulation function and the signal processing function. The Demodulator resides in the CMTS and there is one demodulation function (not necessarily an actual physical demodulator) for each carrier frequency in use. The demodulation function receives all bursts on a given frequency.

The demodulation function of the Demodulator accepts a varying-level signal centered around a commanded power level and performs symbol timing and carrier recovery and tracking, burst acquisition, and demodulation. Additionally, the demodulation function provides an estimate of burst timing relative to a reference edge, an estimate of received signal power, may provide an estimate of signal-to-noise ratio, and may engage adaptive equalization to mitigate the effects of a) echoes in the cable plant, b) narrowband ingress and c) group delay. The signal-processing function of the Demodulator performs the inverse processing of the signal-processing function of the Modulator. This includes accepting the demodulated burst data stream and decoding, etc. The signal-processing function also provides the edge-timing reference and gating-enable signal to the demodulators to activate the burst acquisition for each assigned burst slot. The signal-processing function may also provide an indication of successful decoding, decoding error, or fail-to-decode for each codeword and the number of corrected Reed-Solomon symbols in each codeword. For every upstream burst, the CMTS has a prior knowledge of the exact burst length in modulation intervals (see Sections B.6.2.5, B.6.2.5.1, B.6.2.5.2, B.6.2.6, B.6.2.20 and [MULPI], Annex A.2, MAC Service IDs.

B.6.2.2 **Signal Processing Requirements**

See Section 6.2.2.

B.6.2.3 **Modulation Formats**

See Section 6.2.3.

B.6.2.4 **R-S Encode**

B.6.2.4.1 **R-S Encode Modes**

See Section 6.2.4.1.

B.6.2.4.2 **R-S Bit-to-Symbol Ordering**

See Section 6.2.4.2.

B.6.2.5 **Upstream R-S Frame Structure for DOCSIS 3.0 Multiple Transmit Channel mode Enabled**

See Section 6.2.5.

B.6.2.5.1 *R-S Codeword Length*

See Section 6.2.5.1.

B.6.2.5.1.1 *Burst Size*

See Section 6.2.5.1.1.

B.6.2.5.1.2 *Fixed Codeword Length*

See Section 6.2.5.1.2.

B.6.2.5.1.3 *Shortened Last Codeword*

See Section 6.2.5.1.3.

B.6.2.5.2 *R-S FEC Disabled*

See Section 6.2.5.2.

B.6.2.6 *Upstream R-S Frame Structure for DOCSIS 3.0 Multiple Transmit Channel mode Not Enabled*

See Section 6.2.6.

B.6.2.7 *TDMA Byte Interleaver*

See Section 6.2.7.

B.6.2.7.1 *Byte Interleaver Parameters*

See Section 6.2.7.1.

B.6.2.7.2 *Interleaver Operating Modes*

See Section 6.2.7.2.

B.6.2.7.2.1 *Fixed Mode*

See Section 6.2.7.2.1.

B.6.2.7.2.2 *Dynamic Mode*

See Section 6.2.7.2.2.

B.6.2.8 *Scrambler (Randomizer)*

See Section 6.2.8.

B.6.2.9 *TCM Encoder*

See Section 6.2.9.

B.6.2.9.1 *Byte to TCM Symbol Mapping*

See Section 6.2.9.1.

B.6.2.10 Preamble Prepend

See Section 6.2.10.

B.6.2.11 Modulation Rates

See Section 6.2.11.

B.6.2.11.1 DOCSIS 3.0 Modulation Rates

See Section 6.2.11.1.

B.6.2.11.2 Backward Compatibility Modulation Rates

See Section 6.2.11.2.

B.6.2.12 S-CDMA Framer and Interleaver**B.6.2.12.1 S-CDMA Framing Considerations**

See Section 6.2.12.1.

B.6.2.12.2 Mini-slot Numbering

See Section 6.2.12.2.

B.6.2.12.2.1 Mini-slot Numbering Parameters in UCD

See Section 6.2.12.2.1.

B.6.2.12.2.2 B.6.2.12.2.2 Mini-slot Numbering Examples

See Section 6.2.12.2.2.

B.6.2.12.3 Transmission Time

See Section 6.2.12.3.

B.6.2.12.4 Latency Considerations

See Section 6.2.12.4.

B.6.2.12.5 Spreader-off Bursts for Maintenance on S-CDMA channel

See Section 6.2.12.5.

B.6.2.12.6 Limiting the Number of Codes Assigned to a CM

See Section 6.2.12.6.

B.6.2.13 S-CDMA Framer

See Section 6.2.13.

B.6.2.13.1 Subframe Definition

See Section 6.2.13.1.

B.6.2.13.2 Framing Operation

See Section 6.2.13.2.

B.6.2.13.2.1 Rules for Preamble and Coded TCM Symbols

See Section B.6.2.13.2.1.

B.6.2.13.2.2 Rules for Uncoded Symbols and the Uncoded TCM Subsymbols

See Section 6.2.13.2.2.

B.6.2.13.2.3 Subframe Example

See Section 6.2.13.2.3.

B.6.2.13.2.4 Frame Transmission

See Section 6.2.13.2.4.

B.6.2.14 Symbol Mapping

See Section 6.2.14.

B.6.2.15 S-CDMA Spreader

See Section 6.2.15.

B.6.2.15.1 Code Hopping

See Section 6.2.15.1.

B.6.2.15.1.1 Code Hopping Mode 1

See Section 6.2.15.1.1.

B.6.2.15.1.2 Code Hopping Mode 2

See Section 6.2.15.1.2.

B.6.2.15.1.3 Code Hopping Generator

See Section 6.2.15.1.3

B.6.2.16 Transmit Pre-Equalizer

A transmit pre-equalizer of a linear equalizer structure, as shown in Figure B-1, **MUST** be configured for each upstream channel used by the CM in response to the Ranging Response (RNG-RSP) message transmitted by the CMTS.

There are two modes of operation for the pre-equalizer of a CM: DOCSIS 1.1 mode, and DOCSIS 2.0 pre-equalization mode: In DOCSIS 1.1 mode, the CM **MUST** support a (T)-spaced equalizer structure, with 8 taps. The pre-equalizer **MAY** have 1, 2, or 4 samples per symbol, with a tap length longer than 8 symbols: In DOCSIS 1.1 pre-equalization mode, for backwards compatibility, the CMTS **MAY** support fractionally spaced equalizer format

(T/2 and T/4). In DOCSIS 2.0 pre-equalization mode, the pre-equalizer MUST support a symbol (T)-spaced equalizer structure with 24 taps.

In DOCSIS 1.x-only logical channels (type 1 logical channels [[MULPI]), the CM and the CMTS MUST use DOCSIS 1.1 pre-equalization mode.

In DOCSIS 2.0-only or DOCSIS 3.0-only logical channels (type 3 or type 4 logical channels [[MULPI]), the CM and the CMTS MUST use DOCSIS 2.0 pre-equalization mode.

In DOCSIS 1.x/2.0 mixed logical channels (type 2 logical channels [[MULPI]), the CM and the CMTS MUST use DOCSIS 1.1 pre-equalization mode from initial ranging until DOCSIS 2.0 mode is enabled or until Multiple Transmit Channel mode is enabled in the registration process (if either is enabled). The CM and CMTS MUST use DOCSIS 2.0 pre-equalization mode after DOCSIS 2.0 is enabled for the CM. In the case that a CM is placed into Multiple Transmit Channel mode, the CM and CMTS use DOCSIS 2.0 pre-equalization for all bursts on type 2 logical channels whether the burst descriptor in the UCD for the transmitted burst is a Type 4 or Type 5 burst descriptor.

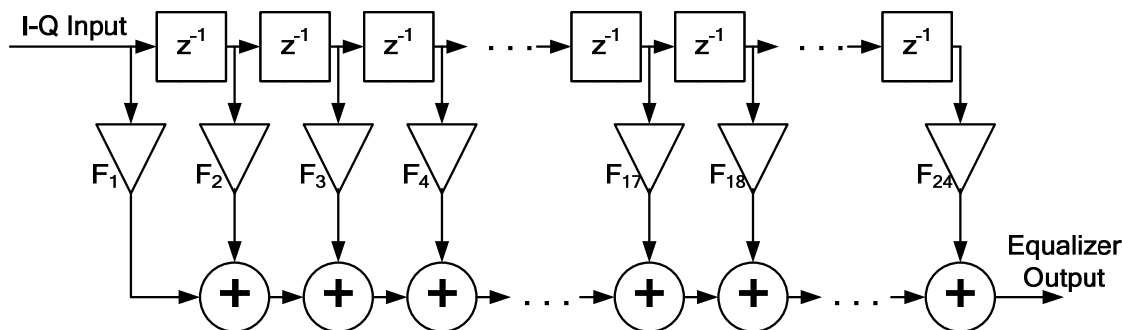


Figure B-1 - Transmit Pre-Equalizer Structure

The RNG-RSP MAC message carries the CM transmit equalization information, and may instruct the CM to either convolve the equalizer coefficients, or (in DOCSIS 2.0 pre-equalization mode only) load them directly. When the CM is instructed to convolve the transmit equalizer coefficients, it MUST convolve the coefficients sent by the CMTS in the RNG-RSP with the existing coefficients, to get the new coefficients. After convolving, the CM MUST truncate the convolution result, such that 24 taps (8 taps in DOCSIS 1.1 pre-equalization mode) remain after the truncation, with the main tap located at the tap designated by the last RNG-RSP received by the CM. The operation of the convolution is formulated by the following equation:

$$F_n^{m+1} = \sum_{k=\max(1-L^{m+1}, n+L^m-L^{m+1}-24)}^{\min(24-L^{m+1}, n+L^m-L^{m+1}-1)} F_{n-k+L^m-L^{m+1}}^m \cdot \hat{F}_{k+L^{m+1}}, \quad n = 1 \dots 24$$

where:

F_n^m are the coefficients prior to the convolution;

F_n^{m+1} are the coefficients after the convolution;

\hat{F}_n are the coefficients sent from the CMTS;

L^m is the main tap location prior to the convolution;

L^{m+1} is the main tap location after the convolution as dictated by the CMTS.

In DOCSIS 2.0 pre-equalization mode, the CMTS MAY instruct the CM to load the transmit pre-equalizer coefficients. When the CM is instructed to load the transmit equalizer coefficients, it MUST load the coefficients sent by the CMTS into the pre-equalizer coefficients after proper normalization, if necessary.

In DOCSIS 1.x-only logical channels, in response to an initial ranging request and periodic ranging requests prior to CM registration, when the CMTS sends the pre-equalizer coefficients, the CMTS MUST compute and send them with an equalizer length of 8 and in T-spaced format, where T is the modulation interval. After registration, the CMTS MAY use a fractionally spaced equalizer format (T/2- or T/4-spaced), with a longer tap length, to match the CM pre-equalizer capabilities that the CMTS learned from the REG-REQ message modem capabilities field.

In DOCSIS 2.0-only or DOCSIS 3.0-only logical channels, the CMTS MUST compute and send the pre-equalizer coefficients with an equalizer length of 24, and in T-spaced format, at all times.

In DOCSIS 1.x/2.0 mixed logical channels, in response to an initial ranging request and periodic ranging requests prior to CM registration, when the CMTS sends the pre-equalizer coefficients, the CMTS MUST compute and send them with an equalizer length of 8 and in T-spaced format. After registration, if the DOCSIS 1.1 pre-equalization mode is enabled, the CMTS MAY use a fractionally spaced equalizer format (T/2- or T/4-spaced), with a longer tap length, to match the CM pre-equalizer capabilities that the CMTS learned from the REG-REQ message modem capabilities field. If DOCSIS 2.0 pre-equalization mode or Multiple Transmit Channel mode is enabled for the CM, the CMTS MUST use a T-spaced equalizer structure with 24 taps. If the first update of the pre-equalizer after the activation of DOCSIS 2.0 pre-equalization mode uses "convolve" mode, the CM MUST zero-pad the existing 8-tap filter to a 24-tap filter, and convolve, according to the rules above.

Prior to making an initial ranging request, and whenever the upstream channel frequency or upstream channel modulation rate changes, the CM MUST initialize the coefficients of the pre-equalizer to a default setting in which all coefficients are zero, except the real coefficient of the first tap (i.e., F1). Whenever the main location is changed, the CM, not the CMTS, MUST compensate for the delay (ranging offset), due to a shift from the previous main tap location, to a new main tap location of the equalizer coefficients, sent by the CMTS (in both "convolve" and "load" operations). The pre-equalizer coefficients are then updated through the subsequent ranging process (unicast initial ranging and periodic ranging).

In DOCSIS 1.1 pre-equalization mode, the CMTS MUST NOT move the main tap location during periodic ranging.

In DOCSIS 1.1 pre-equalization mode, the CMTS MUST NOT instruct the CM to load the transmit equalizer coefficients.

In DOCSIS 2.0 pre-equalization mode, the CMTS MAY move the main tap location during unicast initial ranging or periodic ranging.

Equalizer coefficients may be included in every RNG-RSP message, but typically they only occur when the CMTS determines that the channel response has significantly changed. The frequency of equalizer coefficient updates in the RNG-RSP message is determined by the CMTS.

The CM MUST normalize the transmit equalizer coefficients in order to guarantee proper operation (such as not to overflow or clip). The CM MUST NOT change its target transmit power due to gain or loss of the new coefficients in both "convolve" and "load" operations. The target power is defined in Section B.6.2.19.

In DOCSIS 1.1 mode, if the CM equalizer structure implements the same number of coefficients as assigned in the RNG-RSP message, then the CM MUST NOT change the location of the main tap in the RNG-RSP message. If the CM equalizer structure implements a different number of coefficients than defined in the RNG-RSP message, the CM MAY shift the location of the main tap value. The CM MUST adjust its ranging offset, in addition to any adjustment in the RNG-RSP message, by an amount that compensates for the movement of the main tap location.

B.6.2.17 Spectral Shaping

See Section 6.2.17

B.6.2.17.1 Upstream Frequency Agility and Range

The CM MUST support operation over the frequency range of 5-65 MHz, edge to edge.

Offset frequency resolution MUST be supported per Table B-9.

B.6.2.17.2 Spectrum Format

See Section 6.2.17.2.

B.6.2.18 Relative Processing Delays

See Section 6.2.18.

B.6.2.19 Transmit Power Requirements

The following requirements apply with Multiple Transmit Channel mode enabled. Requirements with Multiple Transmit Channel mode disabled are addressed in Section B.6.2.19.4 below.

The CM MUST 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 [MULPI], Section 6.4.5. Such adjustments by the CM MUST be within the ranges of tolerances described below. A CM MUST confirm that the transmit power per channel limits are met after a RNG-RSP is received or after a UCD change, for each of the CM's active channels, which is referenced. An active channel for a CM is defined as any channel for which the CM has been given a grant, which it will use for channel initialization, or ranging, OR any channel for which the CM is "ranged." The set of "active channels" is also called the Transmit Channel Set. (There will be some mechanism by which a CMTS can command a CM to deactivate an active channel, for the purpose of either reducing the number of active channels on that CM or for the purpose of adding a different active channel as a replacement upstream channel for that CM.) 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) transmitted in the data symbols of a burst, assuming equal likely QAM symbols, measured at the F-connector of the CM. Total transmit power is defined as the sum of the transmit power per channel of each channel transmitting a

burst at a given time. Maximum and minimum transmit power per channel level requirements refer to the CM's target transmit power per channel levels, defined as the CM's estimate of its actual transmit power per channel. The actual transmitted power per channel MUST be within 2 dB of the target power. The target transmit power per channel MUST be variable over the range specified in Table B-10.

With Multiple Transmit Channel mode enabled, let $P_{load} = P_{hi} - P_r$, for each channel, using the definitions for P_{hi} and P_r in the following subsections of B.6.2.19. The channel corresponding to the minimum value of P_{load} is called the highest loaded channel, and its value is denoted as P_{load_1} even if there is only one channel in the Transmit Channel Set. Note that the highest loaded channel is not necessarily the channel with the highest transmit power. Similarly, the channel with the second lowest value of P_{load} is denoted as the second highest loaded channel, and its loading value is denoted as P_{load_2} . The channel with the third lowest value of P_{load} is the third highest loaded channel, and its loading value is denoted as P_{load_3} ; and the channel with the fourth lowest loading value is denoted as the fourth highest loaded channel, and its loading value is P_{load_4} . $P_{load_min_set}$ defines the upper end of the Dynamic Range Window Dynamic Range Window for the CM with respect to P_{hi} for each channel. $P_{load_min_set}$ will limit the maximum power possible for each active channel to a value less than P_{hi} when $P_{load_min_set}$ is greater than zero. $P_{load_min_set}$ is a value commanded to the CM from the CMTS. (In the case of fewer than 4 active channels in the Transmit Channel Set, P_{load_n} will only be valid for the n active upstream channels.) $P_{load_min_set}$, P_{load_n} , P_{hi_n} , P_{r_n} etc. are defined only when Multiple Transmit Channel mode is enabled.

CMTS commands adjusting P_{r_n} MUST be consistent with the $P_{load_min_set}$ previously assigned to the CM and with the following limits:

$$P_{load_min_set} \leq P_{hi_n} - P_{r_n} \leq P_{load_min_set} + 12 \text{ dB}$$

A value is computed, P_{low_multi} , which sets the lower end of the transmit power Dynamic Range Window given the upper end of the range (which is determined by $P_{load_min_set}$).

$$P_{low_multi} = P_{hi} - P_{load_min_set} - 12 \text{ dB}.$$

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

When the CMTS sends a new value of $P_{load_min_set}$ 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 $P_{load_min_set}$ 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_{r_n} may not be implemented immediately upon reception at the CM if the n th channel is bursting. Commanded changes to P_{r_n} may occur simultaneously with the command to change $P_{load_min_set}$. The CMTS SHOULD NOT issue a change in $P_{load_min_set}$ after commanding a change in P_{r_n} until after also providing a sufficient reconfiguration time on the n th channel. The CMTS SHOULD NOT issue a change in $P_{load_min_set}$ after commanding a prior change in $P_{load_min_set}$ until after also providing a global reconfiguration time for the first command. Also, the CMTS SHOULD NOT issue a change in P_{r_n} a) until after providing a global reconfiguration time following a command for a new value of $P_{load_min_set}$, and b) until after providing a sufficient reconfiguration time on the n th channel after issuing a previous change in P_{r_n} . In other words, the CMTS is to avoid sending consecutive changes in P_{r_n} and/or $P_{load_min_set}$ to the CM without a sufficient reconfiguration time for instituting the first command. When a concurrent new value of $P_{load_min_set}$ and change in P_{r_n} are commanded, the CM MAY wait to apply the change in P_{r_n} at the next global reconfiguration time (i.e., concurrent with the institution of the new value of $P_{load_min_set}$) rather than applying the change at the first sufficient reconfiguration time of the n th channel; the value of $P_{load_min_set}$ which applies to the new P_{r_n} is the concurrently commanded $P_{load_min_set}$ value, and thus if the change to P_{r_n} falls outside the Dynamic Range Window of the old $P_{load_min_set}$, then the CM MUST wait for the global reconfiguration time to apply the change in P_{r_n} .

The CM MUST ignore the command to increase the per channel transmit power if such a command would cause P_{load} for that channel to drop below $P_{load_min_set}$. 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 $P_{load_min_set}$ to a smaller value than normal fluctuation of the power per channel in the highest loaded channel is expected to wander. The CM also MUST ignore any commanded change of per channel transmit power which would result in P_r falling below the Dynamic Range Window, determined by $P_{load_min_set}$. The CM also MUST ignore a command changing $P_{load_min_set}$ such that existing values of P_r would fall outside the new Dynamic Range Window.

The spurious performance requirements of Sections B.6.2.22.1, B.6.2.22.1.1, and B.6.2.22.1.2 apply when the CM is operating within certain ranges of values for P_{load_n} , for $n = 1$ to the active number of upstream channels, and for certain ranges of the number of spreading codes being transmitted as a fraction of the total number of active codes on a channel for S-CDMA channels, as detailed in those sections.

Transmit power per channel, for each channel, as reported by the CM in the MIB, is referenced to the 64 QAM constellation. When transmitting with other constellations, a slightly different transmit power will result, depending on the constellation gain in Table B-4.

Table B-5, and Table B-6 below. As an example, if the reported power in a channel is 30 dBmV, 64 QAM will be transmitted with a target power of 30 dBmV in that channel, while QPSK will be transmitted with 28.82 dBmV in that channel.

Table B-4 - Constellation Gains and Per Channel Power Limits with One Channel in the Transmit Channel Set

Constellation	Constellation Gain G_{const} Relative to 64 QAM (dB)	P_{min} (dBmV)			P_{max} (dBmV) TDMA	P_{max} (dBmV) S-CDMA	$P_{min} - G_{const}$ (dBmV)	$P_{max} - G_{const}$ (dBmV) TDMA	$P_{max} - G_{const}$ (dBmV) S-CDMA
		L	M	H					
QPSK	-1.18	17	20	23	61	56	18.18	62.18	57.18
8 QAM	-0.21	17	20	23	58	56	17.21	58.21	56.21
16 QAM	-0.21	17	20	23	58	56	17.21	58.21	56.21
32 QAM	0.00	17	20	23	57	56	17.00	57.00	56.00
64 QAM	0.00	17	20	23	57	56	17.00	57.00	56.00
128 QAM	0.05	17	20	23	N/A	56	16.95	N/A	55.95

(P_{min} is a function of Modulation Rate, with L = 1280 kHz, M = 2560 kHz, and H = 5120 kHz.)

Table B-5 - Constellation Gains and Per Channel Power Limits with Two Channels in the Transmit Channel Set

Constellation	Constellation Gain G_{const} Relative to 64 QAM (dB)	P_{min} (dBmV)			P_{max} (dBmV) TDMA	P_{max} (dBmV) S-CDMA	$P_{min} - G_{const}$ (dBmV)	$P_{max} - G_{const}$ (dBmV) TDMA	$P_{max} - G_{const}$ (dBmV) S-CDMA
		L	M	H					
QPSK	-1.18	17	20	23	58	53	18.18	59.18	54.18
8 QAM	-0.21	17	20	23	55	53	17.21	55.21	53.21
16 QAM	-0.21	17	20	23	55	53	17.21	55.21	53.21
32 QAM	0.00	17	20	23	54	53	17.00	54.00	53.00
64 QAM	0.00	17	20	23	54	53	17.00	54.00	53.00
128 QAM	0.05	17	20	23	N/A	53	16.95	N/A	52.95

(P_{min} is a function of Modulation Rate, with L = 1280 kHz, M = 2560 kHz, and H = 5120 kHz.)

Table B-6 - Constellation Gains and Per Channel Power Limits with Three or Four Channels in the Transmit Channel Set

Constellation	Constellation Gain G_{const} Relative to 64 QAM (dB)	P_{min} (dBmV)			P_{max} (dBmV) TDMA	P_{max} (dBmV) S-CDMA	$P_{min} - G_{const}$ (dBmV)	$P_{max} - G_{const}$ (dBmV) TDMA	$P_{max} - G_{const}$ (dBmV) S-CDMA
		L	M	H					
QPSK	-1.18	17	20	23	55	53	18.18	56.18	54.18
8 QAM	-0.21	17	20	23	52	53	17.21	52.21	53.21
16 QAM	-0.21	17	20	23	52	53	17.21	52.21	53.21
32 QAM	0.00	17	20	23	51	53	17.00	51.00	53.00
64 QAM	0.00	17	20	23	51	53	17.00	51.00	53.00
128 QAM	0.05	17	20	23	N/A	53	16.95	N/A	52.95

(P_{min} is a function of Modulation Rate, with L = 1280 kHz, M = 2560 kHz, and H = 5120 kHz.)

The actual transmitted power per channel, within a burst, MUST 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, theoretically present, due to QAM amplitude modulation, pulse shaping, pre-equalization, and for S-CDMA, spreading and varying number of allocated codes.

The CM MUST support the transmit power calculations defined in Section B.6.2.19.1 and Section B.6.2.19.2.

B.6.2.19.1 TDMA Transmit Power Calculations

In TDMA mode, the CM determines its target transmit power per channel P_t as follows, for each channel which is active. Define for each active channel, for example, upstream channel X:

- P_r = Reported power level (dBmV) of CM in MIB (refers to 64 QAM constellation) for channel X
- P = Power level adjustment (dB); for example, as commanded in ranging response message
- G_{const} = Constellation gain (dB) relative to 64 QAM constellation per Section B.6.2.19
- P_{min} = Minimum target transmit power per channel permitted for the CM per Section B.6.2.19
- P_{max} = Maximum target transmit power permitted for the CM per Section B.6.2.19
- $P_{hi} = \min(P_{max} - G_{const})$ over all burst profiles used by the CM in channel X per Section B.6.2.19
- $P_{low} = \max(P_{min} - G_{const})$ over all burst profiles used by the CM in channel X per Section B.6.2.19
- P_t = Target transmit power per channel level (dBmV) of CM in channel X (actual transmitted power per channel in channel X as estimated by CM)

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

1. $P_s = P_r + P$ //Add power level adjustment (for each channel) to reported power level for each channel
2. $P_r = \min[P_r, P_{hi}]$ //Clip at max power limit per channel
3. $P_r = \max[P_r, P_{low}]$ //Clip at min power limit per channel
Test, IF [$P_r < P_{low_multi}$], ignore command and restore previous P_r //Power per channel from this command would violate the set Dynamic Range Window
4. Test, IF [$P_r > P_{hi} - P_{load_min_set}$], ignore command and restore previous P_r //Power per channel from this command would violate the set Dynamic Range Window

The CM then transmits, in channel X with target power per channel $P_t = P_r + G_{const}$, i.e., the reported power, plus the constellation gain.

Usually the reported power level is a relatively constant quantity, while the transmitted power level in channel X varies dynamically as different burst profiles, with different constellation gains, are transmitted. A CM's target transmit power per channel MUST never be below P_{\min} or above P_{\max} . This implies that in some cases, the extreme transmit power levels (e.g., 61 dBmV for QPSK and 17 dBmV) may not be permitted if burst profiles with multiple constellations are active. Also, if only QPSK is used, the reported power per channel may be greater than 61 dBmV, although the target transmit power per channel will not exceed 61 dBmV.

For example, if with two channels active in channel X, only QPSK and 64 QAM burst profiles are active, $P_{hi} = 54$ dBmV and $P_{low} = 18.2$ dBmV for 1280 kHz modulation rate. P_{low} is dependent on modulation rate. The maximum permitted QPSK transmitted power in channel X is 54 dBmV - 1.2 dB = 52.8 dBmV, the minimum QPSK power in channel X is 18.2 dBmV - 1.2 dB = 17 dBmV (with 1280 kHz modulation rate), the maximum 64 QAM power in channel X is 54 dBmV, and the minimum 64 QAM power in channel X is 18.2 dBmV (with 1280 kHz modulation rate).

B.6.2.19.2 S-CDMA Transmit Power Calculations

See Section 6.2.19.2.

B.6.2.19.2.1 S-CDMA Transmit power Calculations with Maximum Scheduled Codes Not Enabled

In S-CDMA mode when Maximum Scheduled Codes is not enabled, the CM determines its target transmit power per channel P_t as follows, for each channel which is active. Define for each active channel, for example, upstream channel X:

P_r = reported power level (dBmV) of CM in MIB (refers to 64 QAM constellation and all active codes transmitted) for channel X

$P_{hi} = \min[P_{\max} - G_{\text{const}}]$ over all burst profiles used by the CM in channel X per Section B.6.2.19

$P_{low} = \max[P_{\min} - G_{\text{const}}] + 10 \log(\text{number_active_codes} / \text{number_of_codes_per_mini_slot})$ where the maximum is over all burst profiles used by the CM in channel X per Section B.6.2.19

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

1. $P_r = P_r + P$ //Add power level adjustment (for each channel) to reported power level for each channel
2. $P_r = \min[P_r, P_{hi}]$ //Clip at max power limit per channel
3. $P_r = \max[P_r, P_{low}]$ //Clip at min power limit per channel
4. Test, IF [$P_r < P_{low_multi}$], ignore command and restore previous P_r //Power per channel from this command would violate the set Dynamic Range Window
5. Test, IF [$P_r > P_{hi} \cdot P_{load_min_set}$], ignore command and restore previous P_r //Power per channel from this command would violate the set Dynamic Range Window

In a spreader-on frame, the CM then transmits each code i with target power:

$$P_{t,i} = P_r + G_{\text{const},i} - 10 \log(\text{number_active_codes})$$

(i.e., the reported power for channel X, plus the constellation gain $G_{\text{const},i}$ of that code, less a factor taking into account the number of active codes. The total transmit power in channel X, P_t , in a frame is the sum of the individual transmit powers $P_{t,i}$ of each code in channel X, where the sum is performed using absolute power quantities [non-dB domain]).

In a spreader-off frame, the CM target transmit power in channel X is $P_t = P_r + G_{\text{const}}$.

The transmitted power level in channel X varies dynamically as the number of allocated codes varies, and as different burst profiles, with different constellation gains, are transmitted. A CM's target transmit power per channel MUST never be below P_{min} or above P_{max} , including over all numbers of allocated codes and all burst profiles. This implies that in some cases the extreme transmit power levels per channel (e.g., 17 and 56 dBmV) may not be permitted. Also, if, only QPSK is used, the reported power in a channel may be greater than 56 dBmV, although the target transmit power per channel will not exceed 56 dBmV.

If, for example, the Transmit Channel Set contains only one channel, and this channel, channel X, has QPSK and 64 QAM burst profiles, the number of active codes is 128 and the number of codes per mini-slot is 2, then $P_{\text{hi}} = 56$ dBmV and $P_{\text{low}} = 36.24$ dBmV for 1280 kHz modulation rate. P_{low} is dependent on modulation rate. The maximum permitted QPSK transmitted power in channel X is $56 - 1.18 = 54.82$ dBmV when all active codes are transmitted. The minimum QPSK power in channel X is 36.24 dBmV $- 1.18$ dB $- 10\log(128)$ dB $+ 10\log(2)$ dB = 17 dBmV (with 1280 kHz modulation rate), when one mini-slot is transmitted. The last term in the sum is the result of summing the individual powers over 2 codes. Similarly, the maximum 64 QAM power in channel X is 56 dBmV, when all active codes are transmitted and the minimum 64 QAM power in channel X is 36.24 dBmV $- 10\log(128)$ dB $+ 10\log(2)$ dB = 18.18 dBmV (with 1280 kHz modulation rate) when one mini-slot is transmitted. The minimum QPSK power permitted in channel X while transmitting, for example, 2 mini-slots is 20 dBmV (with 1280 kHz modulation rate), and the minimum 64 QAM power permitted in channel X, while transmitting 2 mini-slots, is 21.2 dBmV (with 1280 kHz modulation rate).

The CM needs to implement some form of clipping on the transmitted waveform at the higher output powers in order to prevent peak to average ratio (PAR) issues.

The power received at the CMTS in a spreader-on frame will sometimes be less than the nominal power of a spreader-off frame because of such factors as: 1) broadcast opportunities not used by any CM; 2) unicast grants not used by one or more CMs; or 3) mini-slots assigned to the NULL SID.

B.6.2.19.2.2 S-CDMA Transmit Power Calculations With Maximum Scheduled Codes Enabled.

In S-CDMA mode on channels on which Maximum Scheduled Codes is enabled, the CM determines its target transmit power per channel P_t as follows, for each channel which is active. Define for each upstream channel, for example upstream channel X:

P_r = reported power level (dBmV) of CM in MIB (operational transmit power of the spreader-off ranging burst referenced to 64 QAM modulation) for channel X

$P_{\text{hi}_S} = \min[53 - G_{\text{const}}]$ over all spreader-on burst profiles used by the CM in channel X per Section B.6.2.19

$P_{\text{low}_S} = \max[17 - G_{\text{const}}] + 10 \log(\text{number_active_codes} / \text{number_of_codes_per_mini-slot})$ where the maximum is over all burst profiles used by the CM in channel X per Section B.6.2.19

P_{max_T} = Maximum target transmit power permitted for the CM in channel X in TDMA mode per Section B.6.2.19 for the constellation used in ranging.

$P_{\text{hi}_T} = \min[P_{\text{max}_T} - G_{\text{const}}]$ over all spreader-off burst profiles used by the CM in channel X per Section B.6.2.19.

$P_{\text{on}} = P_r$ clipped at the maximum spreader-on limit.

P_{sf} = CM Power Shortfall.

P_{hr} = S-CDMA Power Headroom in dB. Equivalent to the value of the Ranging Response Message TLV-11 divided by 4.

P = power level adjustment in dB sent from CMTS to CM for channel X

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

1. $P_r = P_r + P$ //Add power level adjustment (for each channel) to reported power level for each channel.
2. $P_r = \min[P_r, P_{hi_r}]$ //Clip at max TDMA power limit per channel.
3. $P_r = \max[P_r, P_{low_s}]$ //Clip at min S-CDMA power limit per channel.
4. Test, IF [$P_r < P_{low_multi}$], ignore command and restore previous P_r //Power per channel from this command would violate the set Dynamic Range Window.
5. Test, IF [$P_r > P_{hi} - P_{load_min_set}$], ignore command and restore previous P_r //Power per channel from this command would violate the set Dynamic Range Window.
6. $P_{on} = \min[P_r, P_{hi_s} - P_{load_min_set}]$ //Clip at max S-CDMA power limit per channel for the multiple channel dynamic range limitation, as set by $P_{load_min_set}$.

In spreader-off frames, the CM transmits in channel off with target power:

$$P_t = P_r + G_{const}$$

Based on the spreader-off transmit power in channel X, the CM updates its power shortfall in channel X according to the following steps:

with only one channel in the Transmit Channel Set

$$P_{sf} = P_t - (56 - P_{load_min_set}) \quad // \text{Difference between spreader-off and max spreader-on target powers in channel X}$$

$$P_{sf} = \max[P_{sf}, 0] \quad // \text{Set } P_{sf} \text{ to 0 if } P_t \text{ is less than } 56 \text{ dBmV} - P_{load_min_set} \text{ in channel X;}$$

and with more than one channel in the Transmit Channel Set

$$P_{sf} = P_t - (53 - P_{load_min_set}) \quad // \text{Difference between spreader-off and max spreader-on target powers in channel X}$$

$$P_{sf} = \max[P_{sf}, 0] \quad // \text{Set } P_{sf} \text{ to 0 if } P_t \text{ is less than } 53 \text{ dBmV} - P_{load_min_set} \text{ in channel X}$$

In spreader-on frames, the CM transmits each code i with target power

$$P_{t,i} = P_{on} + G_{const, i} - 10 \log(\text{number_active_codes}) + P_{hr}$$

i.e., the clipped reported power in channel X plus the constellation gain $G_{const, i}$ of that code, less a factor taking into account the number of active codes, plus the Power Headroom P_{hr} . P_{hr} is the power (in dB) added to account for CMs that have a maximum scheduled code limit and can transmit additional power per code. The total transmit power in channel X, P_t , in a frame is the sum of the individual transmit powers $P_{t,i}$ of each code in channel X, where the sum is performed over all N_{alloc} allocated codes using absolute power quantities (non-dB domain).

$$P_t = 10 \log \sum_{i=1}^{N_{alloc}} 10^{P_{t,i}/10}$$

If, for example, with more than one channel in the Transmit Channel Set, and the burst profile in channel X contains QPSK for IUCs 1, 2, 3, and 4 and 64 QAM for IUCs 9 and 10, the number of active codes is 128, and the number of codes per mini-slot is 2, then $P_{hi_S} = 53$ dBmV, $P_{low_S} = 36.24$ dBmV, and $P_{hi_T} = 58$ dBmV. Assume $P_{load_min_set} = 0$ dB. Assume the CM ranges channel X at spreader-off target transmit power of 57 dBmV. The CM reports for channel X $P_{sf} = 57$ dBmV - 53 dBmV = 4 dB. The CMTS uses P_{sf} to set for channel X (using its vendor-specific algorithm) $max_scheduled_codes = 32$ and $P_{hr} = 6$ dB. (The S-CDMA power headroom may differ from the power shortfall, at the discretion of the CMTS.) The CM sets its transmitted power per code in channel X to:

$$\begin{aligned} P_{t,i} &= P_{on} + G_{const,i} - 10 \log(\text{number_active_codes}) + P_{hr} \\ &= 53 \text{ dBmV} + 0 \text{ dB} - 21 \text{ dB} + 6 \text{ dB} \quad // \text{ For a code with 64 QAM modulation in channel X} \\ &= 38 \text{ dBmV} \end{aligned}$$

A parameter that may be used to illustrate the effect of increased power per code in channel X is the Effective Transmit Power, P_{eff} , the power that would result in channel X hypothetically if all N_a active codes were transmitted. It is computed as:

$$\begin{aligned} P_{eff} &= 10 \log \sum_{i=1}^{N_{act}} 10^{P_{t,i}/10} \\ &= P_{on} + P_{hr} + 10 \log \frac{1}{N_{act}} \sum_{i=1}^{N_{act}} 10^{G_{const,i}/10} \end{aligned}$$

where the last term is the average constellation gain.

For a reference case with all codes transmitted in channel X using 64 QAM modulation ($G_{const} = 0$ dB), the effective transmit power reduces to:

$$P_{eff} = P_{on} + P_{hr}$$

Continuing the above example, the result is:

$$\begin{aligned} P_{eff} &= 53 \text{ dBmV} + 6 \text{ dB} \\ &= 59 \text{ dBmV} \end{aligned}$$

Limiting the number of codes has given the CM an enhanced effective power of 59 dBmV, which is 6 dB above the normal maximum power per channel of 53 dBmV, and 2 dB above the ranging power of 57 dBmV. In this example, the CMTS used its discretion to ask for 2 dB more enhancement in channel X than was needed ($P_{hr} = 6$ dB vs $P_{sf} = 4$ dB), perhaps due to some known impairment in the channel.

The $effective_SNR$ is an SNR estimate for a given code corresponding to the effective transmit power in channel X. It is defined as the measured SNR at the last station maintenance, minus the CM power shortfall, plus the power headroom, plus the difference in constellation gain between the ranging burst and the code under consideration. Its equation is:

$$effective_SNR = measured_SNR - P_{sf} + P_{hr} + (G_{const,i} - G_{const,ranging})$$

Where $G_{const,ranging}$ is the constellation gain of the ranging burst in channel X that resulted in the SNR measurement. In the MIB, $effective_SNR$ corresponds to a reference case with 64 QAM modulation ($G_{const,i} = 0$ dB):

$$effective_SNR = measured_SNR - P_{sf} + P_{hr} - G_{const,ranging}$$

Continuing the example, if the measured SNR in channel X in the last station maintenance was 17 dB, using QPSK modulation ($G_{\text{const, ranging}} = -1.2$ dB), then the effective SNR referenced to 64 QAM modulation is:

$$\text{effective_SNR} = 17\text{dB} - 4\text{dB} + 6\text{dB} + 1.2 \text{ dB} = 20.2 \text{ dB}$$

B.6.2.19.3 Transmit Power Step Size

See Section 6.2.19.3.

B.6.2.19.4 Transmit Power Requirements for Multiple Transmit Channel mode Not Enabled

With Multiple Transmit Channel mode disabled, the CM MUST operate as in [DOCSIS RFI 2.0], Section F.6.2.18, with the exceptions that the minimum upstream power P_{min} MUST be as given in Table B-7 below, and the maximum upstream power P_{max} MUST be as given in Table B-4 above. Note that in Section F.6.2.18.2.2 of [DOCSIS RFI 2.0], the value of 53 in the equation for $P_{\text{hi_s}}$ MUST be replaced with 56 since this value is actually intended to be P_{max} from Table B-4 above.

Table B-7 - Minimum Transmit Power P_{min} for Multiple Transmit Channel mode Not Enabled

Modulation Rate (kHz)	P_{min} (dBmV)	Applicability
160	17	CM MAY support
320	17	CM MAY support
640	17	CM MAY support
1 280	17	CM MUST support
2 560	20	CM MUST support
5 120	23	CM MUST support

B.6.2.20 Burst Profiles

The transmission characteristics are separated into three portions: a) Channel Parameters; b) Burst Profile Attributes; and c) User Unique Parameters. The Channel Parameters include: a) the modulation rate (six rates from 160 kHz to 5120 kHz in octave steps); b) the center frequency (Hz); c) the 1536-bit Preamble Superstring; and d) the S-CDMA channel parameters. The Channel Parameters are shared by all users on a given channel. The Burst Profile Attributes are listed in Table B-8, and these parameters are the shared attributes corresponding to a burst type.

The CM MUST generate each burst at the appropriate time, as conveyed in the mini-slot grants provided by the CMTS MAPs.

The CM MUST support all burst profiles commanded by the CMTS, via the Burst Descriptors in the UCD, and subsequently assigned for transmission in a MAP.

Table B-8 - Burst Profile Attributes

Burst Profile Attributes	Configuration Settings
Modulation	QPSK, 8 QAM, 16 QAM, 32 QAM, 64 QAM, 128 QAM(TCM Only)
Differential Encoding	On/Off
TCM Encoding	On/Off
Preamble Length	0-1536 bits (note Section B.6.2.19, B.6.2.9)
Preamble Value offset	0 to 1534
R-S FEC Error Correction (T)	0 to 16 (0 implies no R-S FEC. The number of codeword parity bytes is 2*T)
R-S FEC Codeword Information Bytes (k)	Fixed: 16 to 253 (assuming R-S FEC on) Shortened: 16 to 253 (assuming R-S FEC on)
Scrambler Seed	15 bits
Maximum Burst Length (mini-slots) ¹	0 to 255
Guard Time	4 to 255 modulation intervals There is no guard time in S-CDMA
Last Codeword Length	Fixed, shortened
Scrambler On/Off	On/Off
Byte Interleaver Depth (I_r) ²	0 to floor(2048/ N_r) ³
Byte Interleaver Block Size (B_r) ⁴	2* N_r to 2048
Preamble Type	QPSK0/QPSK1
S-CDMA Spreader ⁵	On/Off
S-CDMA Codes per Subframe ⁵	1 to 128
S-CDMA Interleaver Step ⁵	1 to (spreading intervals per frame - 1)
S-CDMA Selection Mode for Active Codes and Code Hopping	Mode 1 or Mode 2
S-CDMA Selection String for Active Codes	128-bit string indicating which codes are active in Active Code Selection Mode 2
<p>1. A burst length of 0 mini-slots in the Channel Profile means that the burst length is variable on that channel for that burst type. The burst length, while not fixed, is granted explicitly by the CMTS to the CM in the MAP.</p> <p>2. If depth=1, no interleaving; if depth=0, dynamic mode.</p> <p>3. N_r is the R-S codeword size $k+2T$ as defined in Section B.6.2.5.1.</p> <p>4. Used only in dynamic mode</p> <p>5. Used only for S-CDMA channels.</p>	

The User Unique Parameters may vary for each user even when using the same burst type on the same channel as another user (for example, Power Level), and are listed in Table B-9.

Table B-9 - User Unique Burst Parameters

User Unique Parameter	Adjustment Command	Resulting Parameter Value
Power Level ¹	8-bit two's complement, resolution = 0.25 dB	TDMA: +17 to +54 dBmV (32 QAM, 64 QAM) +17 to +55 dBmV (8 QAM, 16 QAM) +17 to +58 dBmV (QPSK) S-CDMA: +17 to +53 dBmV (all modulations) Resolution = 0.25 dB or better
Offset Frequency ²	Range = ± 32 kHz, resolution = 1 Hz	Frequency Range per Section B.6.2.17.1 Upstream Frequency Agility and Range
Ranging Offset	Integer part: 32-bit two's complement, resolution = $(1 / 10.24 \text{ MHz}) = 6.25 \mu\text{s}/64 = 97.65625 \text{ ns}$ Fractional part: unsigned 8-bit fractional extension, resolution = $6.25 \mu\text{s}/(64*256) = 0.3814697265625 \text{ ns}$	Range: sufficient for maximum cable plant length per Section B.1.2.1 Broadband Access Network Resolution: TDMA: $6.25 \mu\text{s}/64$. S-CDMA: $6.25 \mu\text{s}/(64*256)$
Burst Length (mini-slots) if variable on this channel (changes burst-to-burst)	N/A	1 to 255 mini-slots
Transmit Equalizer Coefficients (See Section B.6.2.16, Transmit Pre-Equalizer)	DOCSIS 3.0 & 2.0 modes: 24 complex coefficients, 4 bytes per coefficient (2 real and 2 imaginary), load and convolve modes DOCSIS 1.1 mode: up to 64 complex coefficients, 4 bytes per coefficient (2 real and 2 imaginary), convolve mode only	DOCSIS 3.0 & 2.0 modes: 24 complex coefficients DOCSIS 1.1 mode: up to 64 complex coefficients
<p>1. The minimum power level limit depends on modulation rate. The maximum power level limit depends on modulation order and the number of upstream channels.</p> <p>2. The CM MUST implement the Offset Frequency Adjustment to effect a change in upstream carrier frequency within ± 10 Hz of the commanded change.</p>		

B.6.2.20.1 Ranging Offset

See Section 6.2.20.1.

B.6.2.20.2 TDMA Reconfiguration Times

The CM MUST be capable of switching burst profiles with no reconfiguration time required between bursts, except for changes in the following parameters: 1) Output Power; 2) Symbol Rate; 3) Offset Frequency; 4) Channel Frequency; 5) Ranging Offset; and 6) Pre-Equalizer Coefficients. In the case of Multiple Transmit Channel mode

enabled, all of these parameters are adjusted independently on each upstream channel and reconfiguration time on the non-adjusted channels is not required by the CM.

In the case of Multiple Transmit Channel mode enabled, global reconfiguration time (simultaneous on all of the active channels) could be required by the CM when the Dynamic Range Window is adjusted (new value of $P_{load_min_set}$), or when any active channel accumulates a transmit power change of more than 3 dB since the last global reconfiguration time (see Section B.6.2.22.1). The CM MUST meet the requirements, on each active channel, for "Dynamic Range Window change" (below) whenever $P_{load_min_set}$ is changed or re-commanded.

For **Dynamic Range Window** change or re-command: If $P_{load_min_set}$ is changed, or re-commanded, the CM MUST be able to implement the change between bursts, as long as the CMTS allocates at least 96 symbols, plus 10 μ s, between the last symbol center of one burst (on any active channel) and the first symbol center of the following burst (on any active channel -- same channel or different channel from the previous burst). (Global "quiet" across all active channels requires the intersection of ungranted burst intervals across all active TDMA channels to be at least 96 symbols plus 10 μ s duration.) The Output Power of the CM MUST be settled to within ± 0.1 dB of its final output power level within 10 μ s from the beginning of the change in $P_{load_min_set}$, and from the beginning of any Output Power change on any channel that accompanies the change or re-command of $P_{load_min_set}$. The Dynamic Range Window, i.e., $P_{load_min_set}$, MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS. The Dynamic Range Window MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted.

For **Output Power** changes: If Output Power is to be changed by 1 dB or less, the CM MUST be able to implement the change between bursts as long as the CMTS allocates at least 96 symbols, plus 5 μ s between the last symbol center of one burst and the first symbol center of the following burst. If Output Power is to be changed by more than 1 dB, the CM MUST be able to implement the change between bursts, as long as the CMTS allocates at least 96 symbols, plus 10 μ s, between the last symbol center of one burst and the first symbol center of the following burst. The maximum reconfiguration time of 96 symbols should compensate for the ramp-down time of one burst and the ramp-up time of the next burst, as well as the overall transmitter delay time, including the pipeline delay and pre-equalizer delay. The Output Power of the CM MUST be settled to within ± 0.1 dB of its final output power level; a) within 5 μ s from the beginning of a change of 1 dB or less; and b) within 10 μ s from the beginning of a change of greater than 1 dB. Output Power MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS. Output Power MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted.

For **Symbol Rate** changes: The CM MUST be able to transmit consecutive bursts as long as the CMTS allows the required time between bursts for UCD parameter changes. Symbol Rate MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS. Symbol Rate MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted.

For **Offset Frequency** changes: The CM MUST be able to transmit consecutive bursts as long as the CMTS allocates at least 96 symbols in between the last symbol center of one burst and the first symbol center of the following burst. The maximum reconfiguration time of 96 symbols should compensate for the ramp-down time of one burst and the ramp-up time of the next burst, as well as the overall transmitter delay time, including the pipeline delay and optional pre-equalizer delay. Offset frequency MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS. Offset frequency MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted.

For **Channel Frequency** changes: The CM MUST be able to implement the change between bursts as long as the CMTS allocates at least 96 symbols plus 100 ms between the last symbol center of one burst and the first symbol of the following burst. The Channel Frequency of the CM MUST be settled within the phase noise and accuracy requirements of Section B.6.2.22.5 and Section B.6.2.22.6 within 100 ms from the beginning of the change. Channel Frequency MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS.

Channel Frequency MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted.

For **Ranging Offset** and/or **Pre-Equalizer Coefficient** changes: The CM MUST be able to transmit consecutive bursts as long as the CMTS allocates at least 96 symbols in between the last symbol center of one burst and the first symbol center of the following burst. The maximum reconfiguration time of 96 symbols should compensate for the ramp-down time of one burst and the ramp-up time of the next burst, as well as the overall transmitter delay time, including the pipeline delay and pre-equalizer delay. Ranging Offset MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS. Ranging Offset MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted.

For **Modulation Type** changes: The CM MUST be able to transmit consecutive bursts with no reconfiguration time between them (except for the minimum guard time). The modulation MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted, EXCLUDING the effect of the transmit equalizer (if present in the CM). (This is to be verified with the transmit equalizer providing no filtering; delay only. Note that if the CMTS has decision feedback in its equalizer, it may need to provide more than the 96 symbol gap between bursts of different modulation type which the same CM may use; this is a CMTS decision.)

B.6.2.20.3 S-CDMA Reconfiguration Times

In S-CDMA mode, for changes in Output Power per mini-slot, Offset Frequency, Pre-equalizer coefficients, and/or Ranging Offset, the CM MUST be able to transmit consecutive bursts as long as the CMTS allocates the time duration of at least one frame in between the bursts. For all other burst profile parameter changes, no reconfiguration is required beyond what is provided by the MAC for such changes. In the case of Multiple Transmit Channel mode enabled, all of these parameters are adjusted independently on each upstream channel and reconfiguration time on the non-adjusted channels is not required by the CM.

In the case of Multiple Transmit Channel mode enabled, global reconfiguration time could be required by the CM when the Dynamic Range Window is adjusted (new value of $P_{load_min_set}$), or when any active channel accumulates a transmit power change of more than 3 dB since the last global reconfiguration time (see Section B.6.2.22.1). (Global "quiet" across all active channels requires the intersection of ungranted burst intervals across all active S-CDMA channels to be at least one frame.) The CM MUST be able to transmit consecutive bursts, even with a change or re-command of $P_{load_min_set}$, as long as the CMTS allocates at least one frame in between bursts, across all channels in the Transmit Channel Set. (From the end of a burst on one channel, to the beginning of the next burst, on any channel, must be at least one frame duration, to provide a "global reconfiguration time" for S-CDMA channels.)

B.6.2.20.4 CM Timing Offsets When Changing Modulation Rate

See Section 6.2.20.4.

B.6.2.21 Burst Timing Convention

See Section 6.2.21.

B.6.2.22 Fidelity Requirements

See Section 6.2.22.

B.6.2.22.1 Spurious Emissions

The noise and spurious power MUST NOT exceed the levels given in Table B-10, Table B-11, and Table B-12. The noise and spurious power SHOULD NOT exceed the levels given in Table B-13 and Table B-14. When

Multiple Transmit Channel mode is enabled, the noise and spurious power requirements for the individual channels are summed (absolute power, NOT in dB) to determine the composite noise floor for the multi-channel transmission condition. Also 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 Transmit Channel Set, and for the ratio of $number_allocated_codes / number_active_codes$ from 1 down to 1/4th, for S-CDMA channels.

For S-CDMA mode, when a modem is transmitting fewer than 1/4th of the active spreading codes on a channel, the spurious emissions requirement limit is the power value (in dBmV), corresponding to the specifications for the power level associated with 1/4th of the active spreading codes (i.e., 6 dB below the channel power corresponding to all active codes for that channel allocated to the CM).

When Multiple Transmit Channel mode is enabled and there are two or more channels in the Transmit Channel Set (TCS), the spurious performance requirements MUST be met only when the following loading ranges are satisfied:

Two or more channels in the TCS: applicable when $P_{load_1} - P_{load_min_set} \leq 3$ dB, AND

Two channels in the TCS: applicable when $P_{load_2} - P_{load_1} \leq 12$ dB.

Three channels in the TCS: applicable when $P_{load_3} - P_{load_1} \leq 8$ dB.

Four channels in the TCS: applicable when:

$$P_{load_2} - P_{load_1} \leq 4 \text{ dB.}$$

$$\text{AND} \quad P_{load_3} - P_{load_2} \leq 4 \text{ dB.}$$

$$\text{AND} \quad P_{load_4} - P_{load_3} \leq 4 \text{ dB.}$$

When a modem is transmitting at power levels falling outside these loading ranges (which is allowed by the 12 dB dynamic range), the spurious emissions requirement limits are the power values (in dBmV rather than dBc), corresponding to the specifications associated with transmissions at the second and higher loaded channels, if the loading of those channels were decreased (transmit power increased) to just meet the applicability ranges.

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. "Global reconfiguration time" is defined as the inactive time interval provided between active transmissions which simultaneously satisfies the requirement in Section B.6.2.20.2 for all TDMA channels in the TCS and the requirement in Section B.6.2.20.3 for all S-CDMA channels in the TCS. With Multiple Transmit Channel mode enabled the CMTS SHOULD provide global reconfiguration time to a CM before the CM has been commanded to change any active upstream channel transmit power by ± 3 dB cumulative since its last global reconfiguration time.

As one example, with three S-CDMA channels in the TCS, at transmit powers per channel of 53 dBmV, 41 dBmV, and 41 dBmV, the absolute spurious emissions requirements, when two or more channels in the TCS are transmitting, correspond to the absolute spurious emissions requirements with transmit power per channel of 53 dBmV, 45 dBmV, and 45 dBmV. The 41 dBmV power levels are raised for the purposes of determining the spurious emissions limits, up to $53 \text{ dBmV} - 8 = 45 \text{ dBmV}$.

As an example with four S-CDMA channels in the TCS, at transmit powers per channel of 53 dBmV, 41 dBmV, 41 dBmV, and 41 dBmV, the absolute spurious emissions requirements, when two or more channels in the TCS are transmitting, correspond to the absolute spurious emissions requirements with transmit power per channel of 53 dBmV, 49 dBmV, 45 dBmV, and 41 dBmV. The 41 dBmV power levels are raised for the purposes of determining the spurious emissions limits, up to $53 \text{ dBmV} - 4 \text{ dB} = 49 \text{ dBmV}$; and the artificial $49 \text{ dBmV} - 4 \text{ dB} = 45 \text{ dBmV}$; and the artificial $45 \text{ dBmV} - 4 \text{ dB} = 41 \text{ dBmV}$.

In Table B–10, inband spurious 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 is equal to the modulation rate (e.g., 1280 to 5120 ksym/s). All requirements expressed in dBc are relative to the actual transmit power that the CM emits in one channel.

The measurement bandwidth is 160 kHz for the Between Bursts (none of the channels in the TCS are bursting) specs of Table B–10, except where called out as 4 MHz or 250 kHz.

The Transmitting Burst specs apply during the mini-slots granted to the CM (when the CM uses all or a portion of the grant), and for 32 modulation intervals before and after the granted mini-slots. The Between Bursts specs apply except during a used grant of mini-slots on any active channel for the CM, and the 32 modulation intervals before and after the used grant.

In TDMA mode, a mini-slot may be as short as 32 modulation intervals, or 6.25 μ s at the 5.12 Msym/s rate, or as short as 25 μ s at the 1.28 Msym/s rate.

Table B–10 - Spurious Emissions

Parameter	Transmitting Burst	Between Bursts
Inband	-40 dBc	-72 dBc
Adjacent Band	See Table B–11 and Table B–12	-72 dBc
Within the upstream operating range of 5-65 MHz (excluding assigned channel, adjacent channels, and carrier-related channels)	See Table B–13 and Table B–14	-72 dBc
For the case where the upstream operating range is 5-65 MHz: CM Integrated Spurious Emission Limits (all in 250 kHz, includes discretetes) 87.5 to 108 MHz	-30 dBmV	-59 dBmV
For the case where the upstream operating range is 5-65 MHz: CM Integrated Spurious Emissions Limits (all in 4.75 MHz, includes discretetes) ¹ 65 to 87.5 MHz 108 to 136 MHz ² 126 to 862 MHz	max(-40 dBc, -26 dBmV) -40 dBmV -45 dBmV	-26 dBmV -45 dBmV max(-45 dBmV, -40 dB ref d/s ³)
For the case where the upstream operating range is 5-65 MHz: CM Discrete Spurious Emissions Limits ¹ 65 to 87.5 MHz 108 to 862 MHz	max(-50 dBc, -36 dBmV) -50 dBmV	-36 dBmV -50 dBmV
1. These spec limits exclude a single discrete spur related to the tuned received channel; this single discrete spur MUST NOT be greater than -40 dBmV. 2. The frequencies from 108 to 136 MHz may be forbidden due to national regulations. 3. "dB ref d/s" is relative to the received downstream signal level. Some spurious outputs are proportional to the receive signal level.		

B.6.2.22.1.1 Adjacent Channel Spurious Emissions

Spurious emissions from a transmitted carrier may occur in an adjacent channel which could be occupied by a carrier of the same or different modulation rate. Table B–11 lists the required adjacent channel spurious emission levels for all combinations of transmitted carrier modulation rates and adjacent channel modulation rates. The measurement is performed in an adjacent channel interval of appropriate bandwidth and distance from the transmitted carrier, based on the modulation rates of the transmitted carrier, and the carrier in the adjacent channel.

Table B–11 - Adjacent Channel Spurious Emissions Requirements Relative to the Per Channel Transmitted Burst Power Level for Each Channel

Specification in the interval	Measurement interval and distance from carrier edge	Adjacent channel carrier modulation rate
-50 dBc	20 kHz to 180 kHz	160 kHz
-50 dBc	40 kHz to 360 kHz	320 kHz
-50 dBc	80 kHz to 720 kHz	640 kHz
-50 dBc	160 kHz to 1440 kHz	1280 kHz
-47 dBc	320 kHz to 2880 kHz	2560 kHz
-44 dBc	640 kHz to 5760 kHz	5120 kHz

In addition, the CM SHOULD meet the requirements in Table B–12, under the same conditions as listed for Table B–11 above. With a lower power channel bursting as a lone bursting channel (at that point in time), 2 dB relaxation will apply to the absolute spurious emissions requirement of Table B–12.

Table B–12 - Adjacent Channel Spurious Emissions Goals Relative to the Per Channel Transmitted Burst Power Level for Each Channel

Specification in the interval	Measurement interval and distance from carrier edge	Adjacent channel carrier modulation rate
-53 dBc	20 kHz to 180 kHz	160 kHz
-53 dBc	40 kHz to 360 kHz	320 kHz
-52 dBc	80 kHz to 720 kHz	640 kHz
-51 dBc	160 kHz to 1440 kHz	1280 kHz
-50 dBc	320 kHz to 2880 kHz	2560 kHz
-48 dBc	640 kHz to 5760 kHz	5120 kHz

B.6.2.22.1.2 Spurious Emissions in the Upstream Frequency Range

Table B–13 lists the possible modulation rates that could be transmitted in an interval, the required spurious level in that interval, and the initial measurement interval at which to start measuring the spurious emissions. Measurements should start at the initial distance and be repeated at increasing distance from the carrier until the upstream band edge is reached. The lower band edge of the upstream is 5 MHz; the upper band edge of the upstream is 65 MHz.

Table B–13 - Spurious Emissions Requirements in the Upstream Operating Frequency Range Relative to the Per Channel Transmitted Burst Power Level for Each Channel

Possible modulation rate in this interval	Specification in the interval	Initial measurement interval and distance from carrier edge
160 kHz	-54 dBc	220 kHz to 380 kHz
320 kHz	-52 dBc	240 kHz to 560 kHz
640 kHz	-50 dBc	280 kHz to 920 kHz
1280 kHz	-50 dBc	360 kHz to 1640 kHz
2560 kHz	-47 dBc	520 kHz to 3080 kHz
5120 kHz	-44 dBc	840 kHz to 5960 kHz

In addition, the CM SHOULD meet the requirements in Table B–14, under the same conditions as listed for Table B–13 above. With a lower power channel bursting as a lone bursting channel (at that point in time), 2 dB relaxation will apply to the absolute spurious emissions requirement of Table B–14.

Table B–14 - Spurious Emissions Goals in the Upstream Operating Frequency Range Relative to the Per Channel Transmitted Burst Power Level for Each Channel

Possible modulation rate in this interval	Specification in the interval	Initial measurement interval and distance from carrier edge
160 kHz	-60 dBc	220 kHz to 380 kHz
320 kHz	-58 dBc	240 kHz to 560 kHz
640 kHz	-56 dBc	280 kHz to 920 kHz
1280 kHz	-54 dBc	360 kHz to 1640 kHz
2560 kHz	-52 dBc	520 kHz to 3080 kHz
5120 kHz	-50 dBc	840 kHz to 5960 kHz

B.6.2.22.2 Spurious Emissions During Burst On/Off Transients

See Section 6.2.22.2.

B.6.2.22.3 Modulation Error Ratio (MER)

See Section 6.2.22.3.

B.6.2.22.3.1 Definitions

See Section 6.2.22.3.1.

B.6.2.22.3.2 Requirements

Unless otherwise stated, the MER MUST meet or exceed the following limits over the full transmit power range of Table B–9 for each modulation, modulation rate, and over the full carrier frequency range, and for S-CDMA, over any valid number of active and allocated codes. The 5-65 MHz carrier frequency range refers more precisely to $[5 \text{ MHz} + \text{modulation rate} * 1.25 / 2]$ to $[65 \text{ MHz} - \text{modulation rate} * 1.25 / 2]$. At the break points between regions, the higher MER specification applies.

Case 1: Flat channel, transmit equalization OFF

Case 1a: for modulation rates 2.56 Msym/s and below for 5 MHz to 65 MHz mode of operation:

- MER_{symb} ≥ 30 dB over 15 to 47 MHz carrier frequency
- MER_{symb} ≥ 27 dB over 10 MHz to 15 MHz and 47 MHz to 54 MHz carrier frequency
- MER_{symb} ≥ 26 dB over 5 MHz to 10 MHz and 54 MHz to 65 MHz carrier frequency

Case 1b: for modulation rate 5.12 Msym/s for 5 MHz to 65 MHz mode of operation:

- MER_{symb} ≥ 27 dB over 15 to 47 MHz carrier frequency
- MER_{symb} ≥ 24 dB over 10 MHz to 15 MHz and 47 MHz to 54 MHz carrier frequency
- MER_{symb} ≥ 23 dB over 5 MHz to 10 MHz and 54 MHz to 65 MHz carrier frequency

Case 2: Flat channel, transmit equalization ON

Case 2a: for TDMA/QPSK, $MER_{\text{symb}} \geq 30$ dB.

Case 2b: for S-CDMA and all TDMA modulations except QPSK, $MER_{\text{symb}} \geq 35$ dB.

Case 2c: for S-CDMA, $MER_{\text{chip}} \geq 33$ dB.

Case 3: Echo channel, transmit equalization ON (Echo channel MUST not be limited by dynamic range and MUST be a true representation described in Table B-3)

Case 3a: In the presence of a single echo selected from the channel micro-reflections defined in Table B-3, the measured MER_{symb} MUST be ≥ 30 dB for TDMA/QPSK, and ≥ 33 dB for S-CDMA and all TDMA modulations, except QPSK.

Case 3b: In the presence of two or three of the echoes defined in Table B-3 (at most one of each specified magnitude and delay), the measured MER_{symb} MUST be ≥ 29 dB.

Since the table does not bound echo delay for the -30 dBc case, for testing purposes, it is assumed that the time span of the echo at this magnitude is less than or equal to 1.5 μs .

The CMTS MUST provide a test mode in which it:

- Accepts equalizer coefficients via an external interface, (e.g., Ethernet),
- Sends the coefficients to the CM's pre-equalizer via ranging response message (both set and convolve modes),
- Does not adjust the CM's frequency, timing or power.

B.6.2.22.4 Filter Distortion

See Section 6.2.22.4.

B.6.2.22.4.1 Amplitude

See Section 6.2.22.4.1.

B.6.2.22.5 Carrier Phase Noise

See Section 6.2.22.5.

B.6.2.22.6 Channel Frequency Accuracy

See Section 6.2.22.6.

B.6.2.22.7 Modulation Rate Accuracy

See Section 6.2.22.7.

B.6.2.22.8 Modulation Timing Jitter**B.6.2.22.8.1 Symbol Timing Jitter for TDMA Upstream Channels in Asynchronous Operation**

See Section 6.2.22.8.1.

B.6.2.22.8.2 Modulation Timing Jitter for Synchronous Operation

See Section 6.2.22.8.2.

B.6.2.23 Upstream Demodulator Input Power Characteristics

See Section 6.2.23.

B.6.2.24 Upstream Electrical Output from the CM

The CM MUST output an RF modulated signal with the characteristics delineated in Table B–15.

Table B–15 - Single Channel Electrical Output from CM

Parameter	Value
Frequency	5 to 65 MHz edge to edge
Level range per channel (Multiple Transmit Channel mode disabled, or Multiple Transmit Channel mode enabled with only one channel in the TCS)	TDMA: P_{min} to +57 dBmV (32 QAM, 64 QAM) P_{min} to +58 dBmV (8 QAM, 16 QAM) P_{min} to +61 dBmV (QPSK) S-CDMA: P_{min} to +56 dBmV (all modulations) where P_{min} = +17 dBmV, 1280 ksym/s modulation rate P_{min} = +20 dBmV, 2560 ksym/s modulation rate P_{min} = +23 dBmV, 5120 ksym/s modulation rate
Level range per channel (two channels in the TCS)	TDMA: P_{min} to +54 dBmV (32 QAM, 64 QAM) P_{min} to +55 dBmV (8 QAM, 16 QAM) P_{min} to +58 dBmV (QPSK) S-CDMA: P_{min} to +53 dBmV (all modulations) where P_{min} = +17 dBmV, 1280 ksym/s modulation rate P_{min} = +20 dBmV, 2560 ksym/s modulation rate P_{min} = +23 dBmV, 5120 ksym/s modulation rate
Level range per channel (three or four channels in the TCS)	TDMA: P_{min} to +51 dBmV (32 QAM, 64 QAM) P_{min} to +52 dBmV (8 QAM, 16 QAM) P_{min} to +55 dBmV (QPSK) S-CDMA: P_{min} to +53 dBmV (all modulations) where P_{min} = +17 dBmV, 1280 ksym/s modulation rate P_{min} = +20 dBmV, 2560 ksym/s modulation rate P_{min} = +23 dBmV, 5120 ksym/s modulation rate
Modulation Type	QPSK, 8 QAM, 16 QAM, 32 QAM, 64 QAM, and 128 QAM

Parameter	Value
Modulation Rate (nominal)	TDMA: 1280, 2560 and 5120 kHz S-CDMA: 1280, 2560 and 5120 kHz Optional pre-3.0 DOCSIS operation: TDMA: 160, 320, 640 kHz
Channel Bandwidth	TDMA: 1600, 3200 and 6400 kHz S-CDMA: 1600, 3200 and 6400 kHz Optional pre-3.0 DOCSIS operation: TDMA: 200, 400, 800 kHz
Output impedance	75 ohms
Output Return Loss	> 6 dB (5 to 65 MHz) > 6 dB (108 to 862 MHz)
Connector	F connector per [ISO/IEC-61169-24] (common with the input)

B.6.2.25 Upstream CM Transmitter Capabilities

The CM advertises its capabilities to the CMTS. These capabilities include:

- Maximum number of active channels, which MUST be 4 or greater.
- Maximum number of 3.2 MHz channels, which MUST be no greater than the maximum number of active channels supported.
- Maximum number of 6.4 MHz channels, which MUST be no greater than the maximum number of 3.2 MHz channels supported, and MUST be 4 or greater.
- Selective Active Code capability – yes/no.

The CM MUST support all the capability it reports to the CMTS.

The CMTS MUST NOT command the CM to operate outside its advertised capabilities.

B.6.2.25.1 Description of CM Upstream Transmit Channel Set Capabilities

See Section 6.2.25.1

B.6.3 Downstream

B.6.3.1 Downstream Protocol and Interleaving Support

The downstream PMD sublayer MUST conform to [EN 300 429]. The downstream RF output electrical requirements including downstream frequency plan, interleaver depths, spectrum format, clock, and symbol requirements are defined in Annex A of the [DRFI] specification. The CM MUST support the interleaver depths defined in Table A-1 of [DRFI].

B.6.3.2 Downstream Electrical Input to CM

The CM MUST be able to accept any number of signals between one to MDBC simultaneously, located in frequency intervals that complies with the CMs TB list and Demod List. A CM MUST be capable of being reconfigured to receive different channels. At a minimum, the CM MUST be able to receive at least 4 downstream channels, located independently within an arbitrary 64 MHz window in the downstream frequency band. The CM

MUST support bonding of any number of downstream channels up to its maximum. The CM MUST be able to accept RF modulated signals with the characteristics defined in Table B–16.

Table B–16 - Electrical Input to CM

Parameter	Value
Center Frequency	112 to 858 MHz \pm 30 kHz
Level Range (one DOCSIS channel)	-17 dBmV to +13 dBmV for 64 QAM -13 dBmV to +17 dBmV for 256 QAM
Modulation Type	64 QAM and 256 QAM
Symbol Rate (nominal)	6.952 Msym/s (64 QAM and 256 QAM)
Bandwidth	8 MHz (alpha = 0.15 square-root raised cosine shaping for 64 QAM and 256 QAM)
Total Input Power (80-862 MHz)	< 33 dBmV
Maximum power of any carrier input to CM	20 dBmV
Input (load) Impedance	75 ohms
Input Return Loss	> 6 dB (5 to 65 MHz) > 6 dB (108 to 862 MHz)
Connector	F connector per [ISO/IEC-61169-24] (common with the output)

B.6.3.3 CM BER Performance

The bit-error-rate performance of a CM MUST be as described in this section. The requirements apply to each individual downstream received channel with the interleaver set to the I = 12, J = 17 mode of interleaving.

B.6.3.3.1 64 QAM

B.6.3.3.1.1 64 QAM CM BER Performance

Implementation loss of the CM MUST be such that the CM achieves a post-FEC BER less than or equal to 10^{-8} when operating at a carrier to noise ratio (E_s/N_0) of 25.5 dB or greater. If it is not possible to measure post-FEC BER directly, Codeword Error Rate, R_C (as defined below) may be used. In this case, the CM MUST achieve a Codeword Error Rate of less than or equal to 9×10^{-7} when operating at a carrier to noise ratio (E_s/N_0) of 25.5 dB or greater.

Computation of Codeword Error Rate, R_C :

$$R_C = \frac{(E_U - E_{U0})}{(E_U - E_{U0}) + (E_C - E_{C0}) + (C - C_0)}$$

Where:

E_U is the value of the count of code words with uncorrectable errors,

E_C is the value of the count of code words with correctable errors and

C is the value of the count of code words without errors.

Sample the values at the beginning of the test interval (indicated by the subscripted values, E_{U0} , E_{C0} , and C_0) and the end of the test interval (indicated by the un-subscripted values, E_U , E_C , and C).

B.6.3.3.1.2 64 QAM Image Rejection Performance

Performance as described in Section B.6.3.3.1.1 MUST be met with analog or digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

B.6.3.3.1.3 64 QAM Adjacent Channel Performance

Performance as described in Section B.6.3.3.1.1 MUST be met with a digital signal at 0 dBc in the adjacent channels.

Performance as described in Section B.6.3.3.1.1 MUST be met with an analog signal at +10 dBc in the adjacent channels.

Performance as described in Section B.6.3.3.1.1, with an additional 0.2 dB allowance, MUST be met with a digital signal at +10 dBc in the adjacent channels.

B.6.3.3.2 256 QAM

B.6.3.3.2.1 256 QAM CM BER Performance

Implementation loss of the CM MUST be such that the CM achieves a post-FEC BER less than or equal to 10^{-8} when operating at a carrier to noise ratio (E_s/N_o) as shown below. If it is not possible to measure post-FEC BER directly, Codeword Error Rate, RC (as defined in Section B.6.3.3.1.1) may be used. In this case, the CM MUST achieve a Codeword Error Rate of less than or equal to 9×10^{-7} when operating at a carrier to noise ratio (E_s/N_o), as shown in the following:

Input Receive Signal Level	E_s/N_o
-13 dBmV to -6 dBmV	34.5 dB or greater
-6 dBmV to +17 dBmV	31.5 dB or greater

B.6.3.3.2.2 256 QAM Image Rejection Performance

Performance as described in Section B.6.3.3.2.1 MUST be met with an analog or a digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

B.6.3.3.2.3 256 QAM Adjacent Channel Performance

Performance as described in Section B.6.3.3.2.1 MUST be met with an analog or a digital signal at 0 dBc in the adjacent channels.

Performance as described in Section B.6.3.3.2.1, with an additional 0.5 dB allowance, MUST be met with an analog signal at +10 dBc in the adjacent channels.

Performance as described in Section B.6.3.3.2.1, with an additional 1.0 dB allowance, MUST be met with a digital signal at +10 dBc in the adjacent channels.

B.6.3.4 Downstream Multiple Receiver Capabilities

See Section 6.3.4.

B.6.3.4.1 Receive Module Parameters

See Section 6.3.4.1.

B.6.3.4.2 Receive Channel Parameters

See Section 6.3.4.2.

B.6.3.4.3 Standard Receive Channel Profile

In order to limit the complexity in the CMTS of configuring arbitrarily complicated Receive Channel Profiles, DOCSIS defines a set of "Standard" Receive Channel Profiles that describe a minimal set of constraints.

A CM reports to the CMTS at least one standard RCP as well as its "Manufacturer" RCP that gives more details of its capabilities and constraints. If the CMTS configures the CM with a Receive Channel Configuration based on a Standard RCP, some CM capabilities represented in the Manufacturer RCP but outside the Standard RCP may not be available.

The Standard Receive Channel Profile for the European technology option named "8-DOCSIS-01" is defined for 8 MHz operation as depicted in Figure B–2. The RF input from the cable plant is connected to a Receive Module with a block of 8 adjacent channels. This Receive Module represents the constraints of a tuner with a 64 MHz capture bandwidth that can be placed at any point in the full DOCSIS frequency range. The demodulation of four channels anywhere within the capture bandwidth is represented by the four Receive Channels, each corresponding to the packet stream from a single QAM channel output to the MAC layer.

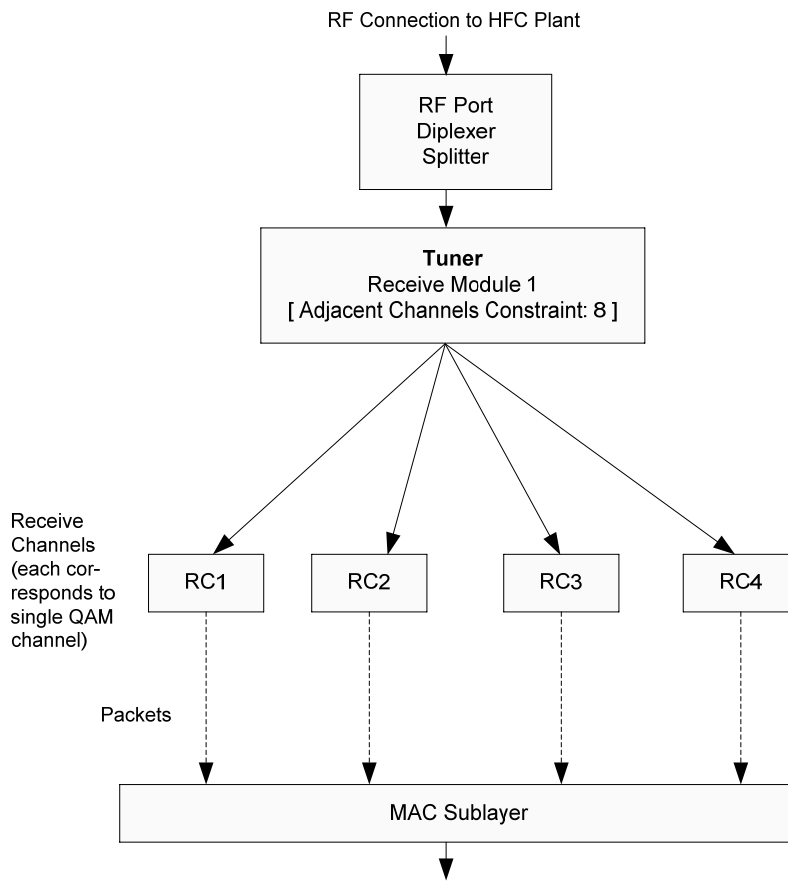


Figure B–2 - 8 MHz DOCSIS Standard Receive Channel Profile 01

For this Standard Receive Channel Profile, the CMTS sends the following parameters to the CM in the Receive Channel Configuration:

- Center frequency of the first channel in the 8-channel Receive Module.
- Center frequency of each Receive Channel within the capture bandwidth of the Receive Module.

B.6.3.4.4 Example Manufacturer Receive Channel Profile

See Section 6.3.4.4.

B.6.3.4.5 Reconfiguration of CM Receive Channels

See Section 6.3.4.5.

B.6.3.5 Non-Synchronous DS Channel Support

See Section 6.3.5.

Annex C MPEG Header Synchronization and Recovery

The downstream bitstream is defined as a continuous series of 188-byte packets. These packets consist of a 4-byte header followed by 184 bytes of payload. This transport stream is specified in MPEG-2 Part 1, Systems [ISO 13818] with modifications described in Section B.4 of [ITU-T J.83-B].

C.1 MPEG Header Synchronization and Recovery in the North American Technology Option

When implementing the first physical layer technology option referred to in Section 1 (1.1 Introduction and Purpose) and specified in Section 6, modifications described in ITU-T J.83 Annex B [ITU-T J.83-B], Section B.4 MPEG-2 transport framing apply to the transport stream format.

The MPEG-2 packet stream SHOULD be declared "in frame" (i.e., correct packet alignment has been achieved) when five consecutive correct parity checksums, each 188 bytes from the previous one, have been received.

The MPEG-2 packet stream SHOULD be declared "out of frame", and a search for correct packet alignment started, when nine consecutive incorrect parity checksums are received.

C.2 MPEG Header Synchronization and Recovery in the European Technology Option

When implementing the second physical layer technology option referred to in Section 1 (1.1 Introduction and Purpose) and specified in Annex B, modifications described in EN 300 429 [EN 300 429] apply to the transport stream format.

The MPEG-2 packet stream SHOULD be declared "in frame" (i.e., correct packet alignment has been achieved) when five consecutive correct sync bytes, each 188 bytes from the previous one, have been received.

The MPEG-2 packet stream SHOULD be declared "out of frame", and a search for correct packet alignment started, when nine consecutive incorrect sync bytes are received.

Appendix I Example Preamble Sequence

I.1 Introduction

A programmable preamble superstring, up to 1536 bits long, is part of the channel-wide profile or attributes, common to all burst profiles on the channel, but with each burst profile able to specify the start location within this sequence of bits and the length of the preamble. The first bit of the Preamble Pattern is designated by the Preamble Value Offset. The first bit of the Preamble Pattern is the first bit into the symbol mapper (Figure 6-2 and Figure 6-3), and is the first symbol of the burst (see Section 6.2.13). As an example, for Preamble Offset Value = 100, the 101st bit of the preamble superstring is the first bit into the symbol mapper, and the 102nd bit is the second bit into the mapper, and is mapped to Q1, and so on. An example 1536-bit-long preamble superstring is given in Appendix I.2.

I.2 Example Preamble Sequence

The following is the example 1536-bit preamble sequence:

Bits 1 through 128:

```
1100 0011 1111 0000 0011 0011 1111 1100 0011 0011 0000 0011 1100 0000 0011 0000
0000 1110 1101 0001 0001 1110 1110 0101 0010 0101 0010 0101 1110 1110 0010 1110
```

Bits 129 through 256:

```
0010 1110 1110 0010 0010 1110 1110 1110 1110 0010 0010 0010 1110 1110 0010
1110 1110 1110 0010 1110 0010 1110 0010 0010 0010 0010 1110 0010 0010 1110 0010
```

Bits 257 through 384:

```
0010 1010 0110 0110 0110 1110 1110 1110 0010 1110 0010 1110 0010 1110 0110 1010
0010 1110 1110 1010 0110 1110 0110 0010 0110 1110 1010 1110 0010 1010 0110 0010
```

Bits 385 through 512:

```
0010 1110 0110 1110 0010 1010 1010 0110 0010 1110 0110 0110 1110 0010 0010 0110
0010 1110 0010 1010 0010 1110 0110 0010 0010 1010 0010 0110 0010 1010 0010 1010
```

Bits 513 through 640:

```
0010 1110 0110 1110 0110 0110 1110 0010 0110 1010 0110 0010 1110 1110 1010 0010
1110 1110 0010 1110 1110 1110 0010 1110 1110 0010 1110 0010 0010 1110 0010 0010
```

Bits 641 through 768:

```
1110 1110 1110 0010 0010 0010 1110 0010 1110 1110 1110 1110 0010 0010 1110 0010
1110 0010 0010 0010 1110 1110 0010 0010 0010 0010 1110 0010 0010 0010 0010 1110
```

Bits 769 through 896:

0011 0000 1111 1100 0000 1100 1111 1111 0000 1100 1100 0000 1111 0000 0000 1100
0000 0000 1111 1111 1111 0011 0011 0011 1100 0011 1100 1111 1100 1111 0011 0000

Bits 897 through 1024:

1100 0011 1111 0000 0011 0011 1111 1100 0011 0011 0000 0011 1100 0000 0011 0000
0000 1110 1101 0001 0001 1110 1110 0101 0010 0101 0010 0101 1110 1110 0010 1110

Bits 1025 through 1152:

0010 1110 1110 0010 0010 1110 1110 1110 1110 1110 0010 0010 0010 1110 1110 0010
1110 1110 1110 0010 1110 0010 1110 0010 0010 0010 0010 1110 0010 0010 1110 0010

Bits 1153 through 1280:

0010 0010 1110 1110 1110 1110 1110 1110 0010 1110 0010 1110 0010 1110 1110 0010
0010 1110 1110 0010 1110 1110 1110 0010 1110 1110 0010 1110 0010 0010 1110 0010

Bits 1281 through 1408

1100 1100 1111 0000 1111 1111 1100 0000 1111 0011 1111 0011 0011 0000 0000 1100
0011 0000 0011 1111 1111 1100 1100 1100 1111 0000 1111 0011 1111 0011 1100 1100

Bits 1409 through 1536:

0011 0000 1111 1100 0000 1100 1111 1111 0000 1100 1100 0000 1111 0000 0000 1100
0000 0000 1111 1111 1111 0011 0011 0011 1100 0011 1100 1111 1100 1111 0011 0000

Appendix II S-CDMA Framing

This appendix is informative. In case of conflict between this section and any normative section of this specification, the normative section takes precedence.

Please note that the pseudo-code in Sections II.1 and II.2 below is specific to the case of a single burst using all spreading codes.

II.1 Coded Subsymbol Numbering

The following code sample contains a short algorithmic description of the operation of the address generator for the coded subsymbols. The address generator for the coded subsymbols fills rows first using the interleaver step size parameter (step in the listing) to step through the spreading intervals within a row. Each step is performed using a modified modulo algorithm which allows the use of interleaver step size and spreading intervals per frame with common divisors. After each row is filled, the next row is begun with the first spreading interval. In the following listings, the index "i" is initialize to the value "1" and coded_col0 is defined as "0".

```
for(row = FIRST_ROW; row ≤ LAST_ROW; row++)
{
    coded_col = 0;
    store_coded( row, coded_col, cod_sym);
    /* Store the coded portion of the symbol (or preamble) to (row, coded_col) */
    for( i = 1; i < framelen; i++ )
    {
        coded_col = coded_col + Interleaver_step_size;

        if( mod( i, framelen / gcd( step, framelen ) ) == 0 )
            coded_col = coded_col + 1; /* gcd is greatest common divisor */

        coded_col = mod( coded_col, framelen );
        store_coded( row, coded_col, cod_sym );
        /* Store the coded portion of the symbol (or preamble) to (row,
        coded_col) */
    }
}
```

II.2 Uncoded Subsymbol Numbering

The following is a short algorithmic description of the operation of the address generator for uncoded subsymbols. The generator fills columns within a subframe first. The row index increments by one for each uncoded subsymbol. At the end of the subframe, the column index is incremented and the row index set to the first row of the subframe. After completing a subframe, the next subframe begins with the next uncoded subsymbol.

```
uncoded_col = 0;
uncoded_row = FIRST_ROW;
while( uncoded_row ≤ LAST_ROW)
{
    if( ( uncoded_row + R ) > LAST_ROW )
        Rprime = LAST_ROW - uncoded_row + 1;
    else
        Rprime = R;

    for( i = 0; i < Rprime; i++)
    {
        /* Check whether (uncoded_row, uncoded_col) is a preamble location.
        * If it is, go to next location */
        if( not_preamble( uncoded_row, uncoded_col ) )
            store_uncoded( uncoded_row, uncoded_col, unc_sym );
        uncoded_row = uncoded_row + 1;
    }
}
```

```

uncoded_row = uncoded_row - Rprime;
uncoded_col = uncoded_col + 1;
if (uncoded_col > framelen)
{
    uncoded_col = 0;
    uncoded_row = uncoded_row + R;
}
}

```

FIRST_ROW and LAST_ROW are, respectively, the first and last row (i.e., code) in each frame spanned by the grant. FIRST_ROW can be between 0 and 127 in the first frame of the allocation and is 0 in any other frames that the grant may span (if any). LAST_ROW can be between 0 to 127 in the last frame of the burst and is 127 for any other frame (if any).

II.3 Framer Output Numbering

The following code sample contains a short algorithmic description of the operation of the address generator for the output symbols. The address generator for the output symbols is used to access both the coded and uncoded subsymbol memories. The output address generator accesses all of the rows (codes) of a spreading interval first followed by subsequent spreading intervals. This algorithmic description is generalized to any allowable number of active codes, where the Number of Active Codes parameter is given by N_a .

```

for( col=0; col < framelen; col++ )
    for( row=(128-Na); row < 128; row++ )
        outsym = get_data( row, col );

```

Note that the Framer row addressing operates independently of the Selectable Active Codes mode. In either Selectable Active Codes Mode 1 or 2, the unused codes are at the lower row address indices of the frame and, therefore, are not used.

II.4 Comments

In the preceding code samples, the number of iterations for the loop is not always correct since an allocation can be less than the number of codes. In II.2, the listing supports the case of a shortened sub-frame.

Appendix III Ambient Temperature and Wind Loading Effects

This appendix discusses possible ambient temperature change and dynamic wind loading effects relevant to operating a system with DOCSIS 2.0 CMs and CMTSs. The intent of this appendix is to describe possible approaches for dealing with these issues. The relationships between the timing variation of the received upstream signal and the rate of change of these ambient plant conditions is discussed. However, measured field data providing the statistics of the ambient conditions used in these relationships is not available, so it is not possible at the time of writing to determine the magnitude or frequency of occurrence of these conditions on operational cable systems. This appendix is not intended to be an exhaustive discussion of either these issues or solutions.

The following issues are discussed in this appendix:

- Synchronization tolerances to plant delay variations;
- Changes in propagation delay due to temperature changes;
- Changes in propagation delay due to wind in the case of aerial cable plant.

III.1 Synchronization Tolerances to Plant Delay Variations

The CMTS receiver synchronization requirements for S-CDMA and Advanced TDMA are identical for the same signal constellation and symbol rates. However, for S-CDMA, burst synchronization is accomplished, to a fine degree, through the ranging process, while for TDMA, burst synchronization is accomplished to a coarse degree, through the ranging process, and then to a fine degree through a receiver burst timing recovery process. In both cases, the degree of timing accuracy required in the receiver is tighter for higher symbol rates and higher-order constellations.

Because S-CDMA requires a fine degree of timing accuracy to be accomplished solely by the ranging process, it is more sensitive to changes in the propagation delay of the cable plant between ranging intervals, which can be as much as 30 s apart. Table III–1 lists plant delay variations that can be accommodated in S-CDMA and TDMA modes for a 1 dB degradation under example conditions.

Table III–1 - Allowable Plant Timing Drift

Constellation	E_s/N_0 for 1e-8 BER(dB)	Allowable peak-peak plant delay variation (ns)S-CDMA mode	Allowable peak-peak plant delay variation (ns)TDMA mode
Fully-coded QPSK	5	90	800
TCM QPSK	9	79	N/A
TCM 8 QAM	12	57	N/A
Uncoded QPSK	15	38	800
Fully-coded 64 QAM	17.7	24	800
TCM 32 QAM	19	18	N/A
Uncoded 16 QAM	22	9	800
Uncoded 32 QAM	25	6	800
TCM 128 QAM	25	6	N/A
Uncoded 64 QAM	28	2	800

Defined conditions:

- 1 dB degradation at $1e-8$ BER;
- Uniform ranging offset over $\pm 1/64$ chip;
- 63 CMs, each with 2 codes;
- E_s/N_o numbers are ideal theoretical values, not including implementation effects;
- 5.12 MHz modulation rate;
- Timing variation over 30 s period;
- TDMA receiver accepts ± 2 symbol coarse timing offset (implementation-dependent).

This channel impairment should be considered along with all of the other upstream channel characteristics highlighted in Table 5-2 or in Table B-3, depending on which of the two technology options is relevant in the particular application (North American or European) as described in Section 1.1.

DOCSIS requires station maintenance at least every 30 s (T4 time out has a maximum value of 35 s). For S-CDMA at a given modulation and symbol rate, if there exists a rapid propagation delay variation such that the resulting delay change cannot be tracked by station maintenance, then one or more of the following performance trade-offs and/or system changes may be enacted: (1) decrease the station maintenance period; (2) decrease the constellation order; (3) decrease the modulation rate; (4) apply additional error correction; (5) apply some combination of 1 through 4, or (6) change the channel to operate in TDMA mode.

The following sections discuss the relationship of temperature changes and wind loading on the propagation delay in coaxial and HFC cable plants.

III.2 Change in Propagation Delay Due to Temperature Changes

III.2.1 Fiber Delay Changes Due to Temperature

In HFC plant design, the number of amplifiers in cascade in the coax portion is kept low in order to keep signal degradation to an acceptable level. As a result, long runs of cable plant are mainly comprised of fiber. A typical value for propagation delay variation due to temperature change of the fiber is 44 ps per km per degree C. The delay variation comes mostly from the change in refractive index of the glass with temperature, not the change in fiber length.

It is assumed that changes in optical cable length due to stretching or expansion will be a negligible factor because optical cables are built to isolate the fiber from stresses in the cable itself. The fiber usually sits loosely in a tube inside the cable and quite a bit of relative movement is possible. This construction allows normal cable handling and aerial deployment without resulting in high stress on the optical fiber.

Assuming 44 ps per km per degree C, any product of optical cable length and temperature change which equals 50, results in approximately a 2 ns change in the fiber propagation delay. For example, 25 km fiber and a 2 degree temperature change will result in a 2 ns change in propagation delay. For the maximum distance between CMTS and CM specified in DOCSIS of 100 miles or roughly 160 km, the temperature change needed for a 2 ns change in one-way propagation delay is 0.3 degree C.

Obviously, the issue is how fast the cable core (where the fiber is) will heat up under ambient temperature changes. For buried or underground cable, there is no issue. For aerial cable solar loading changes should be considered. Black sheathed cable has interior temperatures quite a bit higher than ambient in sunlight, but data is currently unavailable. When the rising sun hits aerial cable on a cold morning, one would expect a temperature change. Similarly, sunlight appearing out of cloud cover may have a similar impact although the size of the shadow of the

cloud moving out of the way has to be considered. The numerical examples above suggest that only long aerial cable runs may have a problem under some combinations of time-of-day and weather.

III.2.2 Coaxial Cable Delay Changes Due to Temperature

The coaxial cable has "blown" foam between the center conductor and the solid shield. The nominal propagation velocity is about 87% of free space velocity. Propagation velocity does not vary markedly with temperature. Given the relatively short lengths of coaxial cable in most HFC plant (a few miles), this seems unlikely to be a significant source of delay variation.

III.2.3 Delay Change Due to Wind

Aerial cable stretches during windy conditions, so it is possible to estimate a propagation delay from the change in length under various wind loads. As mentioned, the construction of optical cable makes it tolerant of stretching, so it is assumed that optical cable stretching, due to wind loading, can be ignored. Wind loading will affect aerial coaxial cables.

Wind loading is difficult to deal with, analytically, because it is unlikely to be uniform along the cable. A delay model derived from a significant body of measured data is needed to investigate this further. Wind loading may be a source of fast delay variation and station maintenance may not occur at intervals small enough for the ranging mechanism to track this variation accurately.

The effects of wind loading on a typical cable was investigated with a publicly available program from a coaxial cable manufacturer. These calculations showed that length changes in the range 0.01% and 0.05% are possible for various amounts of wind loading. This converts to significant propagation delay variation. For example, at 5 miles (8 km) and 0.02% length variation, the change in propagation delay is:

$$(8/3e5)*(1/0.87)*2e-4 \text{ s} = 6 \text{ ns}$$

This is a peak value, but the length of coax is quite short and the wind load is moderate. While the time duration over which this delay variation occurs is unspecified, it may be noted that wind gust data is readily available for most cities, and wind gust will be the primary mechanism for wind-based timing changes on cable plants. For example, in New York City at the time of this writing, wind gusts of up to 40 mph were reported, while average wind speed was about 10 mph. Hence, over a period of 1 to 4 s (the typical wind gust measurement interval), the wind speed changed by 30 mph. Much stronger wind gusts are frequently measured in locations prone to windy conditions.

Appendix IV Description of Upstream Transmit Channel Set Capability: Example Calculations for Reporting and Figuring the Number of Active Channels Supported

As described in the requirements in Section 6.2.25.1, the CM communicates to the CMTS its capabilities in regards to the number of active upstream channels it can support. The method of conveyance must clearly and unambiguously describe the capabilities of the CM to the CMTS.

This appendix provides the definition for the parameters as in Section 6.2.25.1, and also provides examples illustrating the calculations that need to take place at both the CM and the CMTS so that precisely the same possible upstream transmit channel combinations are understood at both.

Let X = number of 6.4 MHz channels, Y = number of 3.2 MHz channels, and Z = number of 1.6 MHz channels that are active at one time (the Transmit Channel Set). Let X_{\max} be the maximum number of 6.4 MHz channels that the CM can support, Y_{\max} be the maximum number of 3.2 MHz channels that the CM can support, and Z_{\max} be the maximum number of 1.6 MHz channels that the CM can support. Z_{\max} is equivalent to the maximum number of upstream transmitters that the CM can support.

The CM reports X_{\max} , Y_{\max} , and Z_{\max} during registration using modem capabilities encodings specified in [MULPI], Annex C, "Modem Capabilities Encoding." It is specified in Section 6.2.25.1 that $X_{\max} \leq Y_{\max} \leq Z_{\max}$.

Communications of X_{\max} , Y_{\max} , and Z_{\max} by a CM provides enough information to describe all the combinations possible for X , Y , and Z .

It is understood that these then directly provide the combinations of (X, Y, Z) which are $(X_{\max}, 0, 0)$; $(0, Y_{\max}, 0)$; and $(0, 0, Z_{\max})$.

From these values the CMTS may calculate $B_{\max} = \max(6.4 * X_{\max}, 3.2 * Y_{\max}, 1.6 * Z_{\max})$.

Then, all values for (X, Y, Z) are allowed which satisfy all inequalities below:

$$6.4 * X + 3.2 * Y + 1.6 * Z \leq B_{\max}$$

$$X \leq X_{\max}$$

$$Y \leq Y_{\max}$$

$$X + Y + Z \leq Z_{\max}$$

Examples:

EX1.

$$X_{\max} = 4; Y_{\max} = 4; Z_{\max} = 4.$$

In this case all combinations with $X + Y + Z \leq 4$ are possible.

EX2.

$$X_{\max} = 4; Y_{\max} = 8; Z_{\max} = 16.$$

In this case we have all combinations where $6.4X + 3.2Y + 1.6Z \leq 25.6$.

This is identical to a receiver which can channelize 25.6 MHz into any combination of the three channel bandwidths without restriction.

EX3.

$$X_{\max} = 4; Y_{\max} = 8; Z_{\max} = 8.$$

This is a subset of Example 2.

With this capability reporting, both conditions a) $6.4X + 3.2Y + 1.6Z \leq 25.6$, and b) $X + Y + Z \leq 8$, must be met.

Thus, all cases in Example 2 with $X + Y + Z > 8$ ARE EXCLUDED with this reporting.

The CM in this case only has to support (and will only be asked to support) a maximum of 8 channels, even if they are all 1.6 MHz bandwidth.

There are no other restrictions on the Transmit Channel Set which must be supported.

EX4.

$$X_{\max} = 4; Y_{\max} = 6; Z_{\max} = 8.$$

This is a subset of Example 2 and a subset of Example 3.

With this capability reporting, both conditions a) $6.4X + 3.2Y + 1.6Z \leq 25.6$, and b) $X + Y + Z \leq 8$, must be met, but in addition, $Y \leq 6$.

Thus, all cases in Example 2 with $X + Y + Z > 8$ ARE EXCLUDED with this reporting, and ALSO EXCLUDED are (0, 8, 0); (0,7,1); {and (0,7,0)}.

The latter cases are EXCLUDED because the cases with $Y > Y_{\max} = 6$ have to be excluded.

Note that (1,6,0); (0,6,2); {and (0, 6, 1) and (0, 6, 0)} are INCLUDED.

Appendix V Description of Upstream Channel Power Control with Multiple Upstream Channels

This appendix provides an explanation of the difference in the power control algorithm with a) multiple upstream channels, compared to b) previous DOCSIS operation. This appendix also explains the extension of some DOCSIS 2.0 parameters and defines new parameters in the upstream power control operation, and also provides examples illustrating the upstream power adjustment and control by the CMTS of CM's which have Multiple Transmit Channel mode enabled.

V.1 DOCSIS 2.0 Parameters Extended to Multiple Transmit Channel mode

The basic parameters from DOCSIS 2.0 such as P_r , P_{hi} , P_{low} , and for S-CDMA with Maximum Scheduled Codes Enabled, P_{hi_T} , P_{low_S} , and P_{on} , maintain their definitions and roles from DOCSIS 2.0. However, these are extended with Multiple Transmit Channel mode enabled so that they are defined for each upstream channel which the CMTS has activated for the CM. These parameters are denoted with a subscript identifying the channel for which they are associated. Therefore, with four upstream channels activated by the CMTS for a particular CM, that CM will have P_{r_1} , P_{r_2} , P_{r_3} , and P_{r_4} for its reported power level (dBmV) for its four upstream channels, respectively.

V.2 New Parameters in DOCSIS 3.0 Upstream Power Control ("loading," P_{load_n} , $P_{load_min_set}$, Dynamic Range Window, P_{low_multi})

In addition to the extension of these DOCSIS 2.0 parameters, there are new parameters introduced in DOCSIS 3.0 upstream power control. Parameters such as $P_{load_min_set}$, P_{load_n} (e.g., P_{load_1} , P_{load_2} , etc.), and P_{low_multi} are added, and new concepts such as Dynamic Range Window and "loading" an upstream digital-to-analog converter with the modulated waveforms to be transmitted upstream are added.

The "loading" of a modulated waveform to be transmitted upstream is a measure of how close to the maximum power level the waveform happens to be (in terms of its average power); this is dependent on the absolute power level of the waveform and of the modulation type of the waveform, and also the number of channels in the Transmit Channel Set. For example, with 64 QAM TDMA modulation the maximum power level allowed is 57 dBmV, 54 dBmV, and 51 dBmV for one, two, and three channels, respectively, in the Transmit Channel Set. A "fully loaded" 64 QAM TDMA waveform with one channel in the Transmit Channel Set would be transmitted at 57 dBmV while a "fully loaded" 64 QAM waveform with three channels in the Transmit Channel Set would be transmitted at 51 dBmV. A 64 QAM TDMA waveform which is "underloaded" by 15 dB would transmit at 57 dBmV - 15 dB = 42 dBmV with one channel in the Transmit Channel Set, and a 64 QAM TDMA waveform which is "underloaded" by 15 dB would transmit at 51 dBmV - 15 dB = 36 dBmV with three channels in the Transmit Channel Set. The parameter P_{load_n} is used to convey the amount by which the nth channel is underloaded. Thus, in the case of the single channel in the Transmit Channel Set, with 42 dBmV for the 64 QAM TDMA transmit level, the channel is underloaded by 15 dB, therefore, $P_{load_1} = 15$ dB. Similarly, in the case of three channels in the Transmit Channel Set, with three 64 QAM TDMA modulations transmitting at 41 dBmV, 38 dBmV, and 36 dBmV, we have $P_{load_1} = 10$ dB, $P_{load_2} = 13$ dB, and $P_{load_3} = 15$ dB, where P_{load_1} corresponds to the channel with the lowest value of P_{load} , or equivalently, the highest or more full loading.

The concept of "loading," and the values of P_{load_n} , are only indirectly tied to the absolute transmit power, as seen in the preceding paragraph, with the absolute transmit power equaling the maximum transmit power (for the modulation and number of channels in the Transmit Channel Set) minus P_{load_n} for the nth channel.

For the case with multiple modulation types across the burst profiles of a channel, such as QPSK TDMA and 64 QAM TDMA burst profiles existing in one channel, the maximum power for the channel is determined by the lower-valued 64 QAM level of 57 dBmV compared to QPSK's 61 dBmV (assuming for this example a single channel in the Transmit Channel Set), just as in DOCSIS 2.0. $P_{hi_1} = 57$ dBmV for this case, as set by the 64 QAM TDMA maximum power. With "full loading" on this channel, 64 QAM TDMA will transmit at 57 dBmV. Note that

QPSK will transmit at $57 \text{ dBmV} - 1.2 \text{ dB} = 55.8 \text{ dBmV}$ in this situation (due to the constellation gain of -1.18 dB for QPSK). While the 64 QAM TDMA modulation is "fully loaded," and $P_{\text{load}_1} = 0 \text{ dB}$ in this example, the QPSK TDMA modulation on this channel is underloaded, in fact, by 5.2 dB . P_{hi_1} is determined by the 64 QAM modulation in this case, and the QPSK modulation is NOT determining P_{hi} (by virtue of the definition of $P_{\text{hi}} = \min(P_{\text{max}} - G_{\text{const}})$ over all burst profiles used by the CM in this channel). Therefore, it is only the loading of the 64 QAM modulation which is described by P_{load_1} . The loading occurring with QPSK in this example is irrelevant in the calculations and parameters involved in power control with Multiple Transmit Channel mode enabled.

The parameter $P_{\text{load_min_set}}$ is commanded to the CM by the CMTS, and sets in place the restriction on P_{load_n} that

$$P_{\text{load_min_set}} \leq P_{\text{load}_n} \leq P_{\text{load_min_set}} + 12 \text{ dB}, \text{ for all channels in the Transmit Channel Set.}$$

Since $P_{\text{load}_n} = P_{\text{hi}_n} - P_{r_n}$, the value of $P_{\text{load_min_set}}$ places the limit on P_{r_n} for each channel that

$$P_{\text{hi}_n} - P_{\text{load_min_set}} \geq P_{r_n} \geq P_{\text{hi}_n} - P_{\text{load_min_set}} - 12 \text{ dB.}$$

(Note that $P_{\text{low_multi}} = P_{\text{hi}_n} - P_{\text{load_min_set}} - 12 \text{ dB}$.)

Therefore, for each channel in the Transmit Channel Set, the range of P_{r_n} for that channel is restricted to operating over a window of 12 dB by the value commanded by the CMTS. This is the Dynamic Range Window for the CM for the n^{th} channel. Note that the actual transmit values, P_{r_n} , for each channel in the Transmit Channel Set may operate over a DIFFERENT 12 dB range, but the P_{load_n} values must all fall within the SAME 12 dB window. Therefore, there IS a unique Dynamic Range Window set for the CM by $P_{\text{load_min_set}}$, however the absolute transmit range for each channel may differ. For example, let there be two channels in the Transmit Channel Set, and one channel has only 64 QAM TDMA modulation and the other has only QPSK TDMA modulation, and further, $P_{\text{load_min_set}} = 15 \text{ dB}$. For each channel the Dynamic Range Window is such that $15 \text{ dB} \leq P_{\text{load}_n} \leq 27 \text{ dB}$. However, the channel with only 64 QAM TDMA modulation will operate with P_{r_n} between $54 \text{ dBmV} - 15 \text{ dB} = 39 \text{ dBmV}$ and $54 \text{ dBmV} - 15 \text{ dB} - 12 \text{ dB} = 27 \text{ dBmV}$, while the channel with only QPSK TDMA modulation will operate with P_{r_n} between $59.2 \text{ dBmV} - 15 \text{ dB} = 44.2 \text{ dBmV}$ and $59.2 \text{ dBmV} - 15 \text{ dB} - 12 \text{ dB} = 32.2 \text{ dBmV}$.

V.3 Example Upstream Power Control with Multiple Transmit Channel mode Enabled

In DOCSIS 2.0, when a change to the power level is commanded to the CM, the new P_r is tested against P_{hi} , and P_{low} , and P_r is clipped at P_{hi} or P_{low} if in fact P_r would have been higher than P_{hi} or lower than P_{low} (per [DOCSIS RFI 2.0], Section 6.2.18). With MSC enabled, P_r is tested and clipped against P_{hi_T} and P_{low_S} , and P_{on} is tested and clipped against P_{hi_S} .

When a CM is placed in Multiple Transmit Channel mode by the CMTS, $P_{\text{load_min_set}}$ must be provided to the CM. This will set the Dynamic Range Window for the CM.

With MTC mode enabled, each channel continues operating with the testing and clipping as described for DOCSIS 2.0, with the following exceptions:

1. An additional test is performed comparing the newly proposed P_{r_n} with $P_{\text{hi}_n} - P_{\text{load_min_set}}$, to restrict the maximum power level for the channel to be $P_{\text{load_min_set}}$ below P_{hi_n} . Effectively, the highest reported power level for each channel is reduced by $P_{\text{load_min_set}}$ (dB).
2. If when comparing the newly proposed P_{r_n} against $P_{\text{hi}_n} - P_{\text{load_min_set}}$, it is found that P_{r_n} is too high, then the power change command for P_{r_n} is to be ignored. (Note that when $P_{\text{load_min_set}}$ is 0 dB there will not be discarding of commands for going over the power range limit; this is because the protocol described in Section 6.2.19 provides that P_{r_n} is first tested and clipped against P_{hi_n} , without discarding if clipping is needed.)
3. The steps of #1 and #2 are performed for checking for the low power levels, too, where now P_{r_n} is not allowed to drop below $P_{\text{hi}_n} - (P_{\text{load_min_set}} + 12 \text{ dB})$.

4. For MSC enabled, $P_{r,n}$ is tested to ensure it is not exceeding $P_{hi,T,n} - P_{load_min_set}$, and also that $P_{r,n}$ is not below $P_{hi,T} - (P_{load_min_set} + 12 \text{ dB})$, and if either occurs the command changing $P_{r,n}$ is ignored.
5. For MSC enabled, P_{on} is clipped so that it will not exceed $P_{hi,S} - P_{load_min_set}$.

Consider an example with four TDMA channels, each with 64 QAM only. In this example, suppose the CMTS begins with $P_{load_min_set} = 3 \text{ dB}$; then for each channel the channel power, $P_{r,n}$, is limited to be $\leq P_{hi,n} - 3 \text{ dB}$, so that each channel is limited to 48 dBmV rather than the four-channel max for 64 QAM which is 51 dBmV. The lowest value allowed for each $P_{r,n}$ is 36 dBmV (12 dB below the highest value.) Consider that each channel is set for $P_{r,n} = 46 \text{ dBmV}$. This allows 2 dB of headroom before $P_{load_min_set}$ would have to be increased to allow $P_{r,n}$ to rise above 48 dBmV.

In this example, consider that the CMTS then commands one of the channels to drop its power level by 2 dB. This change would require a reconfiguration time on that channel, but all other channels may continue to receive grants and be bursting, and all fidelity requirements would have to still be met.

Then consider that the CMTS commands this channel to drop its power another 2 dB, and provides that channel a reconfiguration time. Now this channel has had its power changed by 4 dB since the last global reconfiguration time was provided, so the fidelity requirements are not enforced.

Now consider that this channel is moved back up by 2 dB, so its net change since the last global reconfiguration time is back to -2 dB. The fidelity requirements are still not in place; they were removed when the CM was commanded to move one of its channels by more than 3 dB since the last global reconfiguration time, and the specification states that the fidelity requirements are not back in place until the next global reconfiguration time.

To continue the example, assume again that the channel is dropped another 2 dB, down once again to 42 dBmV, 4 dB below the starting value.

When a global reconfiguration time is subsequently provided for the CM, all conditions are met for full fidelity requirements to be applicable.

Then, this channel is commanded to drop its power another 2 dB, and again it is provided a reconfiguration time such that it can implement this change; no global reconfiguration time is needed or provided. Now the fidelity requirements are slightly relaxed, because this channel loading is 6 dB higher (i.e., lower power) than the next closest channel loading (40 dBmV versus 46 dBmV). (Also note that it is the channel loading, i.e., $P_{r,n}$ compared to $P_{hi,n}$, and NOT the absolute value of $P_{r,n}$, that determines if the 4 dB range for 4 active channels is met.)

Continuing this example further, consider that another channel has its power dropped by 4 dB, and a reconfiguration time is provided on this channel, but until a global reconfiguration time is provided the fidelity requirements are not enforced (due to this 4 dB change since the previous global reconfiguration). Once a global reconfiguration time is provided, the full fidelity requirements will be in place again, since all conditions are met (the $P_{load,n}$ values are 5 dB, 5 dB, 9 dB, and 11 dB and the lowest $P_{load,n}$ value is within 3 dB of $P_{load_min_set}$ which is still 3 dB).

Now consider that the CMTS wishes to raise the power of one of the two remaining high-power channels by 2.5 dB. Unless $P_{load_min_set}$ is reduced by at least 0.5 dB, to 2.5 dB or less from its current value of 3 dB, this power change command would be ignored since it would result in a violation of the Dynamic Range Window for the channel. Let the CMTS send a change in $P_{load_min_set}$ to 1 dB, and at the same time increase the power in one of the two remaining channels by 2.5 dB. The CM will need to wait for a global reconfiguration time to make the dual changes in $P_{load_min_set}$ and in $P_{r,n}$ for the channel moving up to 48.5 dBmV. If all these changes are implemented at the CM, then for each channel, since each has only 64 QAM TDMA, the maximum value of $P_{r,n}$ (the top of the Dynamic Range Window) is 50 dBmV, and the four channels are commanded to 48.5 dBmV, 46 dBmV, 42 dBmV, and 40 dBmV. If the CMTS provides a global reconfiguration time, then these changes can be implemented. However, if the CMTS provides reconfiguration time to the one commanded channel, but this is not a global reconfiguration time, then the CM will not be able to implement the command to increase its power, since global reconfiguration

time is required for the change in the Dynamic Range Window. Note that even if the pending power change did not require the new $P_{\text{load_min_set}}$, since the channel power change was commanded concurrently with the change in the Dynamic Range Window, the CM should not implement the channel power change until the next global reconfiguration time. The CM should wait to implement both changes during the next global reconfiguration time, at which time the CM will still be meeting all fidelity requirements. Note that the P_{load_n} values are now 2.5 dB, 5 dB, 9 dB, and 11 dB.

As a further refinement of this example, consider that the channel at 40 dBmV actually had a QPSK TDMA modulation in its burst profiles, along with the 64 QAM. Let this channel also be commanded to reduce its power by another 2 dB, to $P_{r_n} = 38$ dBmV. P_{load_n} for this channel is thus 51 dBmV - 38 dBmV = 13 dB, which is exactly 12 dB below $P_{\text{load_min_set}}$, and thus P_{r_n} is exactly equal to $P_{\text{low_multi}} = 51$ dBmV - 1 dB - 12 dB = 38 dBmV in this example. There is no other change in this example. However, note that when QPSK is bursted on this channel it is bursting at 38 dBmV - 1.18 dB = 36.8 dBmV. Note that this is BELOW the power level of $P_{\text{low_multi}_n}$ for this channel! This is loading, compared to QPSK TDMA with four channels, which is 55 dBmV - 36.8 dBmV = 18.2 dB, while $P_{\text{load_min_set}}$ is 1 dB. Even though QPSK on this channel is underloaded by 18.2 dB, and this is 17.2 dB from $P_{\text{load_min_set}}$, the P_{load} for this channel is still 12 dB. The fact that the QPSK modulation is more than 17 dB underloaded compared to $P_{\text{load_min_set}}$ is not relevant. The CM must be able to accommodate transmission of the QPSK modulation at the slightly less loaded value than occurs with 64 QAM, since the power of the QPSK bursts compared to the 64 QAM bursts is only differing by 1.18 dB; the fact that the QPSK bursts are more underloaded (by our definitions in these requirements) does not imply that fidelity relief is mandated for this modulation, which is at almost the same absolute power levels for both the 64 QAM and QPSK bursts.

V.4 Examples Regarding Concurrent and Consecutive Changes in P_{r_n} and $P_{\text{load_min_set}}$

When the CMTS sends a new value of $P_{\text{load_min_set}}$ 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 $P_{\text{load_min_set}}$ 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_{r_n} may not be implemented immediately upon reception at the CM if the n^{th} channel is bursting. Commanded changes to P_{r_n} may occur simultaneously with the command to change $P_{\text{load_min_set}}$. The CMTS should not issue a change in $P_{\text{load_min_set}}$ after commanding a change in P_{r_n} until after also providing a sufficient reconfiguration time on the n^{th} channel. The CMTS should not issue a change in $P_{\text{load_min_set}}$ after commanding a prior change in $P_{\text{load_min_set}}$ until after also providing a global reconfiguration time for the first command. Also, the CMTS should not issue a change in P_{r_n} a) until after providing a global reconfiguration time following a command for a new value of $P_{\text{load_min_set}}$, and b) until after providing a sufficient reconfiguration time on the n^{th} channel after issuing a previous change in P_{r_n} . In other words, the CMTS is to avoid sending consecutive changes in P_{r_n} and/or $P_{\text{load_min_set}}$ to the CM without a sufficient reconfiguration time for instituting the first command. When a concurrent new value of $P_{\text{load_min_set}}$ and change in P_{r_n} are commanded, the CM may wait to apply the change in P_{r_n} at the next global reconfiguration time (i.e., concurrent with the institution of the new value of $P_{\text{load_min_set}}$) rather than applying the change at the first sufficient reconfiguration time of the n^{th} channel; the value of $P_{\text{load_min_set}}$ which applies to the new P_{r_n} is the concurrently commanded $P_{\text{load_min_set}}$ value, and thus if the change to P_{r_n} falls outside the Dynamic Range Window of the old $P_{\text{load_min_set}}$, then the CM has to wait for the global reconfiguration time to apply the change in P_{r_n} .

In the following example, P_{load_n} will be discussed along with P_{r_n} , to simplify the example.

In this example, let $P_{\text{load_min_set}}$ be 15 dB, and P_{load_1} is 16 dB and P_{load_2} is 18 dB, and a new command is issued to change $P_{\text{load_min_set}}$ to 14 dB, and concurrently to change P_{load_1} to 14 dB. Now assume the opportunity exists to change the power on Channel 1 but a global reconfiguration time is not provided until a bit later. Then the CM must recognize that it should NOT discard the change in P_{load_1} from 16 dB to 14 dB just because the new $P_{\text{load_min_set}}$ has not had its opportunity to be applied. If the change in P_{r_1} had been commanded BEFORE the change in $P_{\text{load_min_set}}$, then this change command to P_{r_1} places P_{r_1} outside of the Dynamic Range Window, and this command should be ignored by the CM. However, since in the example the change to P_{r_1} and $P_{\text{load_min_set}}$ was concurrent, the CM should bypass the opportunity to change P_{r_1} on Channel 1 when the quiet time on Channel 1 occurs, and the CM should wait for the global reconfiguration time and then change both P_{r_1} and $P_{\text{load_min_set}}$.

If the change in P_{r_1} had been only 1/2 dB, this would NOT have placed P_{r_1} outside the starting Dynamic Range Window, and the CM could have implemented the change in P_{r_1} at the first opportunity, but the CM is allowed in this case to wait and institute the change in P_{r_1} at the same time as the change in $P_{load_min_set}$, since the commands were sent concurrently to the CM.

If the change in P_{r_1} had been only 1/2 dB, and this change was commanded to the CM prior to the command to change $P_{load_min_set}$, then the change to P_{r_1} must be implemented at the first reconfiguration opportunity on Channel 1. If no such opportunity is provided and the change to $P_{load_min_set}$ is commanded, then the CM may ignore this commanded change in $P_{load_min_set}$, even though this command is in compliance with both the existing P_{r_1} and the pending P_{r_1} ; the CMTS should not issue such a command (the CMTS should wait for the CM to change P_{r_1} before sending a new $P_{load_min_set}$). Note that the recommended response of the CM to such an improperly issued change in $P_{load_min_set}$ is to continue with the change in P_{r_1} at the provided Channel 1 reconfiguration time (as it is required to do), and wait for the global reconfiguration time and implement (albeit inappropriately) the commanded $P_{load_min_set}$. Another circumstance in the event of the commanded change in $P_{load_min_set}$ prior to the opportunity for P_{r_1} to change is if the first subsequent opportunity for P_{r_1} to change is also a global reconfiguration time; it IS allowed for the CM to implement the change in P_{r_1} at this opportunity, but NOT to change $P_{load_min_set}$, since this was issued improperly by the CMTS. However, it IS recommended that the CM hold on to this commanded new value of $P_{load_min_set}$ and implement it at the NEXT global reconfiguration time.

Altering the example, if $P_{load_min_set}$ had been commanded to move from 15 dB to 17 dB with no changes in P_{load_1} or P_{load_2} , this change is to be completely discarded since it places P_{r_1} outside the Dynamic Range Window. If the CMTS issues such a command, and before the next global reconfiguration time P_{load_1} is commanded to change from 16 dB to 17 dB, it is still required that the CM ignore the $P_{load_min_set}$ change, since it was in violation of the Dynamic Range Window constraints when it was issued.

Appendix VI Example Spurious Emissions Noise Power Limits with Multiple Channels Bursting

This appendix provides examples illustrating the calculation of the noise power limits set by the spurious emissions requirements for the upstream transmissions when more than one channel is bursting. The spurious emission requirements are levied such that the noise power limits are relative to the channel transmit power; with multiple channels bursting simultaneously, the absolute noise limits from each transmitted channel are added together to produce the composite noise limit of the requirement. There are some conditions in the specification related to channel transmit powers falling below certain levels relative to the other channel transmit powers, where the noise limits are increased (relaxed). In these cases the noise power limit for the channel or channels which are transmitting below the power range determined by the higher power channels corresponds to the transmit power for the low-powered channel(s) if they were increased in power such as to fall just at the lower boundary of the power level range rather than below the range.

As one example, with three S-CDMA channels in the TCS, at transmit powers per channel of 53 dBmV, 41 dBmV, and 41 dBmV, the absolute spurious emissions requirements, when two or more channels in the TCS are transmitting, correspond to the absolute spurious emissions requirements with transmit power per channel of 53 dBmV, 45 dBmV, and 45 dBmV. The 41 dBmV power levels are raised (conceptually but not in actuality), for the purposes of determining the spurious emissions limits, up to $53 \text{ dBmV} - 8 \text{ dB} = 45 \text{ dBmV}$. If the modulation rate is 5120 kHz, and the noise limit is being calculated for a channel which is not adjacent to any of the three bursting channels, then the noise limit corresponds to that which would occur if each bursting channel contributed noise power -44 dB below its bursted signal power (or as in this example, below the artificially raised signal power for channels falling below the ranges for full application of the requirements). Thus in this example, with all three channels bursting, the noise limit for this test, in a non-adjacent transmit channel, is the sum of a) $53 \text{ dBmV} - 44 \text{ dB} = 9 \text{ dBmV}$, b) $45 \text{ dBmV} - 44 \text{ dB} = 1 \text{ dBmV}$, and c) $45 \text{ dBmV} - 44 \text{ dB} = 1 \text{ dBmV}$. Converting to natural numbers (rather than "dB" values) and adding the noise powers, the sum of $9 \text{ dBmV} + 1 \text{ dBmV} + 1 \text{ dBmV} = 10.2 \text{ dBmV}$, which serves as the noise power test limit for the spurious emissions requirements for this case.

As an example with four S-CDMA channels in the TCS, at transmit powers per channel of 53 dBmV, 41 dBmV, 41 dBmV, and 41 dBmV, the absolute spurious emissions requirements, when two or more channels in the TCS are transmitting, correspond to the absolute spurious emissions requirements with transmit power per channel of 53 dBmV, 49 dBmV, 45 dBmV, and 41 dBmV. The 41 dBmV power levels are raised (conceptually but not in actuality), for the purposes of determining the spurious emissions limits, up to $53 \text{ dBmV} - 4 \text{ dB} = 49 \text{ dBmV}$; and the artificial $49 \text{ dBmV} - 4 \text{ dB} = 45 \text{ dBmV}$; and the artificial $45 \text{ dBmV} - 4 \text{ dB} = 41 \text{ dBmV}$. If the modulation rate is 5120 kHz, and the noise limit is being calculated for a channel which is not adjacent to any of the four bursting channels, then the noise limit corresponds to that which would occur if each bursting channel contributed noise power -44 dB below its bursted signal power (or as in this example, below the artificially raised signal power for channels falling below the ranges for full application of the requirements). Thus in this example, with all four channels bursting, the noise limit for this test, in a non-adjacent transmit channel, is the sum of a) $53 \text{ dBmV} - 44 \text{ dB} = 9 \text{ dBmV}$, b) $49 \text{ dBmV} - 44 \text{ dB} = 5 \text{ dBmV}$, c) $45 \text{ dBmV} - 44 \text{ dB} = 1 \text{ dBmV}$, and d) $41 \text{ dBmV} - 44 \text{ dB} = -3 \text{ dBmV}$. Converting to natural numbers (rather than "dB" values) and adding the noise powers, the sum of $9 \text{ dBmV} + 5 \text{ dBmV} + 1 \text{ dBmV} + (-3 \text{ dBmV}) = 11.1 \text{ dBmV}$, which serves as the noise power test limit for the spurious emissions requirements for this case.

Appendix VII Acknowledgements

On behalf of the cable industry and our member companies, CableLabs would like to thank the numerous individuals that contributed to the development of this specification. In particular, we want to extend our sincere appreciation and gratitude to the following PHY focus team members.

Contributor	Company Affiliation	Contributor	Company Affiliation
Tom Ayers	Broadlogic	Jaehoo Lee	ETRI
Zvi Bernstein	ATI	Rick Meller	Cisco Systems
Liran Brecher	Texas Instruments	Jack Moran	Motorola, Inc
L. Alberto Campos	Cable Television Laboratories, Inc.	Richard Prodan	Broadcom Corporation
Phillip Chang	Comcast Corp.	Lakshmi Raman	Cable Television Laboratories, Inc.
Bruce Currivan	Broadcom Corporation	Saifur Rahman	Comcast Corp.
Kirk Erichsen	Adelphia	Vladimir Radionov	Broadlogic
Jeff Finkelstein	Cox Communications	Carey Ritchie	Microtune
Mike Grimwood	ATI	Hal Roberts	BigBand Networks
Tom Kolze	Broadcom Corporation	Greg Taylor	Cisco Systems
Bill Kostka	Cable Television Laboratories, Inc.	Robert Thompson	Motorola, Inc.
David Hull	Conexant	Dan Torbet	Arris International
Volker Lisse	EuroCableLabs	Daniel Wajcer	Texas Instruments
Eungdon Lee	ETRI		

We would particularly like to thank Tom Kolze of Broadcom for his commitment and leadership as the PHY Focus Team editor, Alberto Campos of CableLabs for all his efforts as Focus Team Lead, and all other participants who served as liaisons to other focus teams.

Our sincere appreciation goes out to all the participating vendor companies and their contributing engineers, in particular: Roger Fish of Broadcom Corporation; Xing Gao of Cisco Systems; Bill Kuiper of Microtune; Larry O'Sullivan of Bigband Networks; Ali Somnez of Arris International; and Naor Goldman, Shaul Klein, and Davey Weisberg of Texas Instruments.

Special thanks to Phillip Chang of Comcast and Scott Brown of Arris Interactive for the in-depth specification reviews; as well as to Scott Cummings, Lisa Denney, Victor Hou, and Niki Pantelias of Broadcom; Michael Patrick of Motorola; Ariel Yagil of Texas Instruments; and Eduardo Cardona, Brian Hedstrom, Greg White, and Charlie Bergren of CableLabs for their participation as invited specialists.