International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

J.122
(12/2007)

SERIES J: CABLE NETWORKS AND TRANSMISSION OF TELEVISION, SOUND PROGRAMME AND OTHER MULTIMEDIA SIGNALS

Interactive systems for digital television distribution

Second-generation transmission systems for interactive cable television services – IP cable modems

Recommendation ITU-T J.122

Recommendation ITU-T J.122

Second-generation transmission systems for interactive cable television services – IP cable modems

Summary

This Recommendation describes transmission systems that are consistent with those described in ITU-T Recommendation J.112. Its purpose is to provide a technology through which the demand for symmetrical services on cable networks can be met. This second-generation technology, while using the same RF channel, is backward compatible with that of ITU-T Recommendation J.112. It provides a significant increase in upstream channel capacity with wider channels, higher symbol rate and increased spectral efficiency. It provides for both Synchronous-CDMA and Advanced-TDMA coding and also adds higher order modulation capabilities. It is more robust with regard to noise immunity and provides greater multipath protection than the technology used in ITU-T Recommendation J.112.

Source

Recommendation ITU-T J.122 was approved on 14 December 2007 by ITU-T Study Group 9 (2005-2008) under Recommendation ITU-T A.8 procedure.

FOREWORD

NOTE

INTELLECTUAL PROPERTY RIGHTS

CONTENTS

		Page
1	Scope and purpose	1
	1.1 Scope	1
	1.2 Requirements	1
	1.3 Background	2
	1.4 Conventions for this Recommendation	4
2	References	4
	2.1 Normative references	4
	2.2 Informative references	6
3	Glossary	7
	3.1 Abbreviations and acronyms	7
	3.2 Terms and definitions	9
4	Functional assumptions	17
	4.1 Broadband access network	17
	4.2 Equipment assumptions	18
	4.3 RF channel assumptions	18
	4.4 Transmission levels	20
	4.5 Frequency inversion	21
5	Communication protocols	21
	5.1 Protocol stack	21
	5.2 The MAC forwarder	25
	5.3 Network layer	27
	5.4 Above the network layer	29
	5.5 Data link layer	30
	5.6 Physical layer	30
6	Physical media dependent sublayer specification	31
	6.1 Scope	31
	6.2 Upstream	31
	6.3 Downstream	88
7	Downstream transmission convergence sublayer	94
	7.1 Introduction	94
	7.2 MPEG packet format	95
	7.3 MPEG header for DOCS data-over-cable	95
	7.4 MPEG payload for DOCS data-over-cable	96
	7.5 Interaction with the MAC sublayer	96
	7.6 Interaction with the Physical Layer	98
	7.7 MPEG header synchronization and recovery	98
8	Media access control specification	98
	8.1 Introduction	98

Page

8.2 MAC frame formats 101
8.3 MAC management messages 121
9 Media access control protocol operation 172
9.1 Upstream bandwidth allocation 172
9.2 Support for multiple channels 181
9.3 Timing and synchronization 181
9.4 Upstream transmission and contention resolution 185
9.5 Data link encryption support 187
10 Quality of service and fragmentation 188
10.1 Theory of operation 188
10.2 Upstream service flow scheduling services 203
10.3 Fragmentation 207
10.4 Payload header suppression 214
11 Cable modem – CMTS interaction 221
11.1 CMTS initialization 221
11.2 Cable modem initialization 221
11.3 Standard operation 243
11.4 Dynamic service 250
11.5 Fault detection and recovery 309
12 Supporting future new cable modem capabilities 310
12.1 Downloading cable modem operating software 310
Annex A – Well-known addresses 312
A.1 MAC addresses 312
A.2 MAC service IDs 312
A.3 MPEG PID 313
Annex B – Parameters and constants 314
Annex C – Common radio frequency interface encodings 317
C.1 Encodings for configuration and MAC-layer messaging 317
C.2 Quality-of-service-related encodings 342
C.3 Encodings for other interfaces 366
C.4 Confirmation code 367
Annex D – CM configuration interface specification 371
D.1 CM IP addressing 371
D.2 CM configuration 372
D.3 Configuration verification 376
Annex E – The data-over-cable spanning tree protocol 378
E.1 Background 378
E.2 Public spanning tree 378
E.3 Public spanning tree protocol details 379

Page

E.4	Spanning tree parameters and defaults	380
Annex F – European specification additions		382
F.1	Scope and purpose	382
F.2	References	382
F.3	Glossary	382
F.4	Functional assumptions	382
F.5	Communication protocols	385
F.6	Physical media dependent sublayer specification	386
F.7	Downstream transmission convergence sublayer	407
F.8	Media access control specification	408
Annex G – DOCS 2.0 and 1.0/1.1 interoperability		410
G.1	General interoperability issues	410
G.2	Hybrid devices	413
G.3	DOCS 2.0 TDMA interoperability	414
G.4	DOCS 2.0 S-CDMA interoperability	415
Annex H – The DOCS MAC/PHY interface (DMPI)		416
H.1	Scope	416
H.2	Conventions	416
H.3	Overview	417
H.4	Signals	420
H.5	Protocol	423
H.6	Electrical specifications	428
H.7	Timing specifications	428
H.8	Data format and usage	430
Annex I – (Set aside)		438
Annex J – Japan specification additions		438
J.1	Scope and purpose	438
J.2	References	438
J.3	Glossary	438
J.4	Functional assumptions	438
J.5	Communication protocols	442
J.6	Physical media dependent sublayer specification	442
J.7	Downstream transmission convergence sublayer	465
J.8	Media access control specification	465
J.9	Media access control protocol operation	466
J.10	Quality of service and fragmentation	467
J.11	Cable modem – CMTS interaction	467
J.C	Common radio frequency interface encodings	467
J.E	The data-over-cable spanning tree protocol	468

Page

Appendix I – MAC service definition 469

I.1 MAC service overview 469

I.2 MAC data service interface 471

I.3 MAC control service interface 473

I.4 MAC service usage scenarios 476

Appendix II – Example preamble sequence 477

II.1 Introduction 477

II.2 Example preamble sequence 477

Appendix III – Multiple upstream channels 479

III.1 Single downstream and single upstream per cable segment 479

III.2 Multiple downstreams and multiple upstreams per cable segment 481

Appendix IV – DOCS transmission and contention resolution 485

IV.1 Introduction 485

IV.2 Variable definitions 486

IV.3 State examples 487

IV.4 Function examples 488

Appendix V – IGMP example 490

V.1 Events 490

V.2 Actions 491

Appendix VI – Unsolicited grant services 492

VI.1 Unsolicited grant service (UGS) 492

VI.2 Unsolicited grant service with activity detection (UGS-AD) 494

Appendix VII – S-CDMA framing 498

VII.1 Coded subsymbol numbering 498

VII.2 Uncoded subsymbol numbering 498

VII.3 Framer output numbering 499

VII.4 Comments 499

Appendix VIII – Ambient temperature and wind loading effects 500

VIII.1 Synchronization tolerances to plant delay variations 500

VIII.2 Change in propagation delay due to temperature changes 501

Bibliography 503

Recommendation ITU-T J.122

Second-generation transmission systems for interactive cable television services – IP cable modems

1 Scope and purpose

1.1 Scope

This Recommendation defines the second generation of radio-frequency interface specifications for high-speed data-over-cable systems.

There are differences in the cable spectrum planning practices adopted for different networks in the world. Therefore, three 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 using 6 MHz channelling, and supports upstream transmission in the 5-42 MHz region. The second technology option is based on a multi-programme television distribution using 8 MHz channel spacing and supports upstream in the 5-65 MHz region. The third technology option is based on 6 MHz channel spacing and supports upstream in the 10-55 MHz region. All options have equal status. The first of these options is defined in clauses 4, 6 and 7, whereas the second and third are defined by replacing the content of those clauses with the content of Annex F or Annex J. Compliance with this Recommendation requires compliance with only one of these implementations. It is not required that equipment built to one option shall interoperate with equipment built to another.

These optional physical-layer technologies allow operators some flexibility within any frequency planning, EMC and safety requirements that are mandated for their area of operation. For example, the 6 MHz downstream-based option defined by clauses 4, 6 and 7 might be deployable within an 8 MHz channel plan.

Frequency planning, safety and EMC requirements are a national matter and are not covered by this Recommendation. Compliance remains the operators' responsibility.

All optional physical-layer technologies are required to be backwards compatible with the earlier versions of those options defined in [DOCS9], [DOCS11] and [DOCS12].

NOTE – The structure and content of this Recommendation have been organized for ease of use by those familiar with the original source material; as such, the usual style of ITU-T recommendations has not been applied.

1.2 Requirements

If this Recommendation is implemented, the key words "MUST" and "SHALL" as well as "REQUIRED" are to be interpreted as indicating a mandatory aspect of this Recommendation. The keywords indicating a certain level of significance of a particular requirement that are used throughout this Recommendation are summarized below.

MUST	This word or the adjective "REQUIRED" means that the item is an absolute requirement of this Recommendation.
MUST NOT	This phrase means that the item is an absolute prohibition of this Recommendation.
SHOULD	This word or the adjective "RECOMMENDED" 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 behaviour is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behaviour described with this label.

MAY This word or the adjective "OPTIONAL" 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 Recommendation 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 Recommendation. Support of non-mandatory features and parameter values is optional.

1.3 Background

1.3.1 Service goals

As cable operators have widely deployed high-speed data services on cable television systems, the demand for upstream bandwidth has increased, particularly with the popularity of more symmetric data applications. The current Recommendation has been created for the purpose of increasing channel capacity and improving noise immunity.

The intended service will allow transparent bidirectional transfer of Internet protocol (IP) traffic, between the cable system headend and customer locations, over an all-coaxial or hybrid-fibre/coax (HFC) cable network. This is shown in simplified form in Figure 1-1.

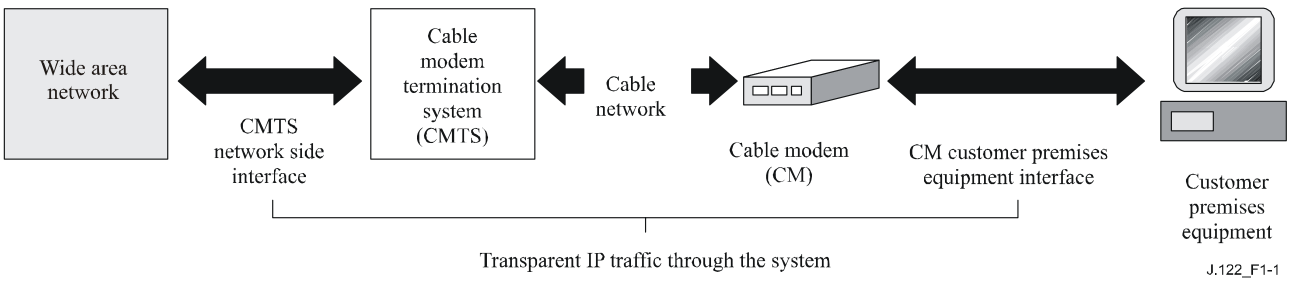

Figure 1-1 – Transparent IP traffic through the data-over-cable system

The transmission path over the cable system is realized at the headend by a cable modem termination system (CMTS), and at each customer location by a cable modem (CM). At the headend (or hub), the interface to the data-over-cable system is called the cable modem termination system-network side interface (CMTS-NSI). At the customer locations, the interface is called the cable-modem-to-customer-premises-equipment interface (CMCI). The intent is for operators to transparently transfer IP traffic between these interfaces, including but not limited to datagrams, DHCP, ICMP and IP group addressing (broadcast and multicast).

1.3.2 Reference architecture

The reference architecture for the data-over-cable services and interfaces is shown in Figure 1-2.

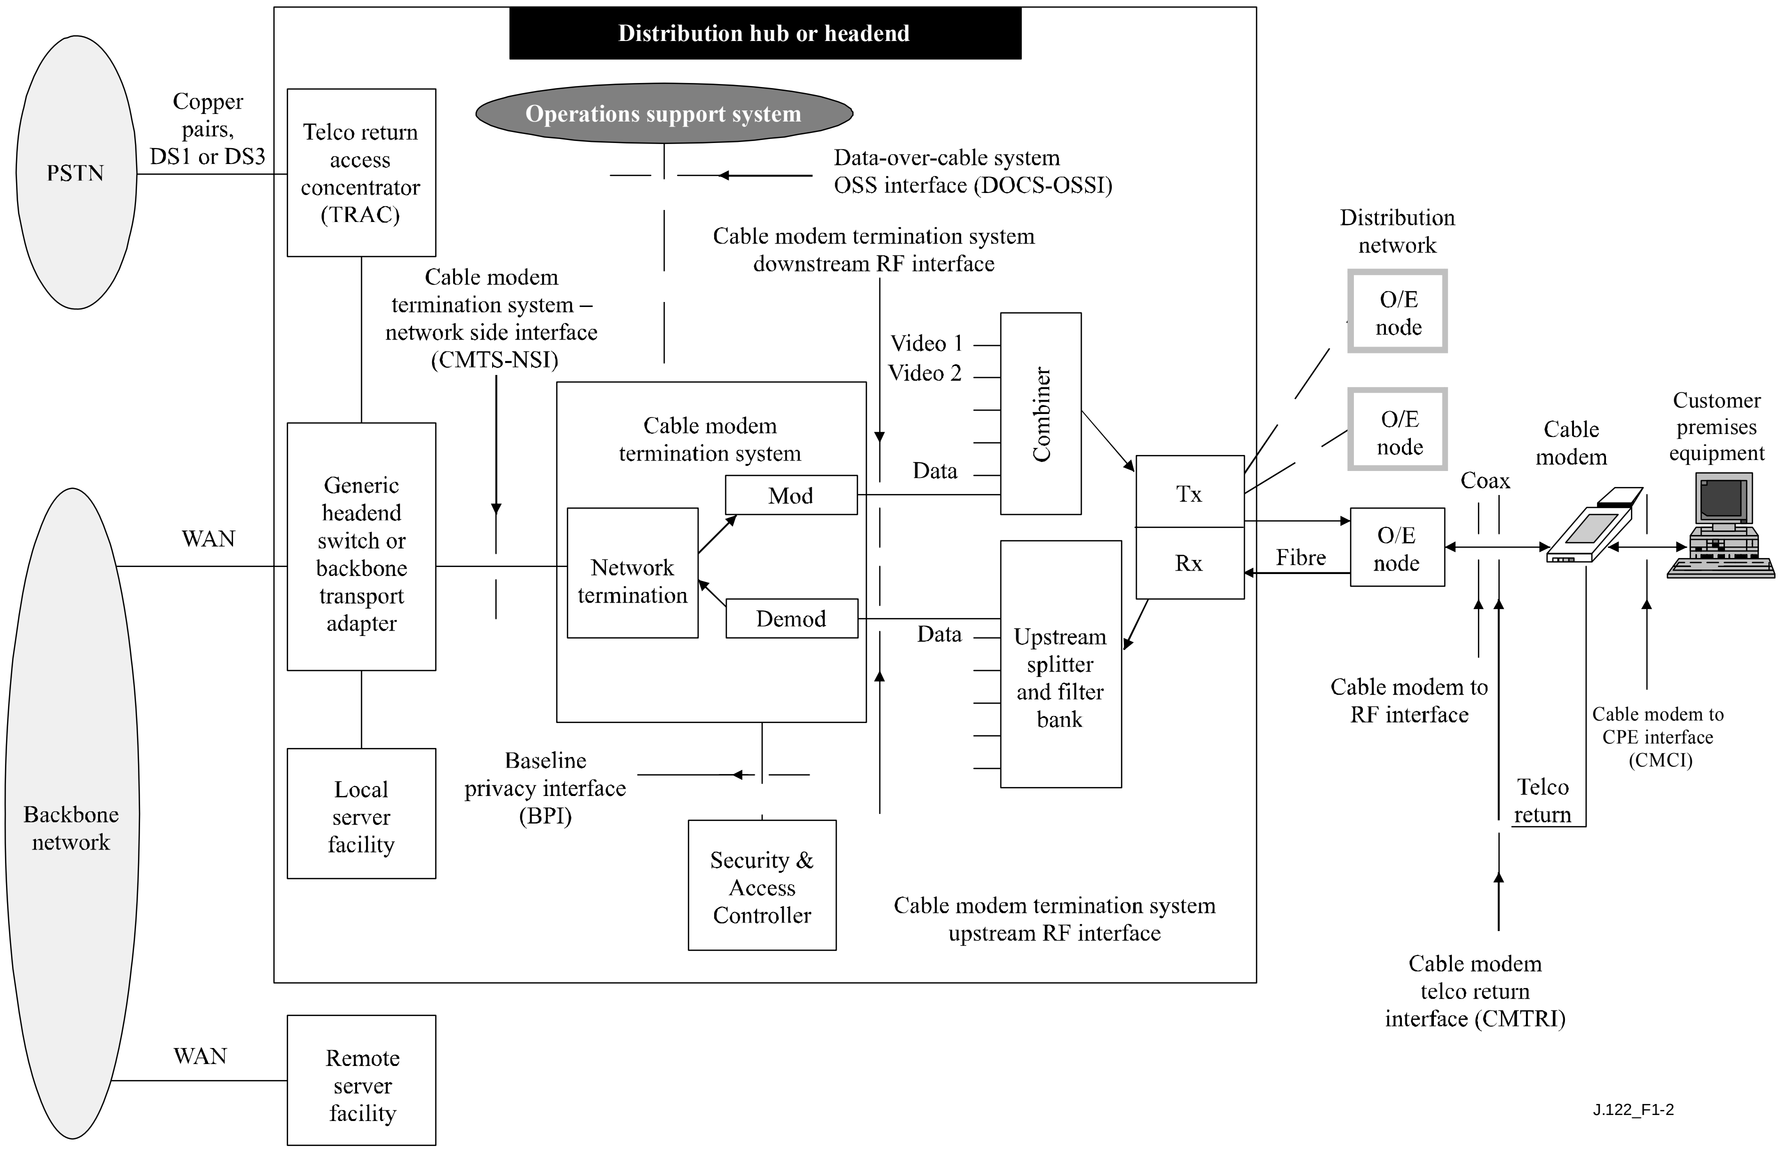

Figure 1-2 – Data-over-cable reference architecture

1.3.3 Categories of interface specification

The basic reference architecture of Figure 1-2 involves four interface categories.

Data interfaces – These are the CMCI and CMTS-NSI, corresponding respectively to the cable-modem-to-customer-premises-equipment (CPE) interface (for example, between the customer's computer and the cable modem), and the cable modem termination system network side interface between the cable modem termination system and the data network.

Operations support systems interfaces – These are network element management layer interfaces between the network elements and the high-level operations support systems (OSSs), which support the basic business processes.

RF interfaces – The RF interfaces defined in this Recommendation are the following:

- between the cable modem and the cable network;
- between the CMTS and the cable network in the downstream direction (traffic toward the customer);
- between the CMTS and the cable network in the upstream direction (traffic from the customer).

Security interfaces – Baseline data-over-cable security is defined in [DOCS8].

1.3.4 Statement of compatibility

This Recommendation specifies an interface, commonly referred to as DOCS 2.0, which is the second generation of the interface specified in [DOCS9], [DOCS11] and [DOCS12], commonly referred to as DOCS 1.x. DOCS 2.0 must be backward- and forward-compatible with equipment built to the previous Recommendations. DOCS 2.0-compliant CMs MUST interoperate seamlessly with DOCS 1.x CMTSs, albeit in the 1.x mode, as the case may be. DOCS 2.0-compliant CMTSs MUST seamlessly support DOCS 1.x CMs.

Refer to Annex G for further interoperability information.

1.4 Conventions for this Recommendation

In this Recommendation, 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 from top to bottom, with the MSB being the first bit so read and the LSB being the last bit so read.

2 References

2.1 Normative references

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

[DOCS3] Data-Over-Cable Service Interface Specifications, *Cable Modem Termination System – Network Side Interface Specification*, SP-CMTS-NSII01-960702.

[DOCS4] Data-Over-Cable Service Interface Specifications, *Cable Modem to Customer Premise Equipment Interface Specification*, SP-CMCI-I04-000714.

[DOCS5] ANSI/SCTE 79-2 2002, *Data-Over-Cable System S 2.0 s, Operations Support System Interface*.

[DOCS8] ITU-T J.112 Annex B (2001), *Data-over-cable service interface specifications, Annex O – Privacy for J.112 Annex B Implementations*.

[DOCS9] SCTE 22-1 2002, *Data-Over-Cable Service Interface Specification, DOCSIS 1.0 Radio Frequency Interface (RFI)*.

[DOCS11] ITU-T Recommendation J.112 Annex B (2001), *Data-over-cable service interface specifications: Radio-frequency interface specification*.

[DOCS12] ITU-T Recommendation J.112 Annex C (2002), *Data-over-cable service interface specifications: Radio-frequency interface specification using QAM technique*.

[IEEE 802] IEEE 802 (2001), *Local and Metropolitan Area Networks: Overview and Architecture.*

[IEC 60169-24] IEC 60169-24 (1991), *Radio-frequency connectors – Part 24: Radio-frequency coaxial connectors with screw coupling, typically for uses in 75 ohm cable distribution systems (Type F)*.

[ISO/IEC 8802-2] ISO/IEC 8802-2:1998 (ANSI/IEEE Std 802.2), *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 2: Logical link control*.

[ISO/IEC 8802-3] ISO/IEC 8802-3:2000 (IEEE Std 802.3), *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications*.

[ISO/IEC 15802-3] ISO/IEC 15802-3:1998 (ANSI/IEEE Std 802.1D), *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Common specifications – Part 3: Media Access Control (MAC) Bridges*.

[ISO/IEC 15802-1] ISO/IEC 15802-1:1995, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Common specifications – Part 1: Medium Access Control (MAC) service definition*.

[ISO/IEC 10039] ISO/IEC 10039:1991, *Information technology – Open systems interconnection – Medium Access Control (MAC) service definition*.

[ITU-T H.222.0] ITU-T Recommendation H.222.0 (2000) | ISO/IEC 13818-1:2000, *Information technology – Generic coding of moving pictures and associated audio information: Systems*.

[ITU-T J.83] ITU-T Recommendation J.83 (2007), *Digital multi-programme systems for television, sound and data services for cable distribution*.

[ITU-T J.162] ITU-T Recommendation J.162 (2007), *Network call signalling protocol for the delivery of time-critical services over cable television networks using cable modems*.

[ITU-T J.163] ITU-T Recommendation J.163 (2007), *Dynamic quality of service for the provision of real-time services over cable television networks using cable modems*.

[ITU-T X.25] ITU-T Recommendation X.25 (1996), *Interface between Data Terminal Equipment (DTE) and Data Circuit-terminating Equipment (DCE) for terminals operating in the packet mode and connected to public data networks by dedicated circuit*.

[ITU-T X.690] ITU-T X.690-series (2002) | ISO/IEC 8825:2002, *Information technology – ASN.1 encoding rules*.

[ITU-T Z.100] ITU-T Recommendation Z.100 (2002), *Specification and Description Language (SDL)*.

[L2VPN] DOCSIS Business Services over DOCSIS, *Layer 2 Virtual Private Networks*, CM-SP-L2VPN-105-070803, August 3, 2007, Cable Television Laboratories, Inc.

[RFC 791] IETF RFC 791 (1981), *Internet Protocol*.

[RFC 826] IETF RFC 826 (1982), *Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet Hardware*.

[RFC 868] IETF RFC 868 (1983), *Time Protocol*.

[RFC 1042] IETF RFC 1042 (1988), *Standard for the transmission of IP datagrams over IEEE 802 networks*.

[RFC 1123] IETF RFC 1123 (1989), *Requirements for Internet Hosts – Application and Support*.

[RFC 1157] IETF RFC 1157 (1990), *Simple Network Management Protocol (SNMP)*.

[RFC 1350] IETF RFC 1350 (1992), *The TFTP Protocol (Revision 2)*.

[RFC 2104] IETF RFC 2104 (1997), *HMAC: Keyed-Hashing for Message Authentication*.

[RFC 2131] IETF RFC 2131 (1997), *Dynamic Host Configuration Protocol*.

[RFC 2132] IETF RFC 2132 (1997), *DHCP Options and BOOTP Vendor Extensions*.

[RFC 2236] IETF RFC 2236 (1997), *Internet Group Management Protocol, Version 2*.

[RFC 2349] IETF RFC 2349 (1998), *TFTP Timeout Interval and Transfer Size Options*.

[RFC 2786] IETF RFC 2786 (2000), *Diffie-Helman USM Key Management Information Base and Textual Convention*.

[RFC 3046] IETF RFC 3046 (2001), *DHCP Relay Agent Information Option*.

[FIPS 180-1] FIPS 180-1 (1995): *Secure Hash Standard*.

2.2 Informative references

[b-CableLabs2] CableLabs2 (1999), *Digital Transmission Characterization of Cable Television Systems*.

[b-EIA 542] CEA Standard 542-B (2003), *Cable Television Channel Identification Plan*.

[b-EN 300 429] ETSI EN 300 429 V1.2.1 (1998), *Digital Video Broadcasting (DVB); Framing structure, channel coding and modulation for cable systems*.

[b-ID-IGMP] FENNER (W.): IGMP-based Multicast Forwarding ("IGMP Proxying"), IETF Internet Draft, 2000. http://www.ietf.org/proceedings/00dec/I-D/draft-fenner-igmp-proxy-03.txt

[b-IEEE 802.1Q] IEEE 802.1Q (2003), *Virtual Bridged Local Area Networks*.

[b-IMA] Internet Multicast Addresses, *Internet Assigned Numbers Authority*. http://www.iana.org/assignments/multicast-addresses.

[b-NCTA] NCTA Recommended Practices for Measurements on Cable Television Systems, *National Cable Television Association*, 2nd Edition, Washington DC, 1993.

[b-RFC 1058] IETF RFC 1058 (1988), *Routing Information Protocol*.

[b-RFC 1493] IETF RFC 1493 (1993), *Definitions of Managed Objects for Bridges*.

[b-RFC 1633] IETF RFC 1633 (1994), *Integrated Services in the Internet Architecture: an Overview*.

[b-RFC 1812] IETF RFC 1812 (1995), *Requirements for IP Version 4 Routers*.

[b-RFC 2212] IETF RFC 2212 (1997), *Specification of Guaranteed Quality of Service*.

[b-RFC 2669] IETF RFC 2669 (1999), *DOCSIS Cable Device MIB Cable Device Management Information Base for DOCSIS compliant Cable Modems and Cable Modem Termination Systems*.

[b-VECCHI] VECCHI (M.), FAHIM (M.): Architectural Model: The Spectrum Management Application (SMA) and the Common Spectrum Management Interface (CSMI), *Time Warner Cable*, Englewood, CO, 1995.

3 Glossary

3.1 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

ARP	Address Resolution Protocol
ATM	Asynchronous Transfer Mode
BPDU	Bridge Protocol Data Unit
C/N or CNR	Carrier-to-Noise Ratio
CCCM	CPE Controlled Cable Modem
CM	Cable Modem
CMCI	Cable Modem to CPE Interface
CMTS	Cable Modem Termination System
CMTS-NSI	Cable Modem Termination System – Network Side Interface
CPE	Customer Premises Equipment
CSO	Composite Second Order beat
CTB	Composite Triple Beat
DCC	Dynamic Channel Change
DHCP	Dynamic Host Configuration Protocol
DOCS	Data-Over-Cable System
FDDI	Fibre Distributed Data Interface
HF	High Frequency
HFC	Hybrid Fibre/Coax system
HRC	Harmonic Related Carrier

ICMP	Internet Control Message Protocol
IE	Information Element
IGMP	Internet Group Management Protocol
IP	Internet Protocol
IRC	Incremental Related Carriers
IUC	Interval Usage Code
LAN	Local Area Network
LLC	Logical Link Control procedure
MAC	Media Access Control procedure
MAP	Bandwidth Allocation Map
MPEG	Moving Picture Experts Group
MSAP	MAC Service Access Point
MTTR	Mean Time to Repair
NTSC	National Television Systems Committee
OSI	Open Systems Interconnection
OUI	Organizationally Unique Identifier
PHS	Payload Header Suppression
PHY	Physical layer
PID	Packet Identifier
PMD	Physical Media Dependent
PSI	Program-Specific Information
PUSI	Payload Unit Start Indicator
QAM	Quadrature Amplitude Modulation
QPSK	Quadrature Phase-Shift Keying
RF	Radio Frequency
RFC	Request for Comments
RIP	Routing Information Protocol
SAID	Security Association Identifier
SAP	Service Access Point
SDU	Service Data Unit
SFID	Service Flow Identifier
SID	Service Identifier
SMS	Spectrum Management System
SNAP	SubNetwork Access Protocol
SNMP	Simple Network Management Protocol
TCP	Transmission Control Protocol
TFTP	Trivial File-Transfer Protocol

TLV	Type/Length/Value
ToS	Type of Service
UCC	Upstream Channel Change
UCD	Upstream Channel Descriptor

3.2 Terms and definitions

This Recommendation defines the following terms:

3.2.1 active service flow: An admitted service flow from the CM to the CMTS which is available for packet transmission.

3.2.2 address resolution protocol (ARP): A protocol of the IETF for converting network addresses to 48-bit Ethernet addresses.

3.2.3 admitted service flow: A service flow, either provisioned or dynamically signalled, which is authorized and for which resources have been reserved but is not active.

3.2.4 allocation: A group of contiguous mini-slots in a MAP which constitute a single transmit opportunity.

3.2.5 American National Standards Institute (ANSI): A US standards body.

3.2.6 asynchronous transfer mode (ATM): A protocol for the transmission of a variety of digital signals using uniform 53-byte cells.

3.2.7 A-TDMA: DOCS 2.0 TDMA mode (as distinguished from DOCS 1.x TDMA).

3.2.8 authorization module: The authorization module is an abstract module that the CMTS can contact to authorize service flows and classifiers. The authorization module tells the CMTS whether the requesting CM is authorized for the resources it is requesting.

3.2.9 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.

3.2.10 bandwidth allocation map: The MAC management message that the CMTS uses to allocate transmission opportunities to CMs.

3.2.11 bridge protocol data unit (BPDU): Spanning tree protocol messages as defined in [ISO/IEC 15802-3].

3.2.12 broadcast addresses: A predefined destination address that denotes the set of all data network service access points.

3.2.13 burst: A single continuous RF signal from the upstream transmitter, from transmitter on to transmitter off.

3.2.14 burst error second: Any errored second containing at least 100 errors.

3.2.15 cable modem (CM): A modulator-demodulator at subscriber locations intended for use in conveying data communications on a cable television system.

3.2.16 cable modem termination system (CMTS): Cable modem termination system, located at the cable television system headend or distribution hub, which provides complementary functionality to the cable modems to enable data connectivity to a wide-area network.

3.2.17 cable modem termination system – network side interface (CMTS-NSI): The interface, defined in [DOCS3], between a CMTS and the equipment on its network side.

3.2.18 cable modem to CPE interface (CMCI): The interface, defined in [DOCS4], between a CM and CPE.

3.2.19 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.

3.2.20 carrier-to-noise ratio (C/N or CNR): The ratio of signal power to noise power in the defined measurement bandwidth. For digital modulation, CNR = E_s/N_0, 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.

3.2.21 CPE controlled cable modem (CCCM): Refer to the DOCS cable modem to customer premises equipment interface (CMCI) specification.

3.2.22 channel: The frequency spectrum occupied by a signal. Usually specified by centre frequency and bandwidth parameters.

3.2.23 chip: Each of the 128 bits comprising the S-CDMA spreading codes.

3.2.24 chip duration: The time to transmit one chip of the S-CDMA spreading code. The inverse of the chip rate.

3.2.25 chip rate: The rate at which individual chips of the S-CDMA spreading codes are transmitted (1280 to 5120 kHz).

3.2.26 classifier: A set of criteria used for packet matching according to TCP, UDP, IP, LLC, and/or 802.1P/Q packet fields. A classifier maps each packet to a service flow. A downstream classifier is used by the CMTS to assign packets to downstream service flows. An upstream classifier is used by the CM to assign packets to upstream service flows.

3.2.27 code-hopping matrix: A shifted version of the reference code matrix (see below) that is used when code hopping is employed to vary the codes used by each CM. The code-hopping matrix is either 128 rows by 128 columns (when all 128 codes are active) or is 127 rows by 128 columns (when less than 128 codes are active in the S-CDMA spreader-on frame). When less than 128 codes are active, code 0 (all ones) is deleted from the matrix, but all remaining codes are still cycled through even if less than 127 codes are active in a frame.

3.2.28 composite second order beat (CSO): The peak of the average level of distortion products due to second-order non-linearities in cable system equipment.

3.2.29 composite triple beat (CTB): The peak of the average level of distortion components due to third-order non-linearities in cable system equipment.

3.2.30 cross-modulation: A form of television signal distortion where modulation from one or more television channels is imposed on another channel or channels.

3.2.31 customer: See end user.

3.2.32 customer premises equipment (CPE): Equipment at the end user's premises; MAY be provided by the end user or the service provider.

3.2.33 data link layer: Layer 2 in the open system interconnection (OSI) architecture; the layer that provides services to transfer data over the transmission link between open systems.

3.2.34 distribution hub: A location in a cable television network which performs the functions of a headend for customers in its immediate area, and which receives some or all of its television programme material from a master headend in the same metropolitan or regional area.

3.2.35 DOCS 1.0: Data-over-cable system defined by [DOCS9].

3.2.36 DOCS 1.1: Data-over cable system defined by [DOCS11].

3.2.37 DOCS 1.x: Abbreviation for "DOCS 1.0 or 1.1".

3.2.38 DOCS 2.0: Data-over-cable system as defined in this Recommendation.

3.2.39 downstream: In cable television, the direction of transmission from the headend to the subscriber.

3.2.40 drop cable: Coaxial cable that connects to a residence or service location from a directional coupler (tap) on the nearest coaxial feeder cable.

3.2.41 dynamic host configuration protocol (DHCP): An Internet protocol used for assigning network-layer (IP) addresses.

3.2.42 dynamic range: The ratio between the greatest signal power that can be transmitted over a multichannel analogue 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.

3.2.43 Electronic Industries Association (EIA): A voluntary body of manufacturers which, among other activities, prepares and publishes standards.

3.2.44 end user: A human being, organization or telecommunication system that accesses the network in order to communicate via the services provided by the network.

3.2.45 errored second: Any one-second interval containing at least one bit error.

3.2.46 extended subsplit: A frequency division scheme that allows bidirectional traffic on a single coaxial cable. Reverse path signals come to the headend from 5 to 42 MHz. Forward path signals go from the headend from 50 or 54 MHz to the upper frequency limit.

3.2.47 feeder cable: Coaxial cables that run along streets within the served area and connect between the individual taps which serve the customer drops.

3.2.48 fibre distributed data interface (FDDI): A fibre-based LAN standard.

3.2.49 fibre node: A point of interface between a fibre trunk and the coaxial distribution.

3.2.50 forward channel: The direction of RF signal flow away from the headend toward the end user; equivalent to downstream.

3.2.51 frame: See MAC frame, S-CDMA frame and MPEG frame.

3.2.52 group delay: The difference in transmission time between the highest and lowest of several frequencies through a device, circuit or system.

3.2.53 guard time: Minimum time allocated between bursts in the upstream referenced from the symbol centre of the last symbol of a burst to the symbol centre of the first symbol of the following burst. The guard time should be at least the duration of five symbols plus the maximum system timing error.

3.2.54 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.

3.2.55 headend: The central location on the cable network that is responsible for injecting broadcast video and other signals in the downstream direction. See also master headend, distribution hub.

3.2.56 header: Protocol control information located at the beginning of a protocol data unit.

3.2.57 high frequency (HF): Used in this Recommendation to refer to the entire subsplit (5-30 MHz) and extended subsplit (5-42 MHz) band used in reverse channel communications over the cable television network.

3.2.58 high return: A frequency division scheme that allows bidirectional traffic on a single coaxial cable. Reverse channel signals propagate to the headend above the downstream passband.

3.2.59 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.

3.2.60 hybrid fibre/coax (HFC) system: A broadband bidirectional shared-media transmission system using fibre trunks between the headend and the fibre nodes, and coaxial distribution from the fibre nodes to the customer locations.

3.2.61 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.

3.2.62 Institute of Electrical and Electronic Engineers (IEEE): A voluntary organization which, among other things, sponsors standards committees and is accredited by the American National Standards Institute.

3.2.63 International Electrotechnical Commission (IEC): An international standards body.

3.2.64 International Organization for Standardization (ISO): An international standards body, commonly known as the International Standards Organization.

3.2.65 Internet control message protocol (ICMP): An Internet network-layer protocol.

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

3.2.67 Internet group management protocol (IGMP): A network-layer protocol for managing multicast groups on the Internet.

3.2.68 impulse noise: Noise characterized by non-overlapping transient disturbances.

3.2.69 information element: The fields that make up a MAP and define individual grants, deferred grants, etc.

3.2.70 Internet protocol (IP): An Internet network-layer protocol.

3.2.71 interval usage code (IUC): A field in MAPs and UCDs to link burst profiles to grants.

3.2.72 latency: The time, expressed in quantity of symbols, taken for a signal element to pass through a device.

3.2.73 layer: A subdivision of the open system interconnection (OSI) architecture, constituted by subsystems of the same rank

3.2.74 link layer: See data link layer.

3.2.75 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.

3.2.76 logical link control (LLC) procedure: In a local area network (LAN) or a metropolitan area network (MAN), that part of the protocol that governs the assembling of data link layer frames and their exchange between data stations, independent of how the transmission medium is shared.

3.2.77 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.

3.2.78 MAC frame: MAC header plus optional PDU.

3.2.79 MAC service access point (MSAP): An attachment to a MAC-sublayer domain. Refer to clauses 5.2 and 8.1.2.2.

3.2.80 master headend: A headend which collects television programme material from various sources by satellite, microwave, fibre and other means, and distributes this material to distribution hubs in the same metropolitan or regional area. A master headend MAY also perform the functions of a distribution hub for customers in its own immediate area.

3.2.81 mean time to repair (MTTR): In cable television systems, the MTTR is the average elapsed time from the moment a loss of RF channel operation is detected up to the moment the RF channel operation is fully restored.

3.2.82 media access control (MAC) address: The "built-in" hardware address of a device connected to a shared medium.

3.2.83 media access control (MAC) procedure: In a subnetwork, that part of the protocol that governs access to the transmission medium independent of the physical characteristics of the medium, but taking into account the topological aspects of the subnetworks, in order to enable the exchange of data between nodes. MAC procedures include framing, error protection and acquiring the right to use the underlying transmission medium.

3.2.84 media access control (MAC) sublayer: 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.

3.2.85 micro-reflections: Echoes in the forward transmission path due to departures from ideal amplitude and phase characteristics.

3.2.86 mid split: A frequency division scheme that allows bidirectional traffic on a single coaxial cable. Reverse channel signals propagate to the headend from 5 to 108 MHz. Forward path signals go from the headend from 162 MHz to the upper frequency limit. The diplex crossover band is located from 108 to 162 MHz.

3.2.87 mini-slot: A "mini-slot" is an integer multiple of 6.25-microsecond increments. The relationship between mini-slots, bytes and time ticks is described in clause 9.3.4.

3.2.88 modulation rate: The signalling rate of the upstream modulator (1280 to 5120 kHz). In S-CDMA, the chip rate. In TDMA, the channel symbol rate.

3.2.89 Moving Picture Experts Group (MPEG): A voluntary body which develops standards for digital compressed moving pictures and associated audio.

3.2.90 multipoint access: User access in which more than one terminal equipment is supported by a single network termination.

3.2.91 multipoint connection: A connection among more than two data network terminations.

3.2.92 National Cable Television 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 United States.

3.2.93 National Television Systems Committee (NTSC): Committee which defined the analogue colour television broadcast standard used today in North America.

3.2.94 network layer: Layer 3 in the open systems interconnection (OSI) architecture; the layer that provides services to establish a path between open systems.

3.2.95 network management: The functions related to the management of data link layer and physical layer resources and their stations across the data network supported by the hybrid fibre/coax system.

3.2.96 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.

3.2.97 open systems interconnection (OSI): A framework of ISO standards for communication between different systems made by different vendors, in which the communication process is organized into seven different categories that are placed in a layered sequence based on their relationship to the user. Each layer uses the layer immediately below it and provides a service to the layer above. Layers 7 through 4 deal with end-to-end communication between the message source and destination, and layers 3 through 1 deal with network functions.

3.2.98 organizationally unique identifier (OUI): A 3-octet IEEE-assigned identifier that can be used to generate universal LAN MAC addresses and protocol identifiers per ANSI/IEEE Std 802 for use in local and metropolitan area network applications.

3.2.99 packet identifier (PID): A unique integer value used to identify elementary streams of a programme in a single- or multi-programme MPEG-2 stream.

3.2.100 partial grant: A grant that is smaller than the corresponding bandwidth request from the CM.

3.2.101 payload header suppression (PHS): The suppression of the header in a payload packet (e.g., the suppression of the Ethernet header in forwarded packets).

3.2.102 payload unit start indicator (PUSI): A flag in an MPEG header. A value of 1 indicates the presence of a pointer field as the first byte of the payload.

3.2.103 physical layer (PHY): 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.

3.2.104 physical media dependent (PMD) sublayer: 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.

3.2.105 primary service flow: All CMs have a primary upstream service flow and a primary downstream service flow. They ensure that the CM is always manageable and they provide a default path for forwarded packets that are not classified to any other service flow.

3.2.106 programme-specific information (PSI): In MPEG-2, normative data necessary for the demultiplexing of transport streams and the successful regeneration of programmes.

3.2.107 programme stream: In MPEG-2, a multiplex of variable-length digital video and audio packets from one or more programme sources having a common time-base.

3.2.108 protocol: A set of rules and formats that determines the communication behaviour of layer entities in the performance of the layer functions.

3.2.109 provisioned service flow: A service flow that has been provisioned as part of the registration process, but has not yet been activated or admitted. It may still require an authorization exchange with a policy module or external policy server prior to admission.

3.2.110 QoS parameter set: The set of service flow encodings that describe the quality of service attributes of a service flow or a service class (refer to clause C.2.2.5).

3.2.111 quadrature amplitude modulation (QAM): A method of modulating digital signals onto a radio-frequency carrier signal involving both amplitude and phase coding.

3.2.112 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.

3.2.113 radio frequency (RF): In cable television systems, this refers to electromagnetic signals in the range 5 to 1000 MHz.

3.2.114 reference code matrix: A 128-by-128 element matrix formed by stacking successive spreading codes on top of each other, i.e., the bottom row of the reference code matrix is code 0 (all ones) and the top row is code 127. The code elements are placed in the matrix from right to left, i.e., the right-most column of the code matrix is the first element of each code, and the left-most column is the last element of each code.

3.2.115 request for comments (RFC): A technical policy document of the IETF; these documents can be accessed on the worldwide web at http://www.rfc-editor.org/.

3.2.116 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.

3.2.117 reverse channel: The direction of signal flow towards the headend, away from the subscriber; equivalent to upstream.

3.2.118 routing information protocol (RIP): A protocol of the IETF for exchanging routing information about IP networks and subnets.

3.2.119 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 $c \times K$ information (QAM) symbols.

3.2.120 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.

3.2.121 security association identifier (SAID): A baseline privacy security identifier between a CMTS and a CM.

3.2.122 service access point (SAP): The point at which services are provided by one layer, or sublayer to the layer immediately above it.

3.2.123 service class: A set of queueing and scheduling attributes that is named and that is configured at the CMTS. A service class is identified by a service class name. A service class has an associated QoS parameter set.

3.2.124 service class name: An ASCII string by which a service class may be referenced in modem configuration files and protocol exchanges.

3.2.125 service data unit (SDU): Information that is delivered as a unit between peer service access points.

3.2.126 service flow: A MAC-layer transport service which:

- provides unidirectional transport of packets from the upper layer service entity to the RF;
- shapes, polices and prioritizes traffic according to QoS traffic parameters defined for the flow.

3.2.127 service flow identifier (SFID): An identifier assigned to a service flow by the CMTS (32 bits).

3.2.128 service flow reference: A message parameter in configuration files and dynamic service MAC messages used to associate classifiers and other objects in the message with the service flow encodings of a requested service flow.

3.2.129 service identifier (SID): An identifier assigned by the CMTS (in addition to a service flow identifier) to an active or admitted upstream service flow (14 bits).

3.2.130 simple network management protocol (SNMP): A network management protocol of the IETF.

3.2.131 spectrum management system (SMS): A system, defined in [b-VECCHI], for managing the RF cable spectrum.

3.2.132 spread symbol or spreading interval: 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) symbol. One spread symbol = one spreading interval = 128 chips = one information (QAM) symbol.

3.2.133 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 CM's transmitter turns on to the time it turns off. There will generally be several spreader-off bursts in a spreader-off frame.

3.2.134 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.

3.2.135 spreading interval: Time to transmit a single complete S-CDMA spreading code, equal to the time to transmit 128 chips. Also, time to transmit a single information (QAM) symbol on an S-CDMA channel. See also spread symbol.

3.2.136 sub-channel: A logical channel sharing the same upstream spectrum (RF centre frequency and RF channel) with other logical channels.

3.2.137 sublayer: A subdivision of a layer in the open system interconnection (OSI) reference model.

3.2.138 subnetwork: Subnetworks are physically formed by connecting adjacent nodes with transmission links.

3.2.139 subnetwork access protocol (SNAP): An extension of the LLC header to accommodate the use of 802-type networks as IP networks.

3.2.140 subscriber: See "end user".

3.2.141 subsplit: A frequency-division scheme that allows bidirectional traffic on a single cable. Reverse path signals come to the headend from 5 to 30 MHz (up to 42 MHz on extended subsplit systems). Forward path signals go from the headend from 50 or 54 MHz to the upper frequency limit of the cable network.

3.2.142 subsystem: An element in a hierarchical division of an open system that interacts directly with elements in the next higher division or the next lower division of that open system.

3.2.143 system clock period: The period of the 10.24 MHz system clock, nominally 97.65625 ns.

3.2.144 systems management: Functions in the application layer related to the management of various open systems interconnection (OSI) resources and their status across all layers of the OSI architecture.

3.2.145 tick: 6.25-microsecond time intervals that are the reference for upstream mini-slot definition and upstream transmission times.

3.2.146 tilt: Maximum difference in transmission gain of a cable television system over a given bandwidth (typically the entire forward operating frequency range).

3.2.147 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.

3.2.148 transmission control protocol (TCP): A transport-layer Internet protocol which ensures successful end-to-end delivery of data packets without error.

3.2.149 transmission convergence sublayer: A sublayer of the physical layer that provides an interface between the data link layer and the PMD sublayer.

3.2.150 transmission link: The physical unit of a subnetwork that provides the transmission connection between adjacent nodes.

3.2.151 transmission medium: The material on which information signals may be carried; e.g., optical fibre, coaxial cable and twisted-wire pairs.

3.2.152 transmission system: The interface and transmission medium through which peer physical layer entities transfer bits.

3.2.153 transmit on/off ratio: In multiple-access systems, the ratio between the signal powers sent to line when transmitting and when not transmitting.

3.2.154 transport stream: In MPEG-2, a packet-based method of multiplexing one or more digital video and audio streams having one or more independent time base(s) into a single stream.

3.2.155 trivial file-transfer protocol (TFTP): An Internet protocol for transferring files without the requirement for user names and passwords that is typically used for automatic downloads of data and software.

3.2.156 trunk cable: Cables that carry the signal from the headend to groups of subscribers. The cables can be either coaxial or fibre depending on the design of the system.

3.2.157 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.

3.2.158 upstream: The direction from the subscriber location toward the headend.

3.2.159 upstream channel descriptor (UCD): The MAC management message used to communicate the characteristics of the upstream physical layer to the cable modems.

4 Functional assumptions

This clause describes the characteristics of cable television plant to be assumed 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 clause.

This clause applies to the first technology option referred to in clause 1.1. For the second and third options, refer to Annexes F and J, respectively.

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

4.1 Broadband access network

A coaxial-based broadband access network is assumed. This may take the form of either an all-coax or hybrid fibre/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 analogue transmission. The key functional characteristics assumed in this Recommendation are the following:

- two-way transmission;
- a maximum optical/electrical spacing between the CMTS and the most distant CM of 100 miles, 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, although this would typically be limited to 15 miles.

4.2 Equipment assumptions

4.2.1 Frequency plan

In the downstream direction, the cable system is assumed to have a passband with a lower edge between 50 and 54 MHz and an upper edge that is implementation-dependent but is typically in the range of 300 to 864 MHz. Within that passband, NTSC analogue television signals in 6-MHz channels are assumed to be present on the standard, HRC or IRC frequency plans of [b-EIA 542], as well as other narrow-band and wideband digital signals.

In the upstream direction, the cable system may have a subsplit (5-30 MHz) or extended subsplit (5-40 or 5-42 MHz) passband. NTSC analogue television signals in 6-MHz channels may be present, as well as other signals.

4.2.2 Compatibility with other services

The CM and CMTS MUST coexist with the other services on the cable network. In particular:

a) they MUST be interoperable in the cable spectrum assigned for CMTS-CM interoperation while the balance of the cable spectrum is occupied by any combination of television and other signals; and

b) they 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.

The latter 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 or other service provider (minimal level of compatibility).

4.2.3 Fault isolation impact on other users

As the data-over-cable system is 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 and other services.

For the interpretation of harmful impact, see clause 4.2.2.

4.2.4 Cable system terminal devices

The CM MUST meet and SHOULD exceed all applicable national regulations for cable system termination devices and cable-ready consumer equipment. None of these national specific requirements may be used to relax any of the specifications contained elsewhere within this Recommendation.

4.3 RF channel assumptions

The data-over-cable system, configured with at least one set of defined physical-layer parameters (e.g., modulation, forward error correction, modulation rate, etc.) from the range of configuration settings described in this Recommendation, MUST be interoperable on cable networks having characteristics defined in this clause 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.

4.3.1 Transmission downstream

The RF channel transmission characteristics of the cable network in the downstream direction are described in Table 4-1. 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 4-1 assume average power in a bandwidth in which the impairment levels are measured in a standard manner for cable TV system. For analogue signal levels, the numbers in Table 4-1 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 Recommendation.

Table 4-1 – Assumed downstream RF channel transmission characteristics (see Note 1)

Parameter	Value
Frequency range	Cable system normal downstream operating range is from 50 MHz to as high as 860 MHz. However, the values in this table apply only at frequencies ≥88 MHz.
RF channel spacing (design bandwidth)	6 MHz
Transit delay from headend to most distant customer	≤0.800 ms (typically much less)
Carrier-to-noise ratio in a 6 MHz band	Not less than 35 dB (see Notes 2 and 3)
Carrier-to-composite triple beat distortion ratio	Not less than 41 dB (see Notes 2 and 3)
Carrier-to-composite second order distortion ratio	Not less than 41 dB (see Notes 2 and 3)
Carrier-to-cross-modulation ratio	Not less than 41 dB (see Notes 2 and 3)
Carrier-to-any other discrete interference (ingress)	Not less than 41 dB (see Notes 2 and 3)
Amplitude ripple	3 dB within the design bandwidth (see Note 2)
Group delay ripple in the spectrum occupied by the CMTS	75 ns within the design bandwidth (see Note 2)
Micro-reflections bound for dominant echo	−20 dBc @ ≤ 1.5 µs, −30 dBc @ > 1.5 µs, −10 dBc @ ≤ 0.5 µs, −15 dBc @ ≤ 1.0 µs (see Note 2)
Carrier hum modulation	Not greater than −26 dBc (5%) (see Note 2)
Burst noise	Not longer than 25 µs at a 10-Hz average rate (see Note 2)
Maximum analogue video carrier level at the CM input	17 dBmV
Maximum number of analogue carriers	121

NOTE 1 – Transmission is from the headend combiner to the CM input at the customer location.

NOTE 2 – Measurement methods defined in [b-NCTA] or [b-CableLabs2].

NOTE 3 – Measured relative to a QAM signal that is equal to the nominal video level in the plant.

4.3.2 Transmission upstream

The RF channel transmission characteristics of the cable network in the upstream direction are described in Table 4-2. All conditions are present concurrently. No combination of the following parameters will exceed any stated interface limit defined elsewhere in this Recommendation.

Table 4-2 – Assumed upstream RF channel transmission characteristics (see Note 1)

Parameter	Value
Frequency range	5 to 42 MHz edge to edge
Transit delay from the most distant CM to the nearest CM or CMTS	≤ 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 (Note 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 (Notes 3 and 4)
Amplitude ripple 5-42 MHz	0.5 dB/MHz
Group delay ripple 5-42 MHz	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
NOTE 1 – Transmission is from the CM output at the customer location to the headend. NOTE 2 – 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. NOTE 3 – Amplitude and frequency characteristics sufficiently strong to partially or wholly mask the data carrier. NOTE 4 – Impulse noise levels more prevalent at lower frequencies (<15 MHz).	

4.3.2.1 Availability

Typical cable network availability is considerably greater than 99%.

4.4 Transmission levels

The nominal power level of the downstream CMTS signal(s) within a 6-MHz channel is targeted to be in the range –10 dBc to –6 dBc relative to analogue video carrier level and will normally not exceed analogue video carrier level. 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.

4.5 Frequency inversion

There will be no frequency inversion in the transmission path in either the downstream or 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.

5 Communication protocols

This clause provides a high-level overview of the communication protocols that must be used in the data-over-cable system. Detailed specifications for the physical media dependent, downstream transmission, and media access control sublayers are provided in clauses 6, 7, and 8, respectively.

5.1 Protocol stack

The CM and CMTS operate as forwarding agents and also as end-systems (hosts). The protocol stacks used in these modes differ as shown below.

The principal function of the cable modem system is to transmit Internet protocol (IP) packets transparently between the headend and the subscriber location. Certain management functions also ride on IP, so that the protocol stack on the cable network is as shown in Figure 5-1 (this does not restrict the generality of IP transparency between the headend and the customer). These management functions include, for example, supporting spectrum management functions and the downloading of software.

5.1.1 CM and CMTS as hosts

CMs and CMTSs operate as IP and LLC hosts in terms of [IEEE 802] for communication over the cable network. The protocol stack at the CM and CMTS RF interfaces is shown in Figure 5-1.

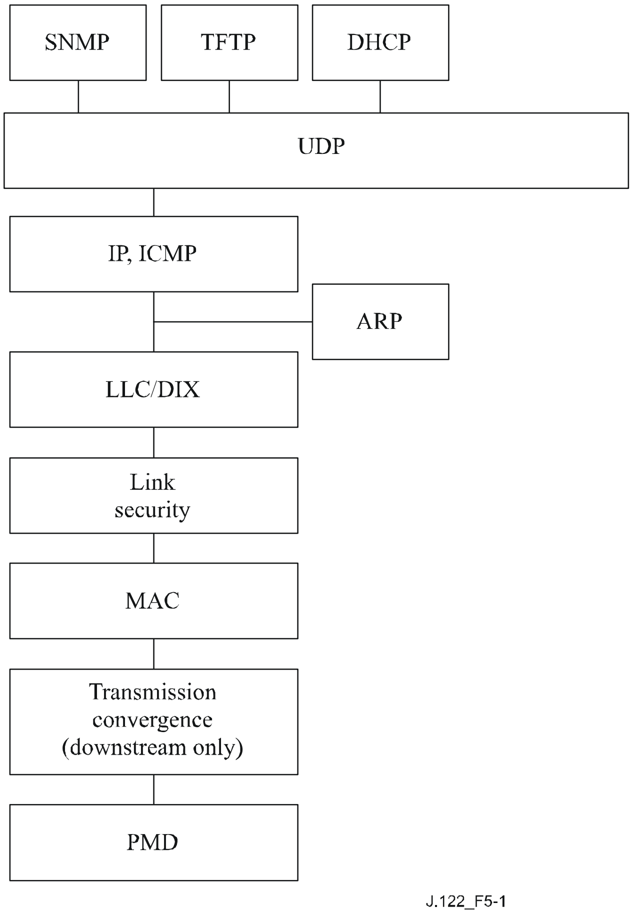

Figure 5-1 – Protocol stack on the RF interface

The CM and CMTS MUST function as IP hosts. As such, the CM and CMTS MUST support IP and ARP over DIX link-layer framing ("DIX link-layer framing" refers to the "type interpretation" of the length/type field in [ISO/IEC 8802-3]). The CMTS MUST NOT transmit frames that are smaller than the DIX 64-byte minimum on a downstream channel[1]. However, the CM MAY transmit frames that are smaller than the DIX 64-byte minimum on an upstream channel.

The CM and CMTS MAY also support IP and ARP over SNAP framing [RFC 1042].

The CM and CMTS also MUST function as LLC hosts. As such, the CM and CMTS MUST respond appropriately to TEST and XID requests per [ISO/IEC 8802-2].

5.1.2 Data forwarding through the CM and CMTS

5.1.2.1 General

Data forwarding through the CMTS MAY be transparent bridging[2], or MAY employ network-layer forwarding (routing, IP switching) as shown in Figure 5-2.

[1] Except as a result of payload header suppression. Refer to clause 10.4.

[2] With the exception that, for packet PDUs less than 64 bytes to be forwarded from the upstream RFI, a CMTS MUST pad out the packet PDU and recompute the CRC.

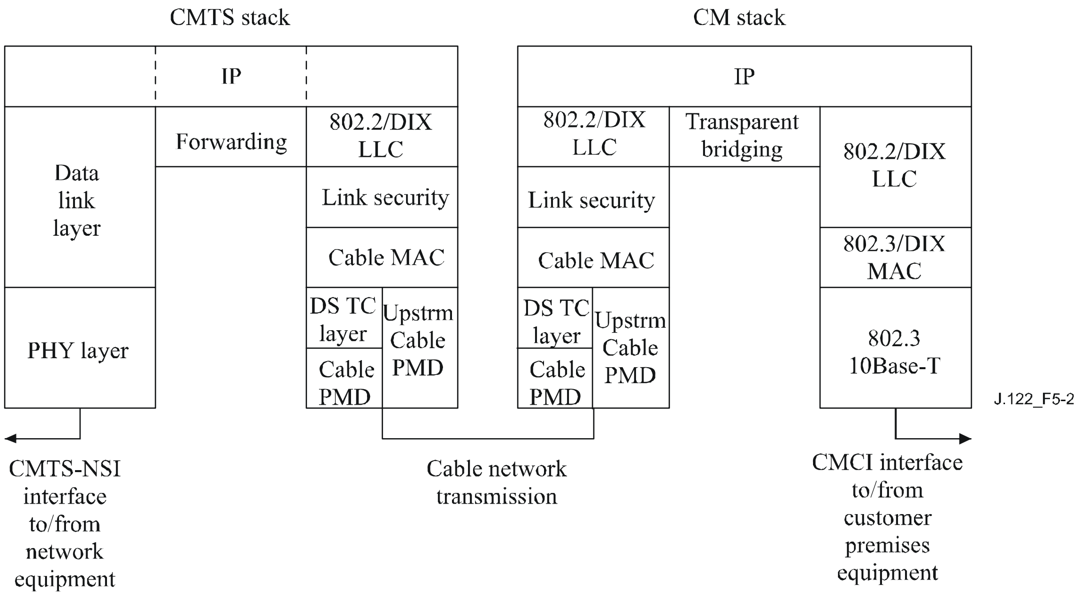

Figure 5-2 – Data forwarding through the CM and CMTS

Data forwarding through the CM is link-layer transparent bridging, as shown in Figure 5-2. Forwarding rules are similar to [b-ISO/IEC 10038] with the modifications described in clauses 5.1.2.2 and 5.1.2.3. This allows the support of multiple network layers.

Forwarding of IP traffic MUST be supported. Other network layer protocols MAY be supported. The ability to restrict the network layer to a single protocol such as IP MUST be supported.

The IEEE 802.1D spanning tree protocol of [b-ISO/IEC 10038], with the modifications described in Annex E, MAY be supported by CMs intended for residential use. CMs intended for commercial use MUST support this version of spanning tree. CMs and CMTSs MUST include the ability to filter (and disregard) IEEE 802.1D BPDUs.

This Recommendation assumes that CMs intended for residential use will not be connected in a configuration which would create network loops such as that shown in Figure 5-3.

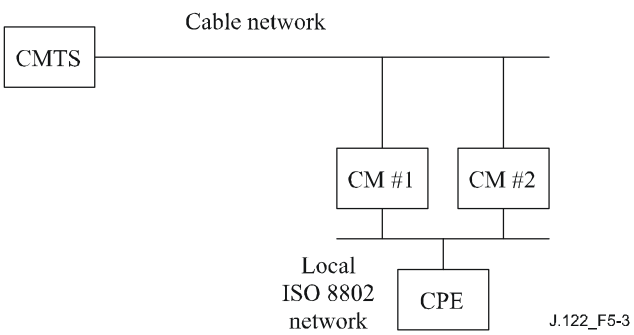

Figure 5-3 – Example condition for network loops

Although provisions exist in this Recommendation for frames to be passed from a higher-layer entity to be forwarded by the cable modem, these frames MUST be treated identically to frames arriving at the CPE port. In particular, all of the forwarding rules defined in clause 5.1.2.3 MUST apply to these frames

5.1.2.2 CMTS forwarding rules

At the CMTS, if link-layer forwarding is used, then it MUST conform to the following general IEEE 802.1D guidelines:

- Link-layer frames MUST NOT be duplicated.
- Stale frames (those that cannot be delivered in a timely fashion) MUST be discarded.
- Link-layer frames, on a given service flow (refer to clause 8.1.2.3), MUST be delivered in the order they are received.

The address-learning and -aging mechanisms used are vendor-dependent.

If network-layer forwarding is used, then the CMTS should conform to IETF router requirements [b-RFC 1812] with respect to its CMTS-RFI and CMTS-NSI interfaces.

Conceptually, the CMTS forwards data packets at two abstract interfaces: between the CMTS-RFI and the CMTS-NSI, and between the upstream and downstream channels. The CMTS MAY use any combination of link-layer (bridging) and network-layer (routing) semantics at each of these interfaces. The methods used at the two interfaces need not be the same.

Forwarding between the upstream and downstream channels within a MAC layer differs from traditional LAN forwarding in that:

- A single channel is simplex, and cannot be considered a complete interface for most protocol (e.g., 802.lD spanning tree, routing information protocol per [b-RFC 1058]) purposes.
- Upstream channels are essentially point-to-point, whereas downstream channels are shared-media.
- Policy decisions may override full connectivity.

For these reasons, an abstract entity called the MAC forwarder exists within the CMTS to provide connectivity between stations within a MAC domain (see clause 5.2).

5.1.2.3 CM forwarding rules

Data forwarding through the CM is link-layer bridging with the following specific rules.

5.1.2.3.1 CPE MAC address acquisition

- The CM MUST acquire Ethernet MAC addresses of connected CPE devices, either from the provisioning process or from learning, until the CM acquires its maximum number of CPE MAC addresses (a device-dependent value). Once the CM acquires its maximum number of CPE MAC addresses, then newly discovered CPE MAC addresses MUST NOT replace previously acquired addresses. The CM must support acquisition of at least one CPE MAC address.
- The CM MUST allow configuration of CPE addresses during the provisioning process (up to its maximum number of CPE addresses) to support configurations in which learning is not practical nor desired.
- Addresses provided during the CM provisioning MUST take precedence over learned addresses.
- CPE addresses MUST NOT be aged out.
- In order to allow modification of user MAC addresses or movement of the CM, addresses are not retained in non-volatile storage. On a CM reset (e.g., power cycle), all provisioned and learned addresses MUST be discarded.

5.1.2.3.2 Forwarding

CM forwarding in both directions MUST conform to the following general IEEE 802.1D guidelines:

- Link-layer frames MUST NOT be duplicated.
- Stale frames (those that cannot be delivered in a timely fashion) MUST be discarded.
- Link-layer frames MUST be delivered in the order that they are received on a given service flow (refer to clause 8.1.2.3). In the upstream direction, the CM may perform one or more frame/packet processing functions on frames received from the CMCI prior to classifying them to a service flow. In the downstream direction, the CM may perform one or more frame/packet processing functions on frames received from the HFC prior to transmitting them on the CMCI. Example processing functions include: DOCSIS protocol filtering as specified in [DOCS5] clause 7.3, a policy-based filtering service as described in clause 10.1.6.1 and Appendix I, and priority-based queuing to support 802.1P/Q services.
- Link-layer frames, on a given service flow (refer to clause 8.1.2.3), MUST be delivered in the order they are received.

Cable-network-to-CMCI forwarding MUST follow the following specific rules:

- Frames addressed to unknown destinations MUST NOT be forwarded from the cable port to the CPE ports.
- Broadcast frames MUST be forwarded to the CPE ports, unless they are from source addresses which are provisioned or learned as supported CPE devices, in which case they MUST NOT be forwarded.
- The forwarding of multicast is controlled by administratively set filters and by either IGMP (refer to clause 5.3.1) or static provisioning (refer to clause C.1.2.12). Multicast frames MUST NOT be forwarded unless both filtering and either IGMP or static multicast provisioning are in a permissive state.

CPE-to-cable-network forwarding MUST follow the following specific rules:

- Frames addressed to unknown destinations MUST be forwarded from all CPE ports to the cable port.
- Broadcast frames MUST be forwarded to the cable port.
- Multicast frames MUST be forwarded to the cable port in accordance with filtering configuration settings specified by the cable operator's operations and business support systems.
- Frames from source addresses other than those provisioned or learned as supported CPE devices MUST NOT be forwarded.
- Other (non-supported) CPE source addresses MUST be learned from all CPE ports and this information used to filter local traffic as in a traditional learning bridge.
- Frames addressed to destination addresses that are learned from all CPE ports MUST be filtered as local traffic.

5.2 The MAC forwarder

The MAC forwarder is a MAC sublayer that resides on the CMTS just below the MAC service access point (MSAP) interface, as shown in Figure 5-4. It is responsible for delivering upstream frames to:

- one or more downstream channels;
- the MSAP interface.

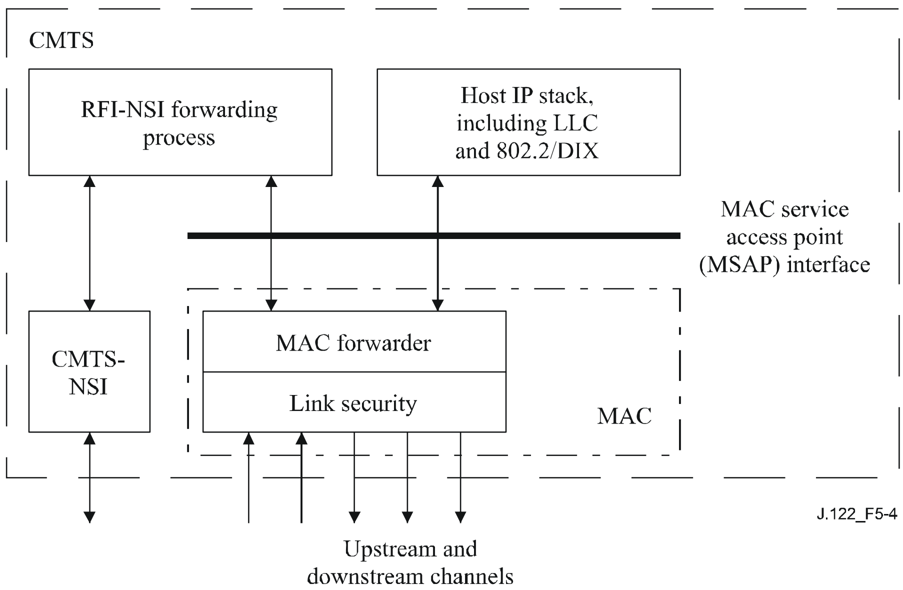

Figure 5-4 – MAC forwarder

In Figure 5-4, the LLC sublayer and link security sublayers of the upstream and downstream channels on the cable network terminate at the MAC forwarder.

The MSAP interface user may be the NSI-RFI forwarding process or the CMTS's host protocol stack.

Delivery of frames may be based on data-link-layer (bridging) semantics, network-layer (routing) semantics, or some combination. Higher-layer semantics may also be employed (e.g., filters on UDP port numbers). The CMTS MUST provide IP connectivity between hosts attached to cable modems, and must do so in a way that meets the expectations of Ethernet-attached customer equipment. For example, the CMTS must either forward ARP packets or it must facilitate a proxy ARP service. The CMTS MAC Forwarder MAY provide service for non-IP protocols.

Note that there is no requirement that all upstream and downstream channels be aggregated under one MSAP as shown above. The vendor could just as well choose to implement multiple MSAPs, each with a single upstream and downstream channel.

5.2.1 Rules for data-link-layer forwarding

The requirements in this clause apply if the MAC forwarder is implemented using only data-link-layer semantics.

Delivery of frames is dependent on the destination address within the frame. The means of learning the location of each address is vendor-dependent, and MAY include:

- transparent-bridging-like source-address learning and aging;
- gleaning from MAC registration request messages;
- administrative means.

If the destination address of a frame is unicast, and that address is associated with a particular downstream channel, then the frame MUST be forwarded to that channel[3].

[3] Vendors may implement extensions, similar to static addresses in 802.1D/ISO/IEC 10038 bridging, that cause such frames to be filtered or handled in some other manner.

If the destination address of a frame is unicast, and that address is known to reside on the other (upper) side of the MSAP interface, then the frame MUST be delivered to the MSAP interface.

If the destination address is broadcast, multicast[4], or unknown, the frame MUST be delivered to both the MSAP and to all downstream channels (with the exception of the multicast forwarding rules in clause 5.3.1.2).

Delivery rules are similar to those for transparent bridging:

- Frames MUST NOT be duplicated.
- Frames that cannot be delivered in a timely fashion MUST be discarded.
- The frame check sequence SHOULD be preserved rather than regenerated.
- Frames, on a given service flow (refer to clause 8.1.2.3), MUST be delivered in the order they are received.

5.3 Network layer

As stated above, the purpose of the data-over-cable system is to transport IP traffic transparently through the system.

The network layer protocol is the Internet protocol (IP) version 4, as defined in [RFC 791], and migrating to IP version 6.

This Recommendation imposes no requirements for reassembly of IP packets.

5.3.1 Requirements for IGMP management

There are two basic modes of IGMP capability that are applicable to a DOCS 2.0 device (CMTS and CM). The first mode is a *passive* operation in which the device selectively forwards IGMP based upon the known state of multicast session activity on the subscriber side (an example of this is described in Appendix V). In *passive* mode, the device derives its IGMP timers based on the rules specified in [DOCS11]. The second mode is an *active* operation in which the device terminates and initiates IGMP based upon the known state of multicast session activity on the subscriber side. One example of the latter, active, mode is commonly referred to as an IGMP-proxy implementation side (as described in [b-ID-IGMP]). A more complete example of an active IGMP device is that of a multicast router.

Active and passive IGMP devices MUST support IGMP v2 [RFC 2236].

5.3.1.1 IGMP timer requirements

The following IGMP timer requirements apply only when the device (CMTS/CM) is operating in passive IGMP mode:

- The device MUST NOT require any specific configuration for the associated multicast timer values and MUST be capable of adhering to the timers specified in this clause.
- The device MAY provide configuration control that overrides the default values of these timers.
- The device MUST derive the membership query interval by looking at the inter-arrival times of the membership query messages. Formally: If $n < 2$, MQI = 125, else MQI = MAX (125, $MQ_n - MQ_{n-1}$), where MQI is the membership query interval in seconds, n is the number of membership queries seen, and MQ_n is the epoch time at which the nth membership query was seen to the nearest second.

4 All multicasts, including 802.1D/ISO 10038 spanning tree bridge BPDUs, MUST be forwarded.

- The query response interval is carried in the membership query packet. The query response interval MUST be assumed to be 10 seconds if not otherwise set (or set to 0) in the membership query packet.

5.3.1.2 CMTS rules

- If link-layer forwarding of multicast packets is used, the CMTS MUST forward all membership queries on all downstream channels using the appropriate 802.3 multicast group (e.g., 01:00:5E:xx:xx:xx where xx:xx:xx are the low-order 23 bits of the multicast address expressed in hex notation. Refer to [b-IMA]).
- The CMTS MUST forward the first copy of solicited and unsolicited membership reports for any given group received on its upstream RF interface to all of its downstream RF interfaces. However, if membership is managed on a per downstream RF interface basis, membership reports and IGMP v2 leave messages MAY be forwarded only on the downstream interface to which the reporting CPE's CM is connected.
- The CMTS SHOULD suppress the transmission of additional membership reports (for any given group) downstream for at least the query response interval. If the CMTS uses data-link-layer forwarding, it MUST also forward the membership report out all appropriate network side interfaces.
- The CMTS SHOULD suppress the downstream transmission of traffic to any IP multicast group that does not have subscribers on that downstream RF interface (subject to any administrative controls).
- If the CMTS performs network-layer forwarding of multicast packets, it MUST support active IGMP mode.
- If link-layer forwarding of multicast packets is used, the CMTS SHOULD support passive IGMP mode and MAY support Active IGMP mode.

5.3.1.3 CM rules

The CM MUST support IGMP with the cable-specific rules specified in this clause.

The CM MUST implement the passive IGMP mode. Additionally, the CM MAY implement the active IGMP mode. If it implements the active IGMP mode, the CM MUST support a capability to switch between modes.

5.3.1.3.1 Multicast forwarding requirements

The following requirements apply to both passive and active modes of IGMP operations:

- The CM MUST NOT forward membership queries from its CPE interface to its RF interface.
- The CM MUST NOT forward membership reports or IGMP v2 leaves received on its RF interface to its CPE interface.
- The CM MUST NOT forward multicast traffic from its RF interface to its CPE interface unless a device on its CPE interface is a member of that IP multicast group.
- The CM MUST forward multicast traffic from its CPE interface to its RF interface unless administratively (via configuration or other mechanism) prohibited.
- As a result of receiving a membership report on its CPE interface, the CM MUST begin forwarding traffic for the appropriate IP multicast group. The CM MUST stop forwarding multicast traffic from the RF to the CPE side whenever the CM has not received a membership report from the CPE side for more than the membership interval, which is $(2 \times MQI) + QRI$, where MQI is the membership query interval and QRI is the query response interval.

- The CM MAY stop forwarding traffic from the RF to the CPE side for a particular multicast group prior to the expiration of the membership interval (see above) if it can determine (for example, via an IGMP LEAVE message and the appropriate protocol exchange) that there are no CPE devices subscribed to that particular group.

The following requirements apply only when the CM is operating in passive IGMP mode:

- The CM MUST forward traffic for the ALL-HOSTS multicast group from its RF interface to its CPE interface unless administratively prohibited. The CPE MUST always be considered a member of this group. In particular, the CM MUST forward ALL-HOSTS group queries that pass permit filters on its RF interface to its CPE interface.
- Upon receiving a membership report on its CPE interface, the CM MUST start a random timer between 0 and 3 seconds. During this time period, the CM MUST discard any additional membership reports received in its CPE interface for the associated multicast group. If the CM receives a membership report on its HFC interface for the associated multicast group, the CM MUST discard the membership report received on its CPE interface. If the random timer expires without the reception of a membership report on the CMs HFC interface, the CM MUST transmit the membership report received on its CPE interface.

The following requirements apply only when the CM is operating in active IGMP mode:

- The CM MUST implement the host portion of the IGMP v2 protocol [RFC 2236] on its RF interface for CPEs with active groups and MUST NOT act as a querier on its RF interface.
- The CM MUST act as an IGMP v2 querier on its CPE interface.
- If the CM has received a membership report on its downstream RF interface for groups active on the CMs CPE interface within the query response interval, it MUST suppress transmission on its upstream RF interface of such membership reports.
- The CM MUST suppress all subsequent membership reports for this group until such time as the CM receives a membership query (general or specific to this group) on its RF interface or a IGMP v2 leave is received for this group from the CPE interface.
- The CM MUST treat unsolicited membership reports (IGMP JOINs) from its CPE interface as a response to a membership query received on its RF interface. Upon receipt of this unsolicited JOIN from its CPE interface, the CM MUST start a random timer according to the host state diagram, specified in [RFC 2236], and MUST use a query response interval of 3 seconds. As specified above, if the CM receives a membership report on its RF interface for this group during this random time period, it MUST suppress transmission of this join on its upstream RF interface.
- On startup, the CM SHOULD send one or more general queries on its CPE interface (as described in [RFC 2236]) in order to quickly and reliably determine membership information for attached CPEs.

NOTE – Nothing in this clause would prohibit the CM from being specifically configured to not forward certain multicast traffic as a matter of network policy.

5.4 Above the network layer

The subscribers will be able to use the transparent IP capability as a bearer for higher-layer services. Use of these services will be transparent to the CM.

In addition to the transport of user data, there are several network management and operation capabilities which depend upon the network layer. These include:

- Simple network management protocol (SNMP) [RFC 1157] MUST be supported for network management.

- Trivial file transfer protocol (TFTP) [RFC 1350], a file transfer protocol, MUST be supported for downloading operational software and configuration information, as modified by TFTP timeout interval and transfer size options [RFC 2349].
- Dynamic host configuration protocol (DHCP) [RFC 2131], a framework for passing configuration information to hosts on a TCP/IP network, MUST be supported.
- Time protocol, [RFC 868], MUST be supported to obtain the time of day.

DHCP, TFTP and ToD client messages generated by the CM MUST only be sent via the RF Interface. DHCP, TFTP and ToD client messages include DHCPDISCOVER, DHCPREQUEST, DHCPDECLINE, DHCPRELEASE, DHCPINFORM, TFTP-RRQ, TFTP-ACK and ToD request.

The CM's DHCP, TFTP and ToD client MUST ignore DHCP, TFTP and ToD server messages received on the CMCI port. DHCP, TFTP and ToD server messages include: DHCPOFFER, DHCPACK, DHCPNAK, TFTP-DATA and ToD time message.

5.5 Data link layer

The data link layer is divided into sublayers in accordance with [IEEE 802], with the addition of link-layer security in accordance with [DOCS8]. The sublayers, from the top, are:

- logical link control (LLC) sublayer (class 1 only);
- link-layer security sublayer;
- media access control (MAC) sublayer.

5.5.1 LLC sublayer

The LLC sublayer MUST be provided in accordance with [ISO/IEC 10039]. Address resolution MUST be used as defined in [RFC 826]. The MAC-to-LLC service definition is specified in [ISO/IEC 10039].

5.5.2 Link-layer security sublayer

Link-layer security MUST be provided in accordance with [DOCS8].

5.5.3 MAC sublayer

The MAC sublayer defines a single transmitter for each downstream channel – the CMTS. All CMs listen to all frames transmitted on the downstream channel upon which they are registered and accept those where the destinations match the CM itself or CPEs reached via the CMCI port. CMs can communicate with other CMs only through the CMTS.

The upstream channel is characterized by many transmitters (CMs) and one receiver (the CMTS). Time in the upstream channel is slotted, providing for time division multiple access at regulated time ticks. The CMTS provides the time reference and controls the allowed usage for each interval. Intervals may be granted for transmissions by particular CMs, or for contention by all CMs. CMs may contend to request transmission time. To a limited extent, CMs may also contend to transmit actual data. In both cases, collisions can occur and retries are used.

Clause 8 describes the MAC-sublayer messages from the CMTS which direct the behaviour of the CMs on the upstream channel, as well as messaging from the CMs to the CMTS.

5.5.3.1 MAC service definition

The MAC sublayer service definition is in Appendix I.

5.6 Physical layer

The physical (PHY) layer is comprised of two sublayers:

- transmission convergence sublayer (present in the downstream direction only);

• physical media dependent (PMD) sublayer.

5.6.1 Downstream transmission convergence sublayer

The downstream transmission convergence sublayer exists in the downstream direction only. It provides an opportunity for additional services over the physical-layer bitstream. These additional services might include, for example, digital video. Definition of any such additional services is beyond the scope of this Recommendation.

This sublayer is defined as a continuous series of 188-byte MPEG [ITU-T H.222.0] packets, each consisting of a 4-byte header followed by 184 bytes of payload. The header identifies the payload as belonging to the data-over-cable MAC. Other values of the header may indicate other payloads. The mixture of payloads is arbitrary and controlled by the CMTS.

The downstream transmission convergence sublayer is defined in clause 7.

5.6.2 PMD sublayer

The physical media dependent sublayer is defined in clause 6.

5.6.2.1 Interface points

Three RF interface points are defined at the PMD sublayer:

a) downstream output on the CMTS;

b) upstream input on the CMTS;

c) cable in/out at the cable modem.

Separate downstream output and upstream input interfaces on the CMTS are required for compatibility with typical downstream and upstream signal combining and splitting arrangements in headends.

6 Physical media dependent sublayer specification

6.1 Scope

This Recommendation 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 Recommendation 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 Recommendation to imply any specific implementation.

This clause applies to the first technology option referred to in clause 1.1. For the second and third options, refer to Annexes F and J, respectively.

Whenever any reference in this clause to spurious emissions conflicts with any legal requirement for the area of operation, the latter shall take precedence.

6.2 Upstream

6.2.1 Overview

The upstream physical media dependent (PMD) sublayer uses a 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.

Frequency division multiple access (FDMA) indicates that multiple RF channels are assigned in the upstream band. A CM transmits on a single RF channel unless reconfigured to change channels. Time division multiple access (TDMA) indicates that upstream transmissions have a burst nature. A given RF channel is shared by multiple CMs via the dynamic assignment of time slots. Synchronous code division multiple access (S-CDMA) 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 Recommendation, the following naming conventions are used. For TDMA, the term "modulation rate" refers to the RF channel symbol rate (160 to 5120 ksymb/s). For S-CDMA, the term "chip rate", which is the modulation rate (1280 to 5120 kHz) of a single bit of the S-CDMA spreading code, may be used interchangeably with "modulation rate". 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 entity 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 are used to pad data blocks; this indicates data with zero-valued binary bits, which result in non-zero transmitted RF energy. In other cases a numerical zero is used; this denotes, for example, symbols that 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 centre of the last symbol of one burst and the centre 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 clause 6.2.19. 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 centred 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) narrow-band 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 clauses 6.2.5.1, 6.2.19 and A.2).

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.

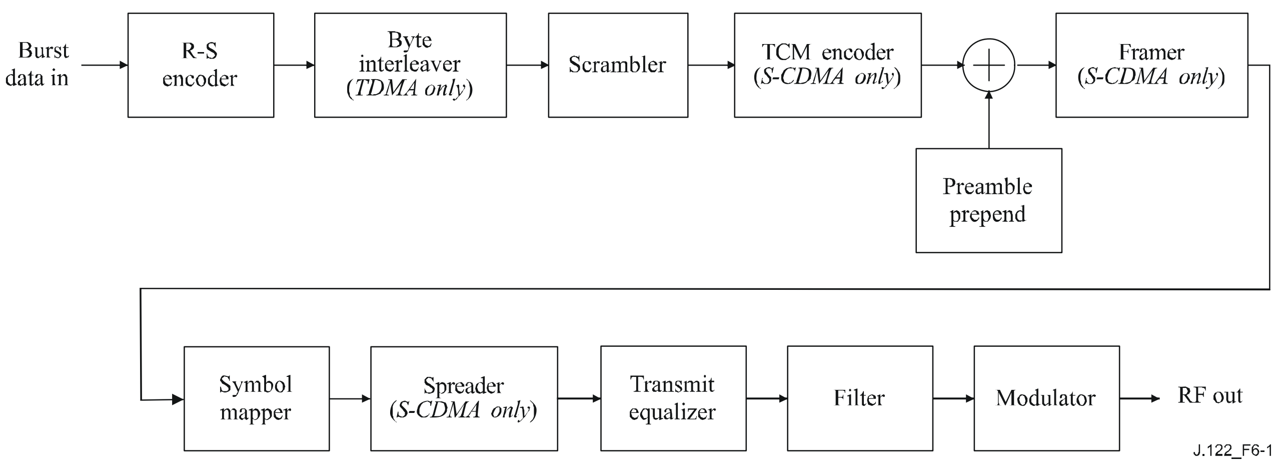

Figure 6-1 – Upstream signal-processing sequence

	Packet stream input
	↓
Block the data	Separate packet into information blocks (= data bytes in one codeword)
	↓
R-S encode	R-S (Reed-Solomon) encode each information block, using shortened codeword for last block if needed. R-S FEC can be turned off.
	↓
Byte interleave	R-S Byte interleave. R-S Byte interleaver can be turned off.
	↓
Scramble	Scramble (see Figure 6-7)
	↓
Preamble prepend	Prepend preamble symbols
	↓
Symbol map	Map the data stream into modulator symbols
	↓
Transmit equalize	Pre-equalize the symbol stream
	↓
Filter	Filter symbol stream for spectral shaping
	↓
Modulate	Modulate at precise times (QPSK; 8QAM; 16QAM; 32QAM; 64QAM)
	↓
	Output RF waveform burst

Figure 6-2 – TDMA upstream transmission processing

	Packet stream input
	↓
Block the data	Separate packet into information blocks (= data bytes in one codeword)
	↓
R-S encode	R-S (Reed-Solomon) encode each information block, using shortened codeword for last block if needed. R-S FEC can be turned off.
	↓
Scramble	Scramble (see Figure 6-7)
	↓
TCM encode	TCM (trellis-coded modulation) encode the bytes. TCM can be turned off.
	↓
Preamble prepend	Prepend preamble symbols
	↓
S-CDMA framer	Frame and interleave the data into mini-slots
	↓
Symbol map	Map the data stream into modulator symbols
	↓
S-CDMA spreader	Spread the symbols. For spreader-off on S-CDMA channels, spreader can be turned off.
	↓
Transmit equalize	Pre-equalize the symbol stream
	↓
Filter	Filter symbol stream for spectral shaping
	↓
Modulate	Modulate at precise times (QPSK; 8QAM; 16QAM; 32QAM; 64QAM, 128QAM/TCM only)
	↓
	Output RF waveform burst

Figure 6-3 – S-CDMA upstream transmission processing

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 clause 6.2.8. The S-CDMA framer maps mini-slots into code resources and provides interleaving of data symbols and is described in clause 6.2.11. The S-CDMA spreader spreads S-CDMA framed symbols for transmission and is described in clause 6.2.14.

6.2.3 Modulation formats

The modulation formats listed here specify requirements for J.122-compliant equipment. Cable operators are free to configure the modulation format to best address their system characteristics and application requirements.

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

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

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

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

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

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

The upstream demodulator MAY support QPSK, 8QAM, 16QAM, 32QAM, 64QAM, and 128QAM 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:

$$g(x) = \left(x + \alpha^0\right)\left(x + \alpha^1\right)\ldots\left(x + \alpha^{2T-1}\right)$$

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 codewords 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 clause 6.2.5.1.3.

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 per clause 8.2.1.3.

6.2.5 R-S frame structure

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 codewords. Example 1 illustrates the fixed codeword-length mode, and example 2 illustrates the shortened last codeword mode. These modes are defined in clause 6.2.5.1.

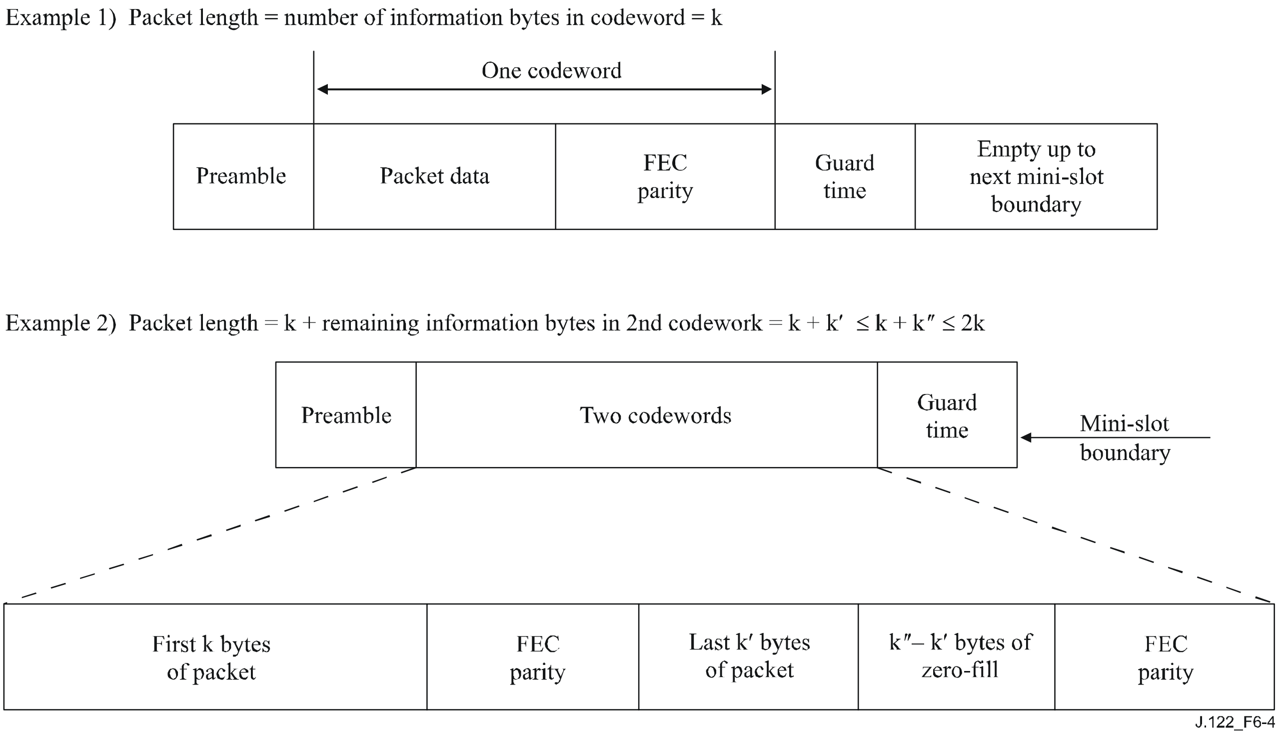

Figure 6-4 – Example frame structures with flexible burst length mode

6.2.5.1 R-S codeword length

When R-S FEC is enabled, the CM operates in either fixed-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 clauses is to define rules and conventions such that CMs request the proper number of mini-slots and 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 clauses 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 preamble, 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 (refer to Annex A).
4-6, 9-11	Number of mini-slots allocated.

6.2.5.1.2 Fixed codeword length

With the fixed-length codewords, after all the data are encoded, zero-fill will occur in this codeword if necessary to reach the assigned k data bytes per codeword. Additionally, zero-fill MUST continue up to the point when no additional fixed-length codewords 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) codewords. 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 zero-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 codewords 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 zero-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´ zero-fill bytes are necessary with $16 \le k'' \le k$ and $k' \le k''$.

More generally, the CM MUST zero-fill data until the point when no additional fixed-length codewords 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), and then, if possible, a shortened last codeword of zero-fill MUST be inserted to fit into the last mini-slot.

If, after zero-fill of additional codewords 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 shall not create this last shortened codeword.

6.2.5.2 R-S FEC disabled

When T = 0 (no FEC parity bytes), the R-S encoder SHOULD zero-fill in full bytes to the end of the burst specified in clause 6.2.5.1.1 above, accounting for preamble, return-to-zero bits, and guard-time symbols (if any).

6.2.6 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 codewords, 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).

6.2.6.1 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 floor (2048/N_r) for 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 $N_r \times I_r \leq 2048$ (in other words, for a given N_r, the maximal value of I_r is $I_{r,max}$ = 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 = l, 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.

6.2.6.2 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.

6.2.6.2.1 Fixed mode

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

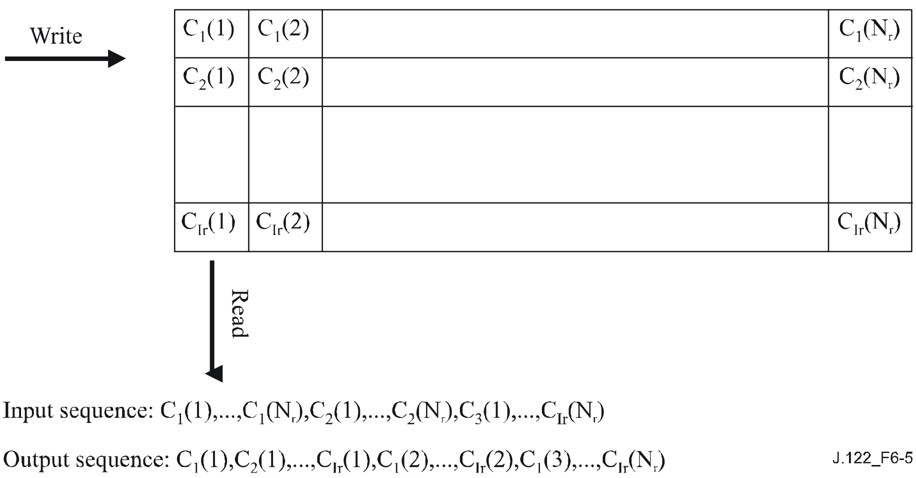

Figure 6-5 – Byte interleaver operation

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.

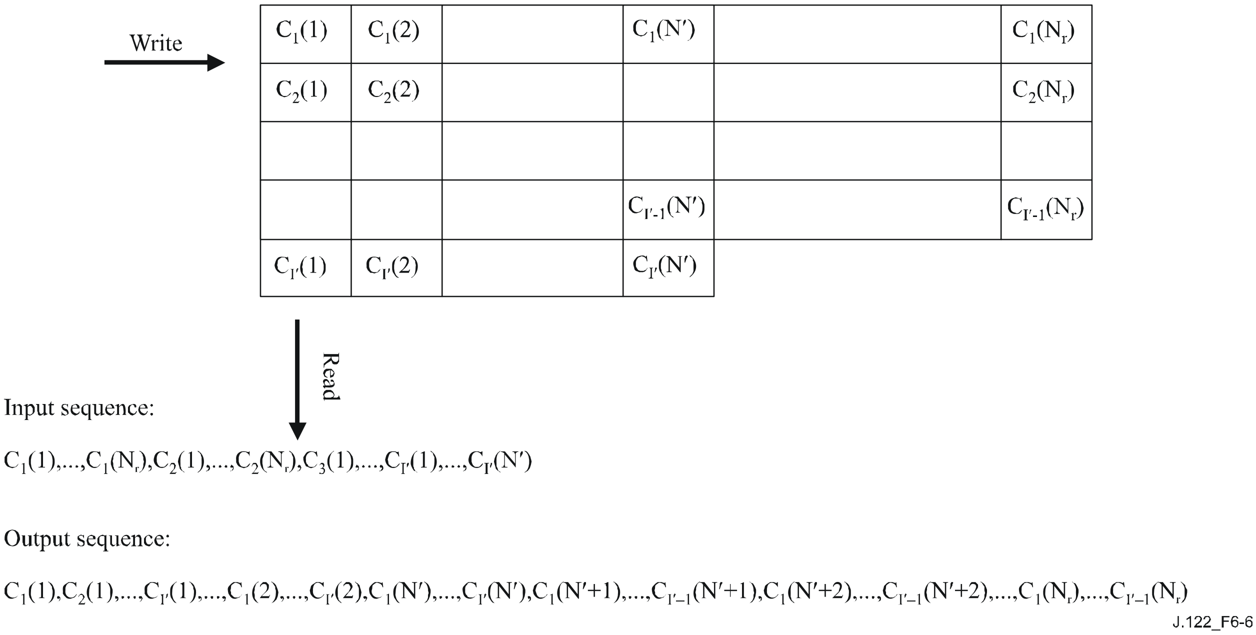

Figure 6-6 – Interleaver operation for last interleaver block (with shortened last codeword)

6.2.6.2.2 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 $N_r \times I_r^{(i)}$ bytes (i.e., a block of $I_r^{(i)}$ R-S codewords). 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 clause 6.2.6.2.2.1).

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′).

6.2.6.2.2.1 Dynamic mode calculations

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

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

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

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

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

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

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

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

6.2.7 Scrambler (randomizer)

The upstream modulator MUST implement a scrambler (shown in Figure 6-7) where the 15-bit seed value MUST be arbitrarily programmable.

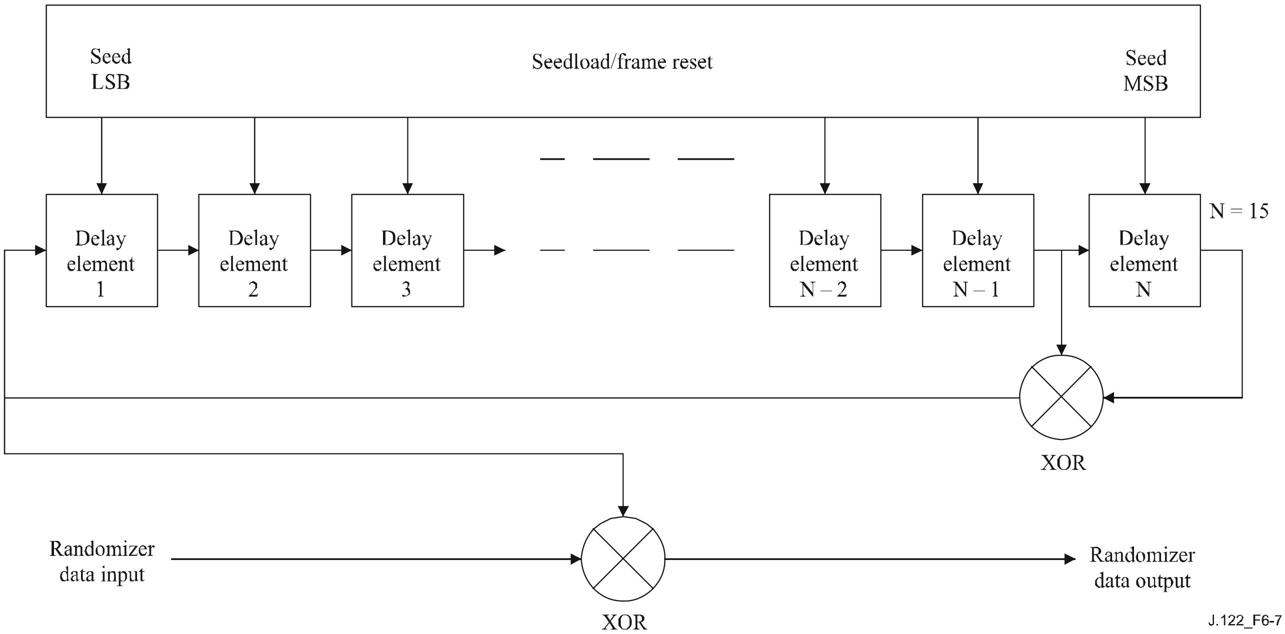

Figure 6-7 – Scrambler structure

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 $x^{15} + x^{14} + 1$.

6.2.8 TCM encoder

R-S symbol interleaving is commonly included between the TCM and R-S blocks to preserve coding gain in the presence of bursty 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, 8QAM, 16QAM, 32QAM, 64QAM and 128QAM constellations, respectively. Support of TCM in the CMTS is optional.

Figure 6-8 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.

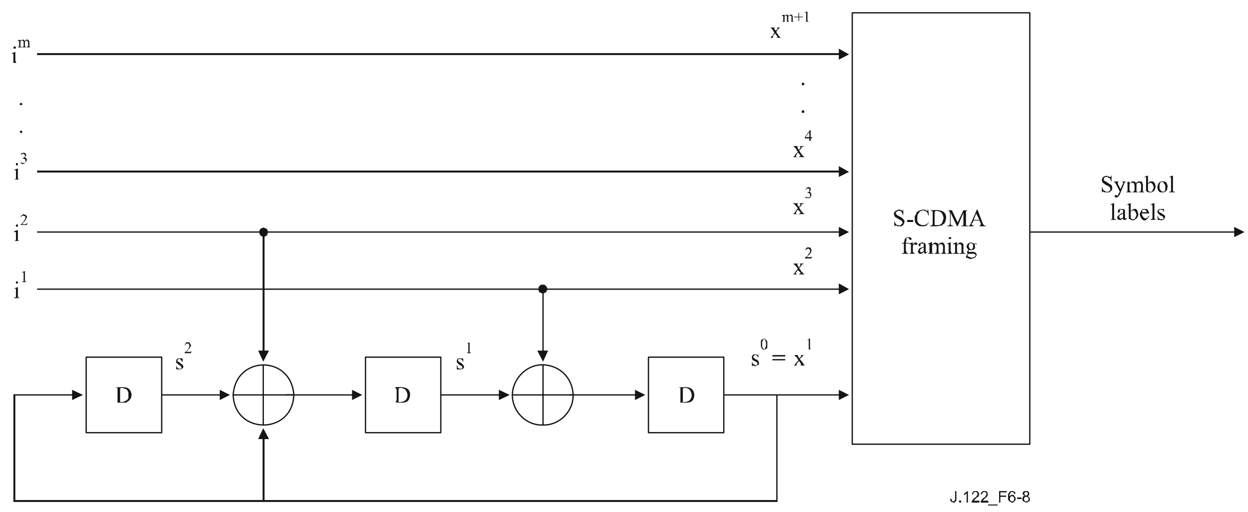

Figure 6-8 – Convolutional encoder

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 6-8, 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^2i^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 ≥ 3, the uncoded bits i^m,..., 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 ($0 \leq n_t \leq 2$).

6.2.8.1 Byte to TCM symbol mapping

The mapping of R-S bytes to TCM symbols is done such that each byte is mapped entirely to the uncoded bits i^m, ...,i^3, or entirely to the convolutional encoder input bits i^2i^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 6-9 for m = 1 to 6. In the figure bit i^m is at the top and bit i^1 is at the bottom.

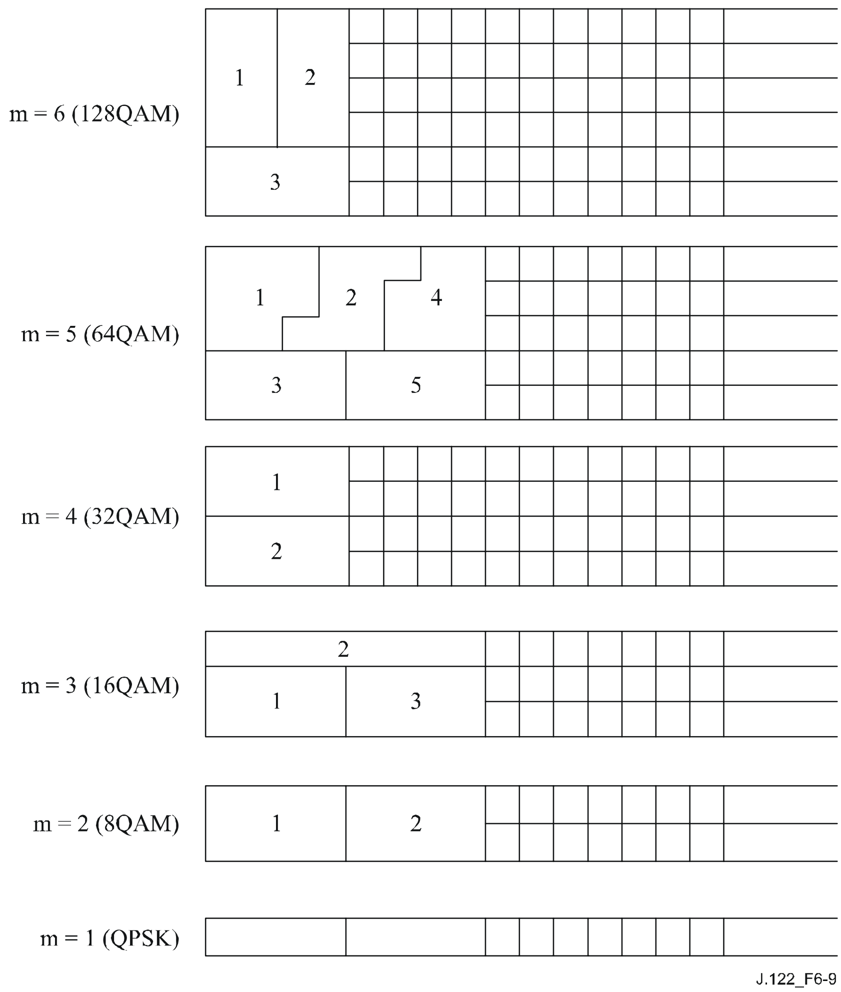

Figure 6-9 – Repetitive patterns of byte mapping to symbol map bits for TCM

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 6-10 illustrates the byte assignments for Trellis-coded 64QAM 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.

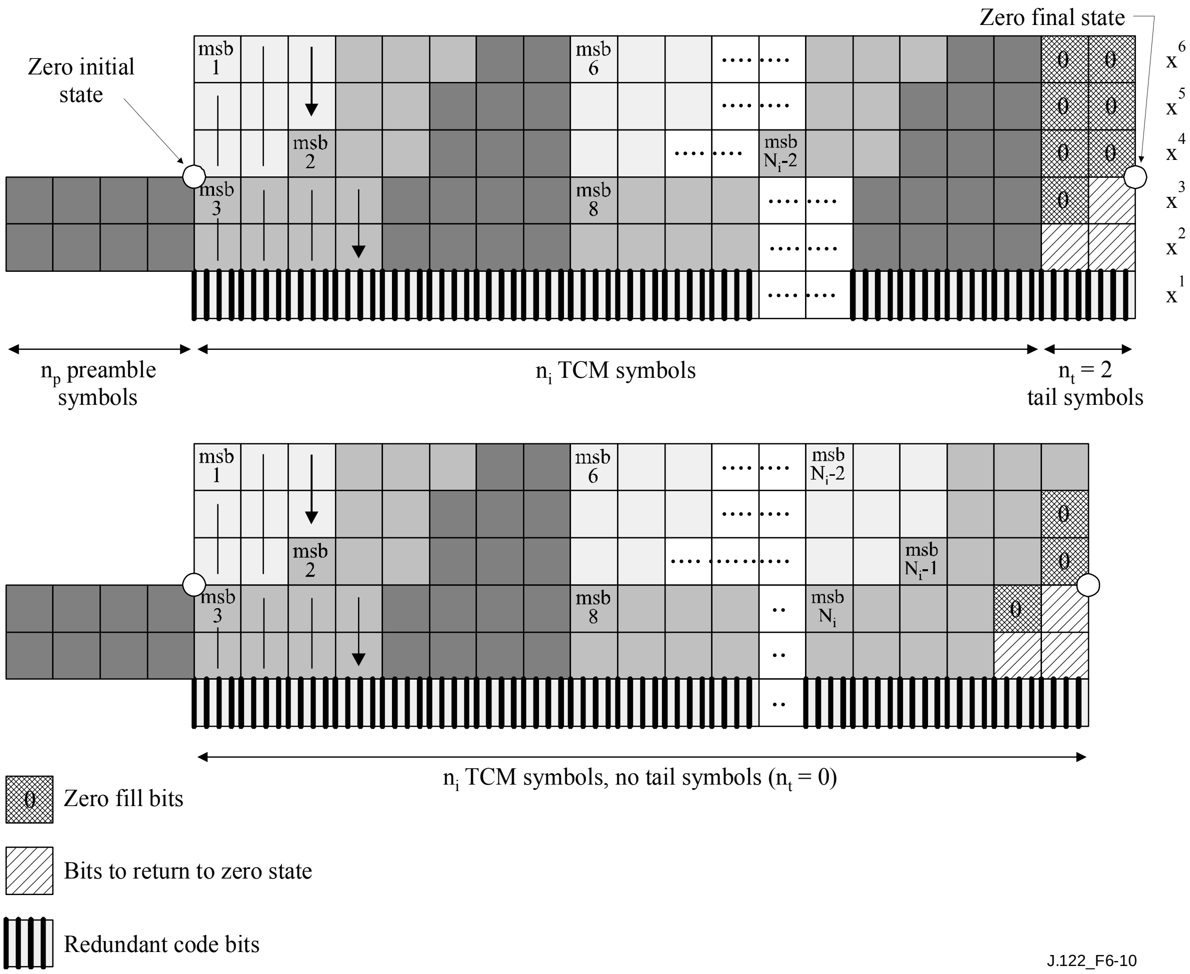

Figure 6-10 – Example byte to bit assignment for 64QAM

The CM MUST place the return-to-zero bits right after the last coded data sub-symbol, that is, the last coded sub-symbol corresponding to the parity bytes of the last shortened or fixed codeword including any zero-filled codeword in the grant. The rest of the coded bits MUST be filled with zeros.

Figure 6-11 illustrates the placing of return-to-zero bits for 64QAM 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.

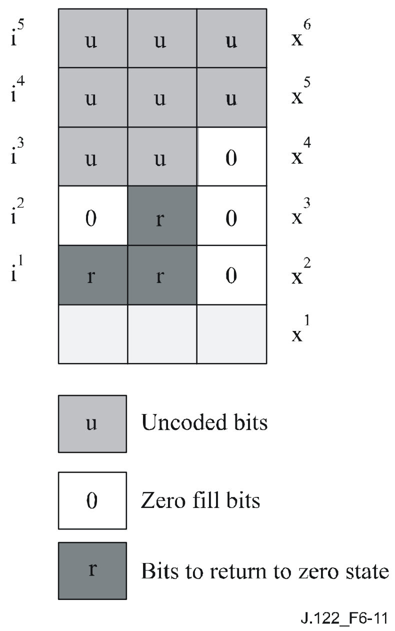

Figure 6-11 – Example of return-to-zero bits followed by "0"

6.2.9 Preamble prepend

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

The first bit of the preamble pattern is the first bit into the symbol mapper (see clause 6.2.13). The first bit of the preamble pattern is designated by the preamble value offset as described in Table 8-19. The preamble is interleaved by the framer in S-CDMA mode.

The preamble sequence MUST be programmable. For DOCS 2.0 bursts (bursts encoded using a type 5 burst descriptor per clause 8.3.3), the preamble MUST use the QPSK0 or QPSK1 constellation (per Figures 6-18 and 6-19) with preamble length 0, 2, 4, 6,..., or 1536 bits (maximum 768 QPSK symbols). For DOCS 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 DOCS 1.x compatible bursts (type 4 burst descriptor) that use 16QAM modulation, the preamble and data MUST use the 16QAM constellation with preamble length 0, 4, 8, 12,..., or 1024 bits (maximum 256 16QAM symbols).

The preamble length and value MUST be configured in response to the upstream channel descriptor message transmitted by the CMTS.

6.2.10 Modulation rates

In TDMA mode, the CM upstream modulator MUST provide all modulations at 160, 320, 640, 1280, 2560 and 5120 ksymb/s.

In S-CDMA mode, the CM upstream modulator MUST provide all modulations at 1280, 2560 and 5120 ksymb/s.

In TDMA mode, the CMTS upstream demodulator MUST be able to support demodulation at 160, 320, 640, 1280, 2560 and 5120 ksymb/s. In S-CDMA mode, the CMTS upstream demodulator MUST be able to support demodulation at 1280, 2560 and 5120 ksymb/s.

This variety of modulation rates, and flexibility in setting upstream carrier frequencies, permits operators to position carriers in gaps in the pattern of narrow-band ingress, as discussed in Annex G.

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 S-CDMA framer and interleaver

6.2.11.1 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 must be 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. 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 clause 6.2.14, "S-CDMA spreader".

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.

6.2.11.2 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: (these messages are described in clauses 8.3.2 and 8.3.3).

Mini-slots are mapped onto frames starting at the first active code (usually code number 0), are numbered sequentially through the remainder of the frame (code number 127), and then wrap to the next sequential frame. 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 MUST identify frame boundaries relative to the timestamp counter on a periodic basis. This is called the *timestamp snapshot* and must be sent in the UCD for each upstream S-CDMA channel.

The CMTS MUST maintain a frame counter and a mini-slot counter, and MUST sample these values along with the timestamp, on a frame boundary, as shown in Figure 6-12. The CMTS MUST obtain a new sample prior to sending each UCD message.

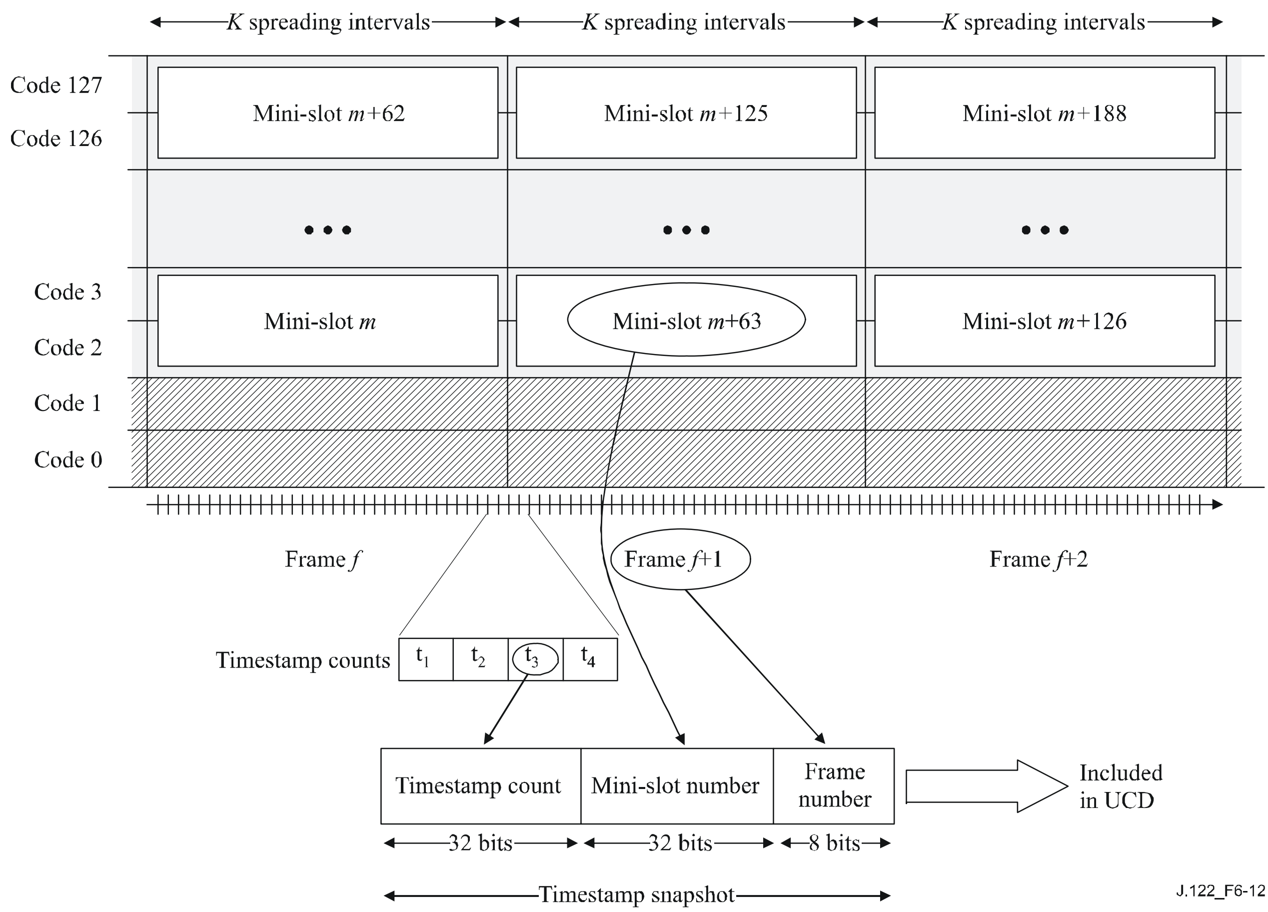

Figure 6-12 – Timestamp snapshot

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 clause 6.2.14.1.

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 clause 6.2.12.

6.2.11.2.1 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 signalling rate), $1/T_s$, define the time duration of an S-CDMA frame, T_{fr}.

$$T_{fr} = K \times 128 \times T_s$$

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:

$$S_{ms} = K \times C_{ms}$$

The lower limit on mini-slot capacity is 16 symbols (refer to Annex B). However, the mini-slot must also be large enough to allow the transmission of the largest-sized data PDU (including physical layer overhead) in 255 mini-slots (see clause 8.3.3). The upper limit on mini-slot capacity is not specifically constrained, but in general is governed by channel efficiency and MAC performance issues. The valid range of the codes per mini-slot parameter is 2 to 32.

Number of active codes: The number of active codes parameter allows the number of codes used to carry data to be less than or equal to 128. When the number of active codes is less than 128, low numbered codes starting with code 0 are not used, as shown in Figure 6-14.

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 coloured 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.

The valid range of the *number of active codes* parameter is 64 to 128.

A CMTS MUST support 126 and 128 active codes.

A CM MUST support any nonprime number of active codes in the range of 64 to 128, inclusive.

NOTE – When the number of active codes is 64 or greater, the S-CDMA frame must consist of more than one mini-slot, since the number of codes per mini-slot must be in the range 2 to 32. This implies that the number of active codes must be nonprime.

The prime numbers between 64 and 128 are 67, 71, 73, 79, 83, 89, 97, 101, 103, 107, 109, 113 and 127.

6.2.11.2.2 Mini-slot numbering examples

A typical mini-slot numbering example is shown in Figure 6-13. 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.

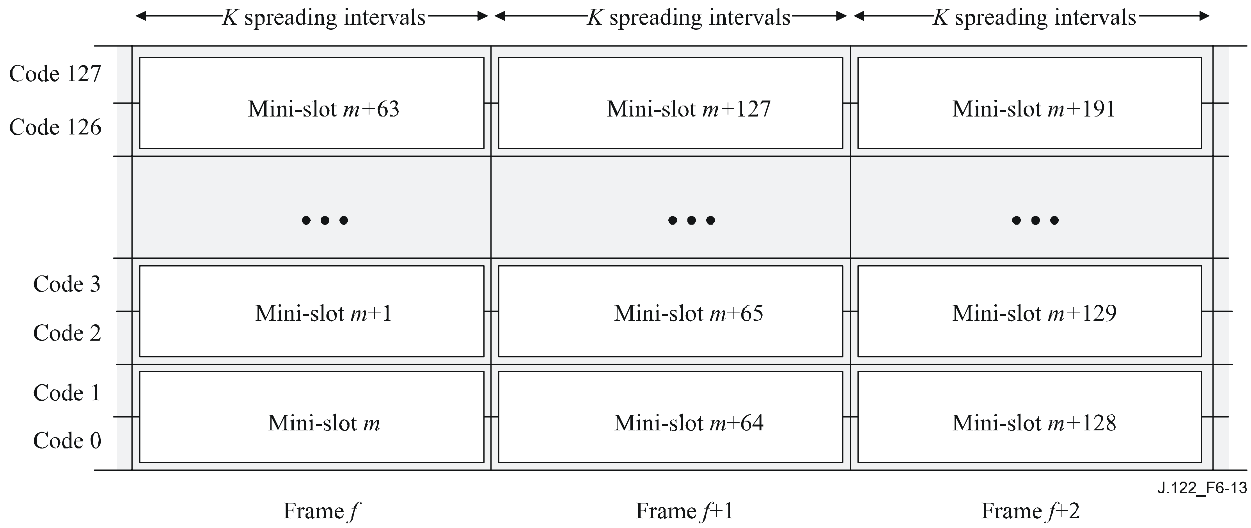

Figure 6-13 – Mini-slot mapping with two codes per mini-slot, 128 active codes

A second example, using three codes per mini-slot, is shown in Figure 6-14. Since it is required that there be an integral number of mini-slots per frame, the number of active codes has been restricted to 126 codes. In this example, a trade-off has been made to increase mapping flexibility at the expense of a small reduction in channel capacity (2/128).

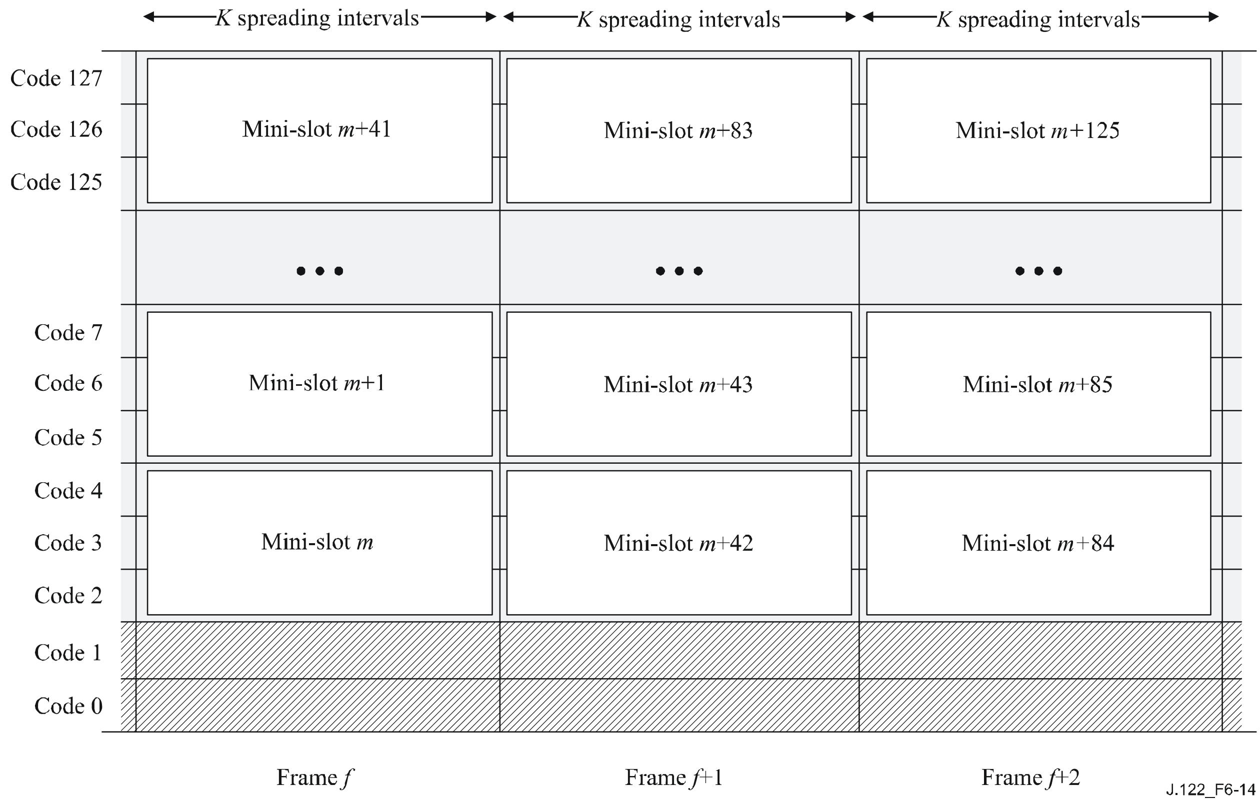

Figure 6-14 – Mini-slot mapping with three codes per mini-slot, 126 active codes

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.

6.2.11.3 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.

6.2.11.4 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 signalling rates and the fact that there are 128 signalling periods in a spreading interval, the S-CDMA frame time MUST always be a multiple of 25 µs.

Selecting the number of spreading intervals per frame and the signalling rate therefore exactly define the S-CDMA frame duration. As a specific example, a burst profile defined with 10 spreading intervals per frame with a signalling rate of 2.56 Mbaud would result in a frame duration of 500 µs.

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 clause 6.2.17.

6.2.11.5 Spreader-off bursts for maintenance on S-CDMA channel

Spreader-off bursts are defined as bursts on an S-CDMA channel whose attributes specify that 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 $C_{ms} \times K$ symbols.

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 $C_{ms} \times K$ 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 clause 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 clause 6.2.5.1.1. Differential encoding and R-S byte interleaving MUST NOT be used with spreader-off bursts on S-CDMA channels.

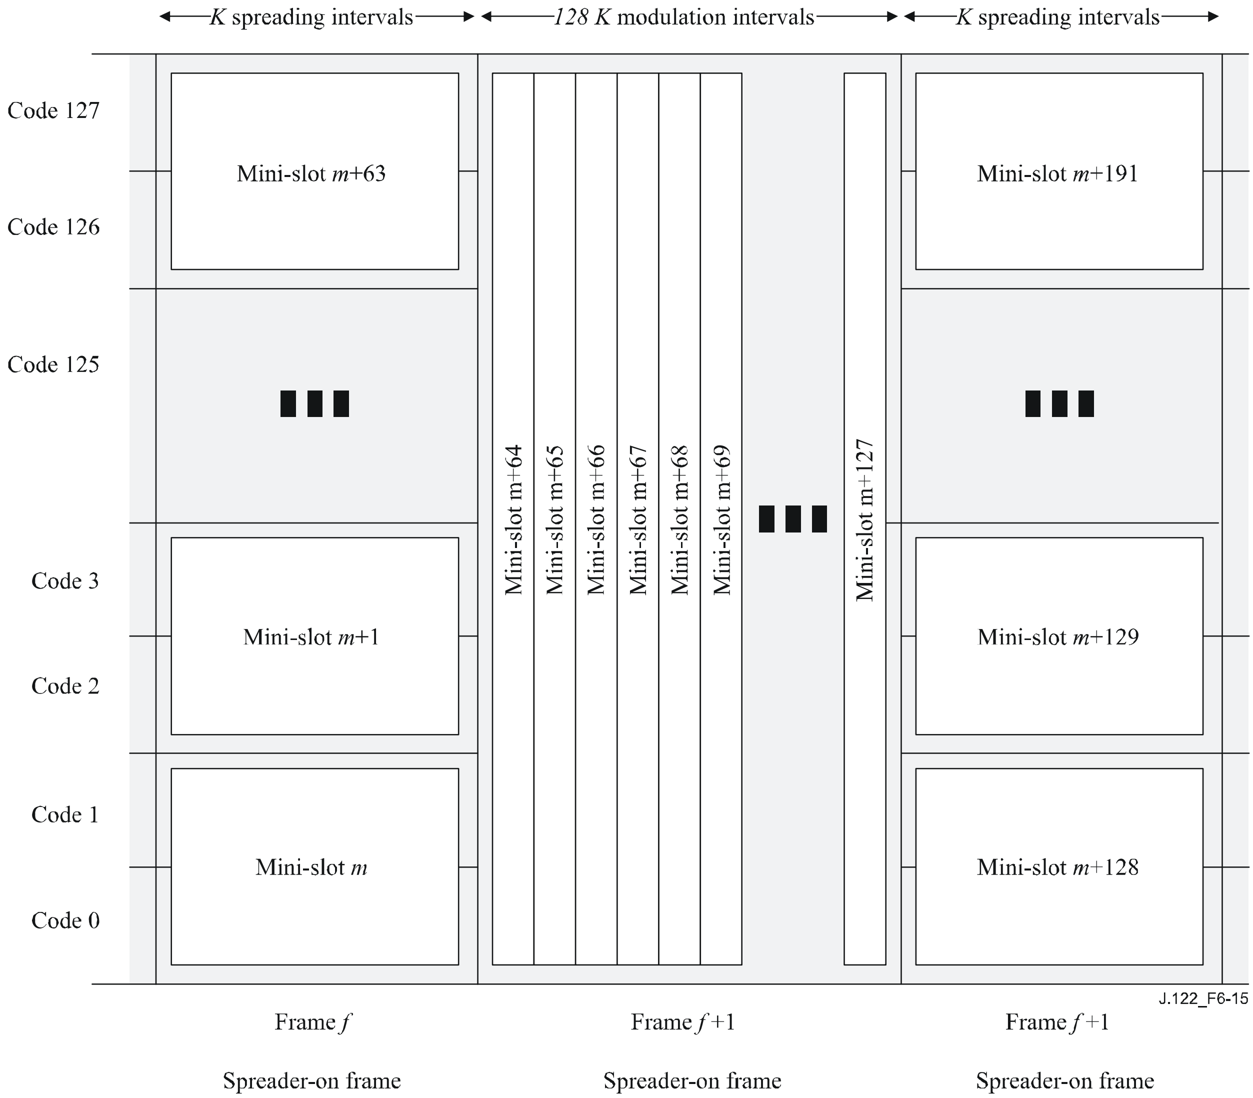

Figure 6-15 – S-CDMA and spreader-off intervals

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, and 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.

6.2.11.6 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 [DOCS5]. 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.

6.2.12 S-CDMA framer

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 clause 6.2.11. 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 clause 6.2.13.

6.2.12.1 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 codewords to enhance protection from impulse noise.

Given an S-CDMA frame which is N active codes 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 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, and 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′ ≤ R. Figure 6-16 shows a subframe consisting of R rows and K spreading intervals within an S-CDMA frame.

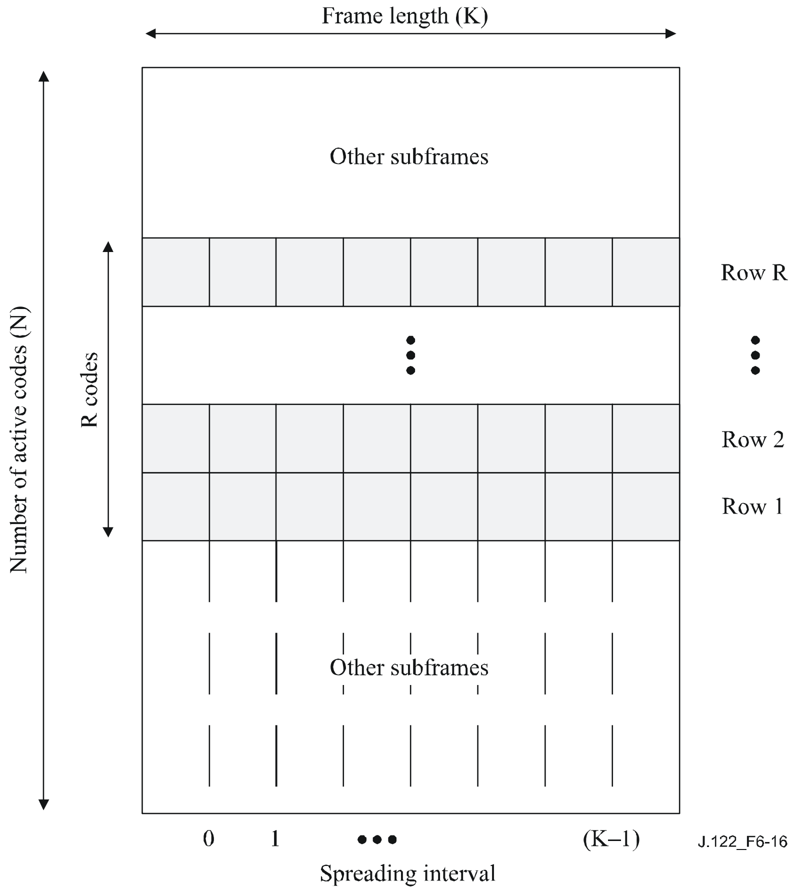

Figure 6-16 – Subframe structure

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 1 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.

6.2.12.2 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 clauses.

6.2.12.2.1 Rules for preamble and coded TCM symbols

The 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 6-16 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 + 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 + 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 clause 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.

6.2.12.2.2 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 clause 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.

6.2.12.2.3 Subframe example

Figure 6-17 below shows an example which follows the above specified rules. Each box in the figure represents a symbol that 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.

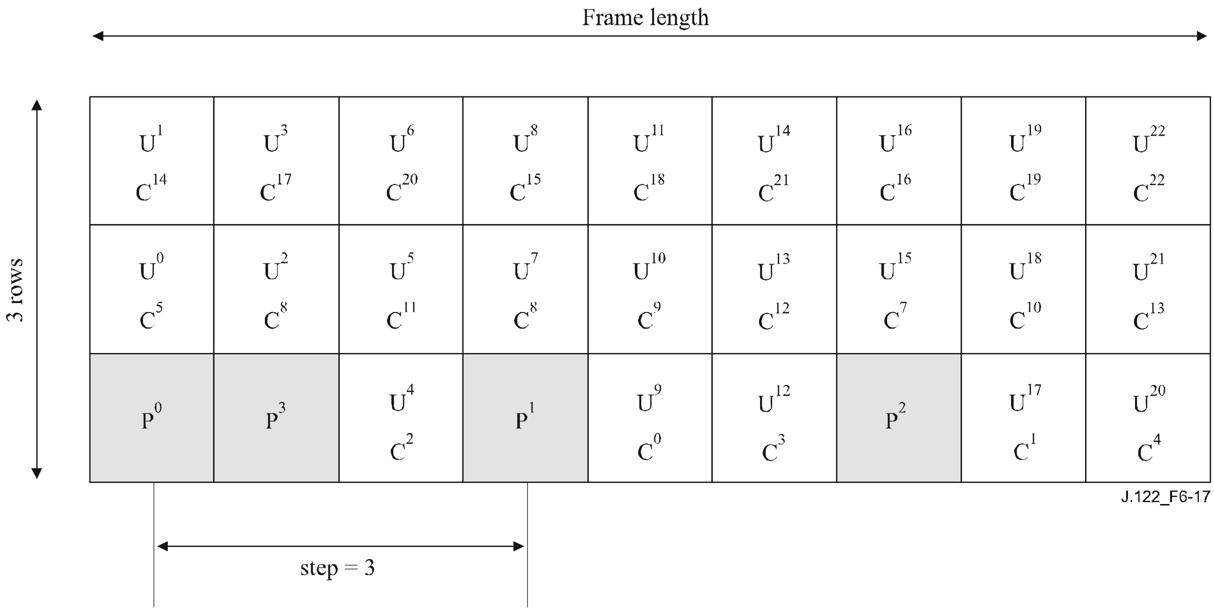

Figure 6-17 – Symbol numbering without TCM

6.2.12.2.4 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 described in Figure 6-16, 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.

6.2.13 Symbol mapping

The modulation mode is configurable via MAC messages. Differential encoded QPSK and 16QAM are available for TDMA channels. QPSK, 8QAM, 16QAM, 32QAM and 64QAM are available for TDMA and S-CDMA channels. TCM-encoded QPSK, 8QAM, 16QAM, 32QAM, 64QAM and 128QAM 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, 8QAM, 16QAM, 32QAM, 64QAM and 128QAM, respectively. The MSB MUST be the first bit in the serial data into the symbol mapper and it 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.

Table 6-3 – I/Q mapping

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

All constellations are defined on a common integer grid in Figure 6-18. 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 6-18, 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 64QAM, G_{const}, are given. The QPSK0 constellation is employed for low-power preamble and QPSK data symbols. Use of QPSK1 is restricted to high-power preamble symbols.

The upstream symbol constellations MUST be as shown in Figure 6-18.

The upstream QPSK Gray-coded and differential symbol mapping MUST be as shown in Figure 6-19.

The upstream 8QAM symbol mapping MUST be as shown in Figure 6-20.

The upstream 16QAM Gray-coded symbol mapping MUST be as shown in Figure 6-21.

The upstream 16QAM differential symbol mapping MUST be as shown in Figure 6-21.

The upstream 32QAM symbol mapping MUST be as shown in Figure 6-22.

The upstream 64QAM Gray-coded symbol mapping MUST be as shown in Figure 6-23.

The TCM symbol mappings used for S-CDMA are shown in Figure 6-24 through Figure 6-26.

The upstream QPSK TCM symbol mapping MUST be as shown in Figure 6-24.

The upstream 8QAM TCM symbol mapping MUST be as shown in Figure 6-24.

The upstream 16QAM TCM symbol mapping MUST be as shown in Figure 6-25.

The upstream 32QAM TCM symbol mapping MUST be as shown in Figure 6-25.

The upstream 64QAM TCM symbol mapping MUST be as shown in Figure 6-26.

The upstream 128QAM TCM symbol mapping MUST be as shown in Figure 6-26.

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 16QAM 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 16QAM 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

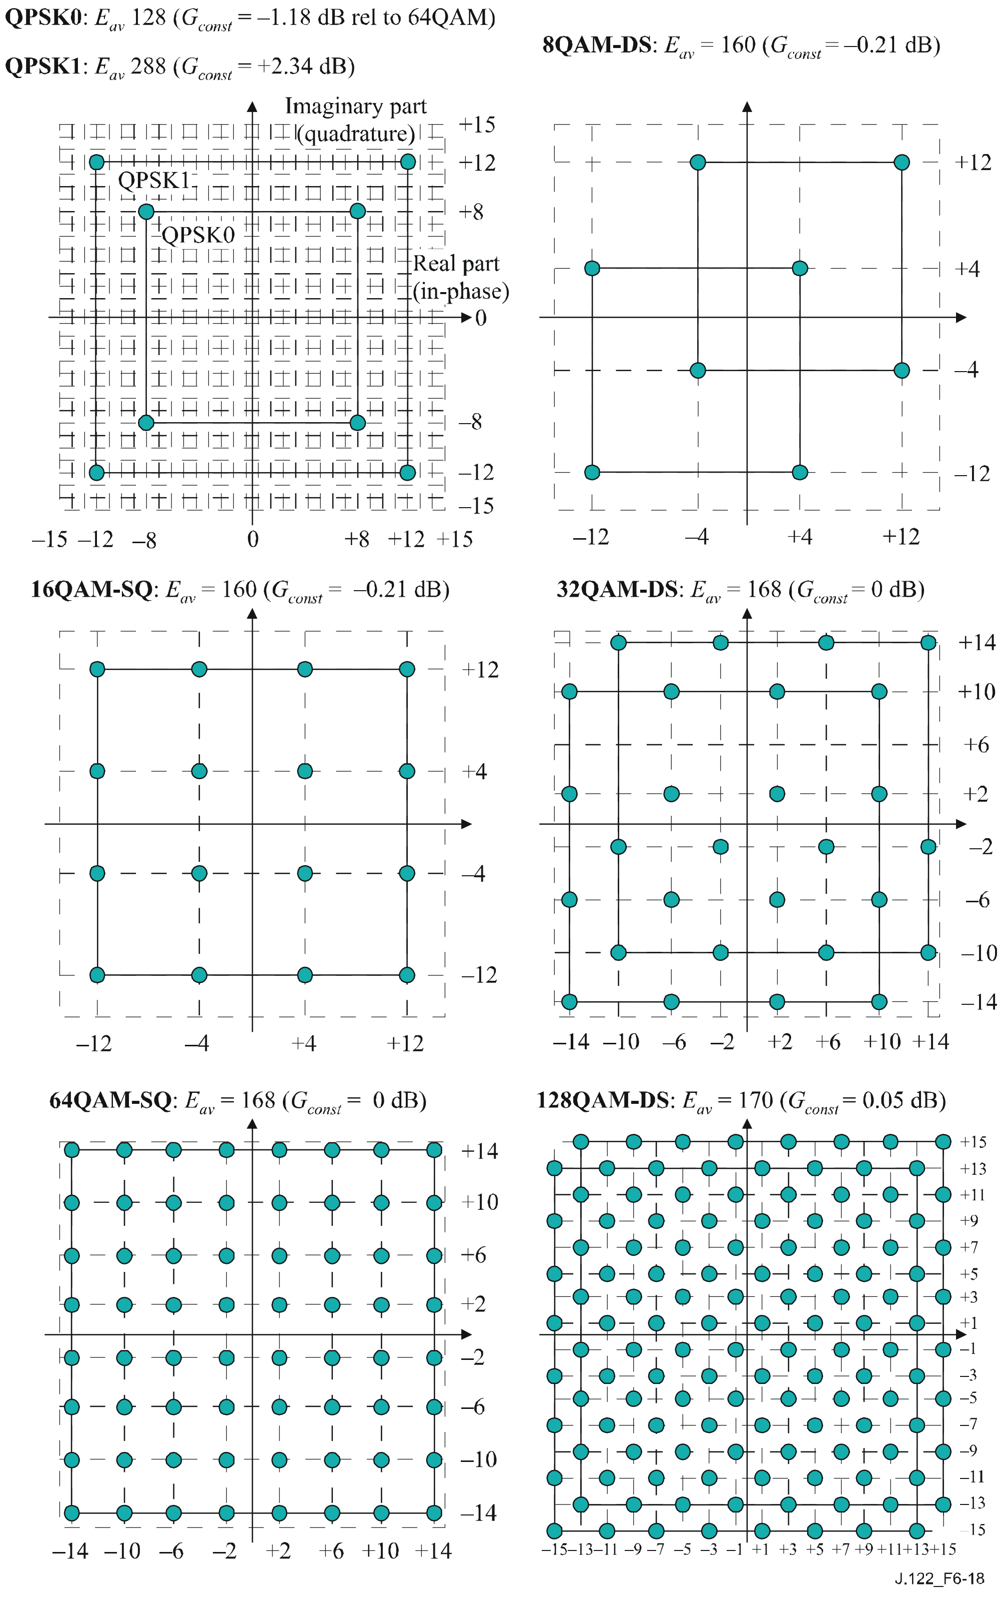

Figure 6-18 – Symbol constellations

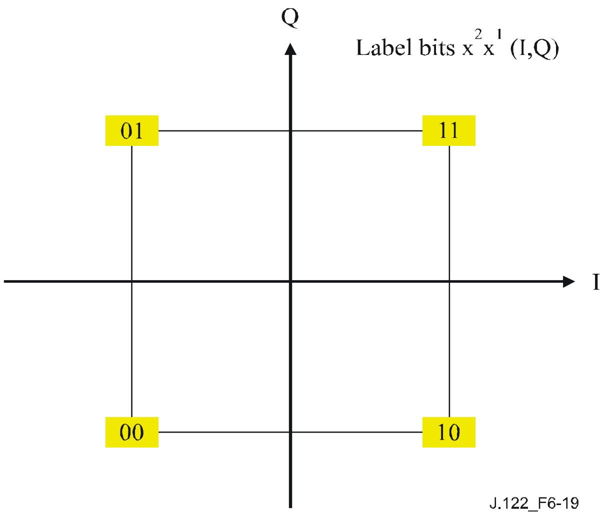

Figure 6-19 – QPSK Gray-coded and differential symbol mapping

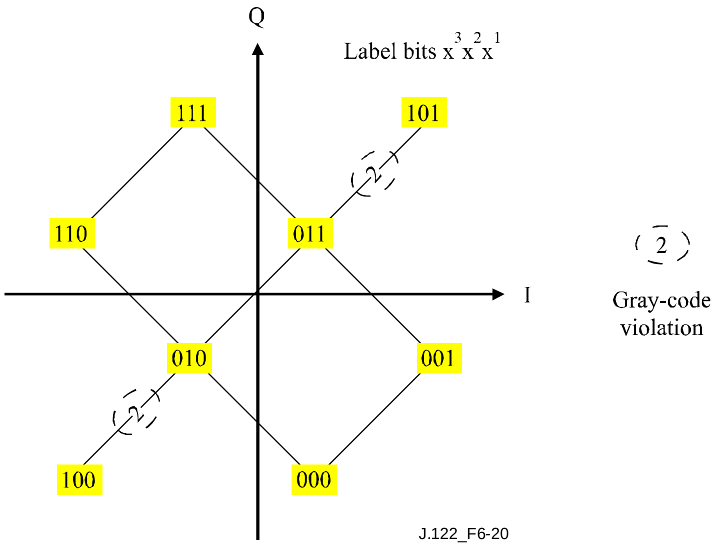

Figure 6-20 – 8QAM symbol mapping

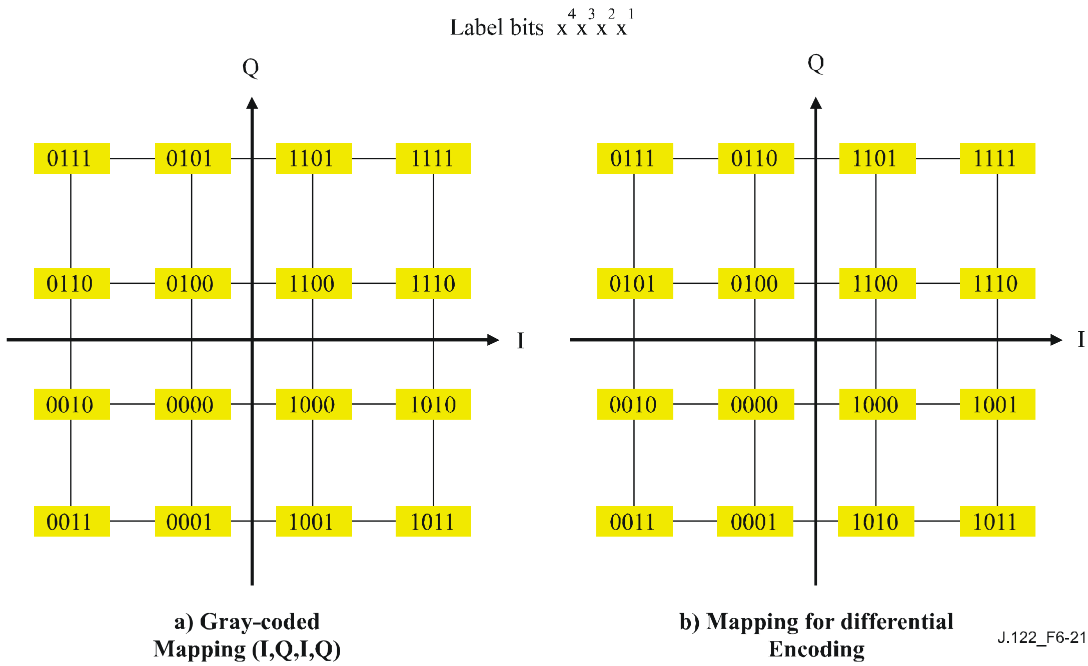

Figure 6-21 – 16QAM symbol mapping

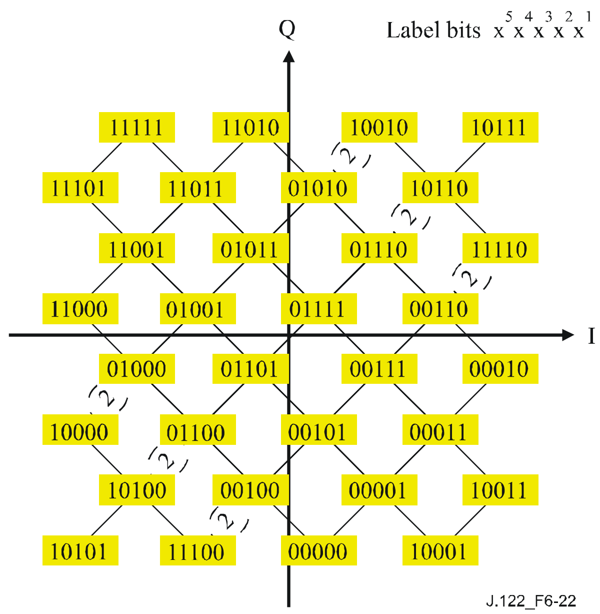

Figure 6-22 – 32QAM symbol mapping

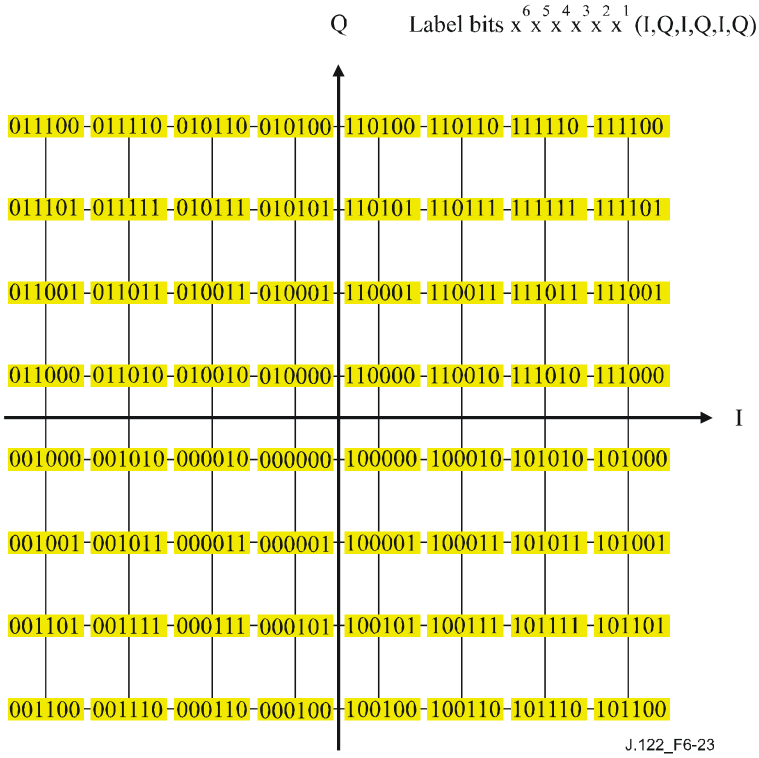

Figure 6-23 – 64QAM symbol mapping

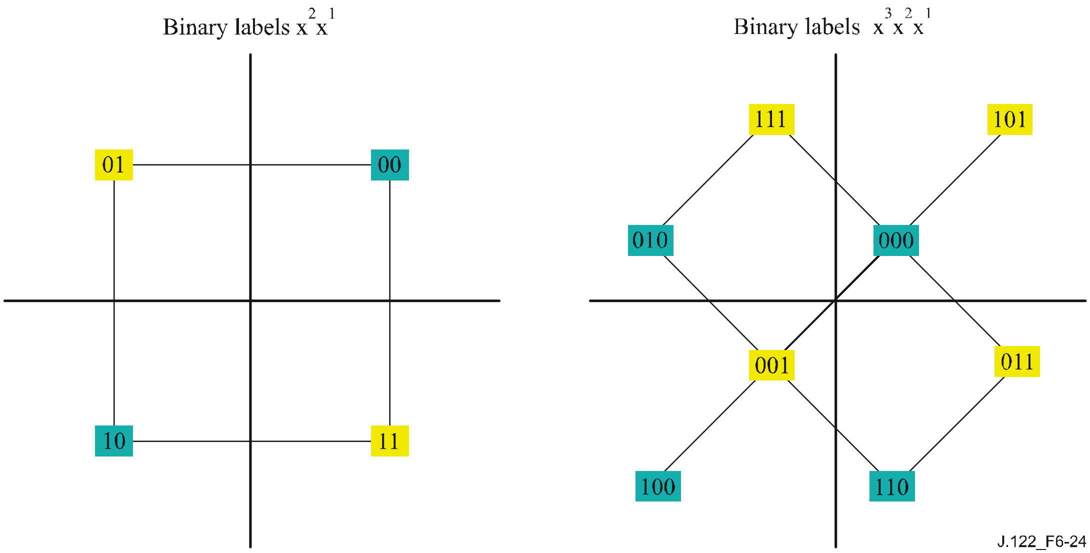

Figure 6-24 – QPSK and 8QAM TCM symbol mapping

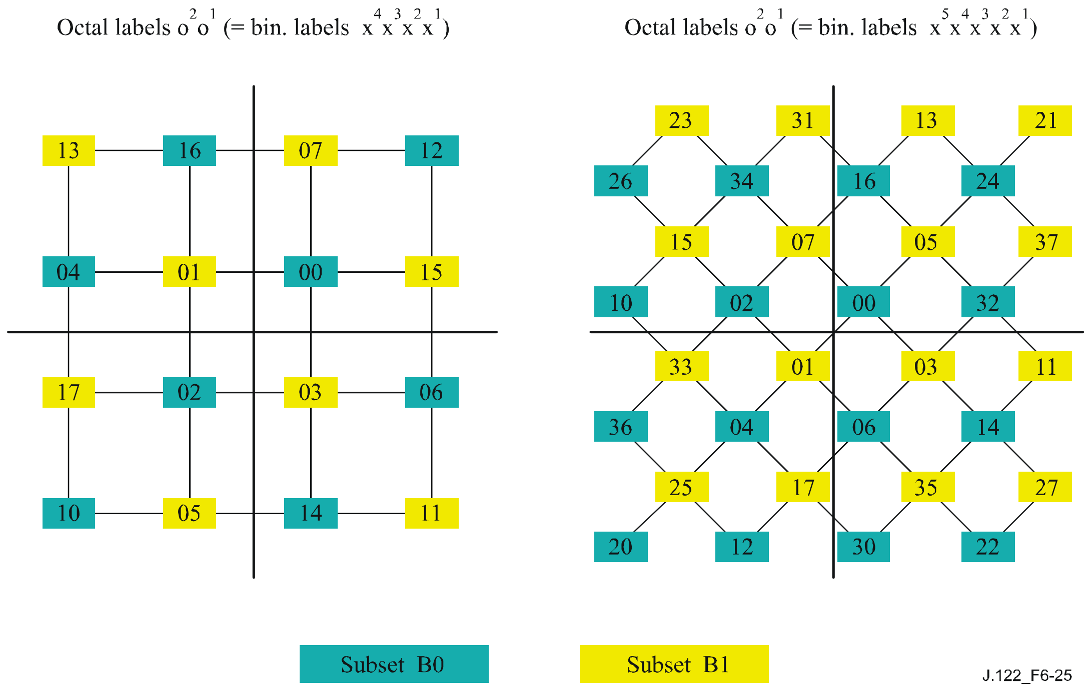

Figure 6-25 – 16QAM and 32QAM TCM symbol mapping

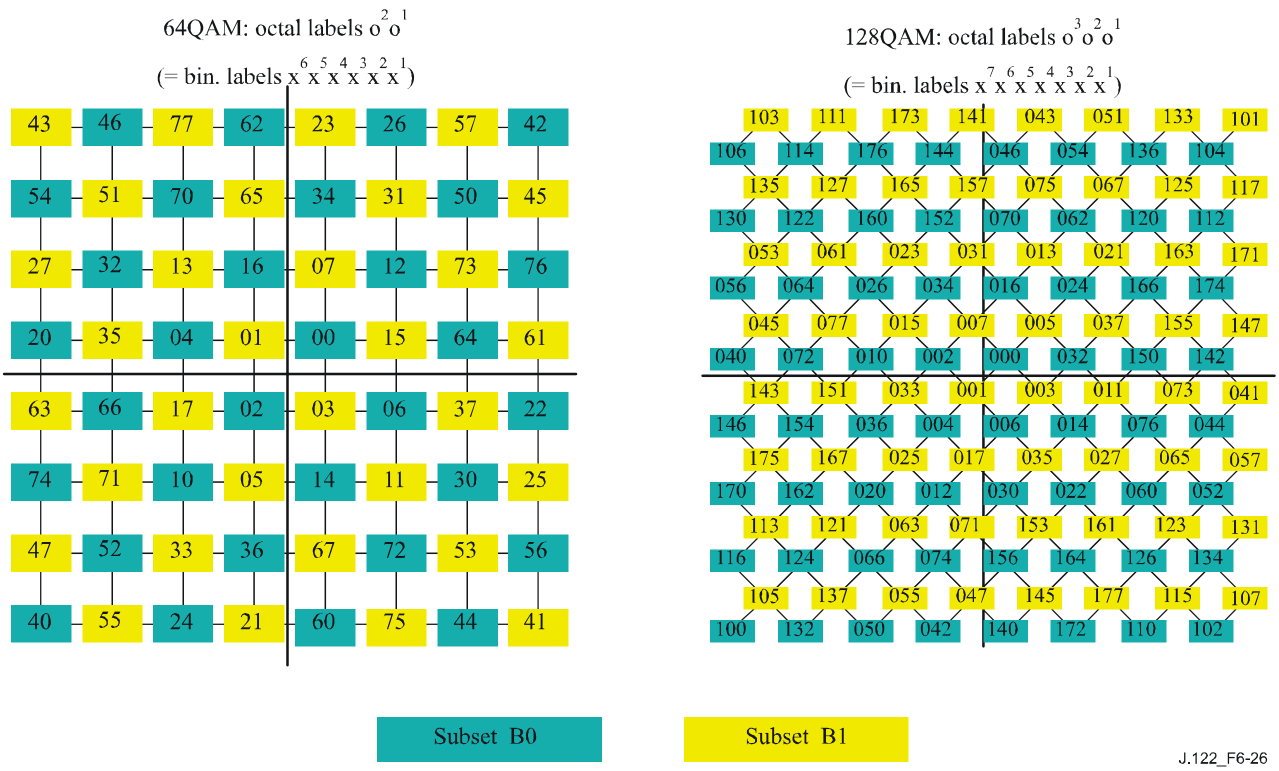

Figure 6-26 – 64QAM and 128QAM TCM symbol mapping

6.2.14 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, $\overline{P}_k$, is transmitted such that:

$$\overline{P}_k = \overline{S}_k \times C$$

where $\overline{S}_k$ is a vector, $[s_{k,127}, s_{k,126}, ..., s_{k,0}]$, of modulation symbols on the integer grid of clause 6.2.13 to be transmitted in spreading interval *k,* and *C* is a matrix:

$$C = \begin{bmatrix} c_{127,127} & c_{127,126} & \cdots c_{127,0} \\ c_{126,127} & c_{126,126} & \cdots c_{126,0} \\ \ldots & \ldots & \ldots \quad \ldots \\ c_{0,127} & c_{0,126} & \cdots c_{0,0} \end{bmatrix} = \begin{bmatrix} x_1 & x_{127} & \ldots & x_2 & -1 \\ x_2 & x_1 & \ldots & x_3 & -1 \\ \ldots & \ldots & \ldots & \ldots & -1 \\ x_{127} & x_{126} & \ldots & x_1 & -1 \\ 1 & 1 & 1 & 1 & 1 \end{bmatrix}$$

where the rows of *C* are the 128 spreading codes such that $Code(j) = [c_{j,127}, c_{j,126}, ..., c_{j,0}]$. The result of the spreading operation is the transmission vector $\overline{P}_k$, which has 128 elements, $[P_{k,127}, P_{k,126}, ... P_{k,0}]$, where each element is transmitted at the signalling 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 that are either +1 or −1. *Code*(0) consists of 128 elements all of which have a value of +1. For each of the other spreading codes *Code*(*j*)*,* the element $c_{j,0}$ 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 clause.

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*(*j*) is obtained by cyclic shift to the left (in the direction of increasing indices) of *Code*(*j*−1) where the element, $c_{j,0}$, 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 clause 6.2.13.

If a CM has not been assigned to use a particular code, j, at a spreading time interval, *k,* then in its computation of its transmission vector $\overline{P}_k$, it will set $s_{k,j}$ 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 clause 6.2.12.

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

It is also important to note that in the matrix multiplication 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 $\overline{P}_k$ 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 6-8).

6.2.14.1 Code hopping

Code hopping refers to a systematic re-ordering of the rows of the spreading matrix, C, at each spreading interval, *k*. The code-hopping algorithm uses a pseudo-random number, *lfsr_out(k)*, to determine a cyclic shift of the rows of the matrix C. When the number of active codes equals 128, the code-hopping algorithm uses all codes. When the number of active codes is less than 128, the code-hopping algorithm hops only over the cyclic codes (*Code*(0), the all 1s code, is excluded). The generalization of the spreading matrix at spreading interval, k, is given by (where matrix elements, $c_{j,i}$, are as defined above in clause 6.2.14):

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

where:

$$f(k,i) = \begin{cases} \text{modulo}[(128 - lfsr\_out(k) + i),128] & \text{with } 0 \le i \le 127 \text{ for 128 active codes} \\ \text{modulo}[(126 - lfsr\_out(k) + i),127] + 1 & \text{with } 1 \le i \le 127 \text{ for} < 128 \text{ active codes} \end{cases}$$

In S-CDMA mode, the CM MUST support code hopping.

Note that when the number of active codes is less than 128, the unused codes are those starting with matrix index 0. In this case, the code hopping continues to "hop" over all of the codes except for *Code*(*0*), even if the number of active codes is less than 127.

The pseudo-random number generator which determines the spreading matrix reordering is the linear-feedback shift register (LFSR) which is shown in Figure 6-27. In order to align the CM's code-hopping pseudo-random sequence with that at the CMTS, the pseudo random generator must output the following value at the first spreading interval of each frame:

lfsr_out(frame_number * spreading_interval_per_frame)

where lfsr_out(k) is the value of lfsr_out after k shifts following the code hopping seed load into the LFSR.32 (the description of the frame counter and the procedures for its synchronization are contained in clause 6.2.11.2, "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. Then at each subsequent spreading interval, k, a new bit is shifted into the LFSR producing a new 7-bit value, *lfsr_out*(*k*). This value is used, with bit 7 as the MSB, to compute the spreading matrix indices as given by the equation above. Note that the code-hopping mechanism (LFSR, spreading interval index) is advanced every spreading interval (128 modulation intervals) in both spreader-on and spreader-off frames.

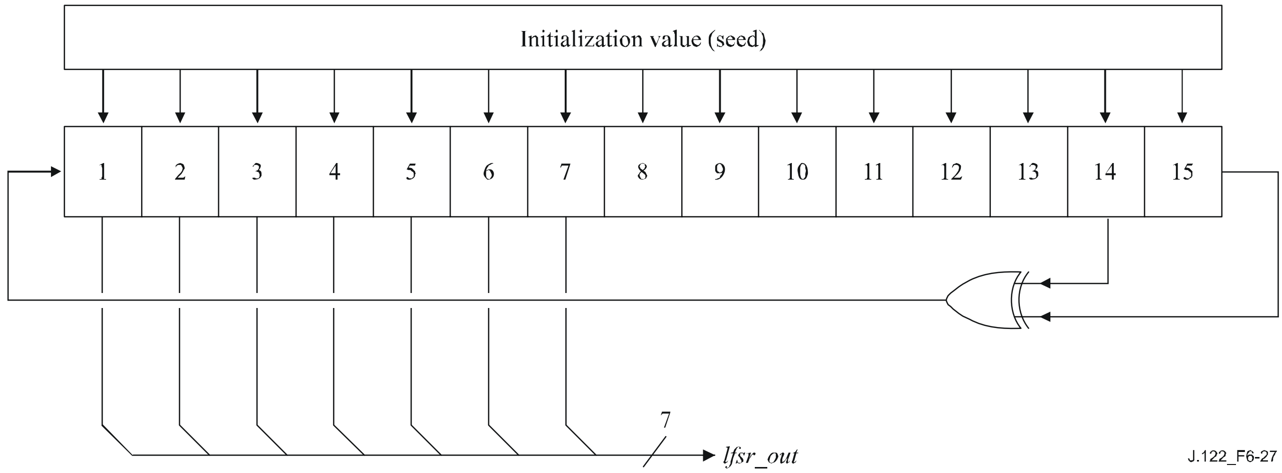

Figure 6-27 – Code-hopping random number generator

The 15-bit seed value is configured in response to the upstream channel descriptor message from the CMTS.

6.2.15 Transmit pre-equalizer

A transmit pre-equalizer of a linear equalizer structure, as shown in Figure 6-28, MUST be configured by the CM in response to the ranging response (RNG-RSP) message transmitted by the CMTS.

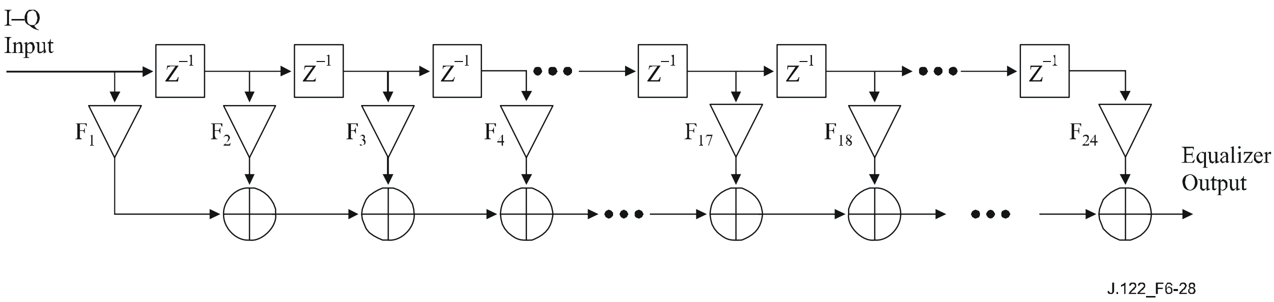

Figure 6-28 – Transmit pre-equalizer structure

There are two modes of operation for the pre-equalizer of a CM: DOCS 1.1 mode, and DOCS 2.0 mode. In DOCS 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 DOCS 1.1 mode, for backwards compatibility, the CMTS MAY support fractionally spaced equalizer format (T/2 and T/4). In DOCS 2.0 mode, the pre-equalizer MUST support a symbol (T)-spaced equalizer structure with 24 taps.

In DOCS 1.x-only logical channels, the CM and the CMTS MUST use DOCS 1.1 mode.

In DOCS 2.0-only logical channels, the CM and the CMTS MUST use DOCS 2.0 mode.

In DOCS 1.x/2.0 mixed logical channels, the CM and the CMTS MUST use DOCS 1.1 mode from initial ranging until DOCS 2.0 is activated in the registration process (if it is activated), and MUST use DOCS 2.0 mode after DOCS 2.0 is activated.

The RNG-RSP MAC message carries the CM transmit equalization information, and may instruct the CM to either convolve the equalizer coefficients or (in DOCS 2.0 mode only) load them directly (refer to clause 8.3.6.1, "Encodings"). 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 DOCS 1.1 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:

$$F_n^{m+1} = \sum_{k=\max(1-L^{m+1}, n+L^m-L^{m+1}-24)}^{\min(24-L^m-L^{m+1}-1)} F^m_{n-k+L^m-L^{m+1}} \cdot F_{k+L^{m+1},}\ n=1\ldots24$$

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

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 DOCS 1.x-only 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. See clause 8.3.8.1.1 for proper use of the modem capabilities field.

In DOCS 2.0-only 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 DOCS 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 DOCS 1.1 mode is activated, 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 DOCS 2.0 is activated, the CMTS MUST use a T-spaced equalizer structure with 24 taps. If the first update of the pre-equalizer after the activation of DOCS 2.0 uses "convolve" mode, the CM MUST zero-pad the existing 8-tap filter to a 24-tap filter, and then 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 DOCS 1.1 mode, the CMTS MUST NOT move the main tap location during periodic ranging.

In DOCS 1.1 mode, the CMTS MUST NOT instruct the CM to load the transmit equalizer coefficients.

In DOCS 2.0 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 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 clause 6.2.18.

In DOCS 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. Again, in doing so, 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.

6.2.16 Spectral shaping

The upstream transmitter MUST approximate a Nyquist square-root raised-cosine pulse-shaping filter with roll-off factor $\alpha = 0.25$. The –30 dB transmitted bandwidth MUST NOT exceed the channel width values in Table 6-5. The channel width values are given analytically by:

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

Table 6-5 – Maximum channel width

Modulation rate (kHz)	Channel width (kHz)
160	200
320	400
640	800
1280	1600
2560	3200
5120	6400

6.2.16.1 Upstream frequency agility and range

The upstream PMD sublayer MUST support operation over the frequency range of 5-42 MHz edge to edge.

Offset frequency commands MUST be supported per Table 6-8.

6.2.16.2 Spectrum format

The upstream modulator MUST provide operation with the format

$$s(t) = I(t) \times \cos(\omega t) - Q(t) \times \sin(\omega t)$$

where t denotes time and ω denotes angular frequency.

6.2.17 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 specified below.

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

$$D_p = 200 + \frac{M}{5.12} \mu s \ ,$$

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

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 DOCS 1.x mode, $M = 0$. Note that in the above equations, the values for B_r and $I_r \times N_r$ 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 1025 µs, assuming a modulation rate of 5.12 MHz.

NOTE 1 – The CM MAP processing delay does not include downstream FEC de-interleaving delay.

NOTE 2 – 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.

6.2.18 Transmit power requirements

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

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

The protocol by which power adjustments are performed is defined in clause 11.2.4. Such adjustments by the CM MUST be within the ranges of tolerances described below. A CM MUST confirm that the transmit power limits are met after a RNG-RSP is received or after a UCD change.

Transmit power is defined as the average RF power in the occupied bandwidth (channel width) transmitted in the data symbols of a burst, assuming equally likely QAM symbols, measured at the F-connector of the CM. Maximum and minimum transmit power level requirements refer to the CM's target transmit power level, defined as the CM's estimate of its actual transmit power. The actual transmitted power MUST be within ±2 dB of the target power. The target transmit power MUST be variable over the range specified in Table 6-8.

Transmit power as reported by the CM in the MIB is referenced to the 64QAM constellation. When transmitting with other constellations, a slightly different transmit power will result, depending on the constellation gain in Table 6-6 (see clause 6.2.13). As an example, if the reported power is 30 dBmV, 64QAM will be transmitted with a target power of 30 dBmV, while QPSK will be transmitted with 28.82 dBmV.

Table 6-6 – Constellation gains and power limits

Constellation	Constellation gain G_{const} relative to 64QAM (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
QPSK	–1.18	8	58	53	9.18	59.18	54.18
8QAM	–0.21	8	55	53	8.21	55.21	53.21
16QAM	–0.21	8	55	53	8.21	55.21	53.21
32QAM	0.00	8	54	53	8.00	54.00	53.00
64QAM	0.00	8	54	53	8.00	54.00	53.00
128QAM	0.05	8	N/A	53	7.95	N/A	52.95

The actual transmitted power within a burst MUST be constant to within 0.1 dB peak to peak. 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 clauses 6.2.18.1 and 6.2.18.2.

6.2.18.1 TDMA transmit power calculations

In TDMA mode, the CM determines its target transmit power P_t as follows. Define:

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

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

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

P_{min} = Minimum target transmit power permitted for the CM per clause 6.2.21.1 (see Table 6-6)

P_{max} = Maximum target transmit power permitted for the CM per clause 6.2.21.1 (see Table 6-6)

P_{hi} = $\min[P_{max} – G_{const}]$ over all burst profiles used by the CM (see Table 6-6)

P_{low} = $\max[P_{min} – G_{const}]$ over all burst profiles used by the CM (see Table 6-6)

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

The CM updates its reported power by the following steps:

1) $P_r = P_r + \Delta P$ //Add power level adjustment to reported power level

2) $P_r = \min[P_r, P_{hi}]$ //Clip at max. power limit

3) $P_r = \max [P_r, P_{low}]$ //Clip at min. power limit

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

Usually, the reported power level is a relatively constant quantity, while the transmitted power level varies dynamically as different burst profiles, with different constellation gains, are transmitted. A CM's target transmit power MUST never be below P_{min} or above P_{max}. This implies that in some cases the extreme transmit power levels (e.g., 58 dBmV for QPSK and 8 dBmV) may not be permitted if burst profiles with multiple constellations are active. Also, if only QPSK is used, the reported power may be greater than 58 dBmV, although the target transmit power will not exceed 58 dBmV.

For example, if only QPSK and 64QAM burst profiles are active, P_{hi} = 54 dBmV and P_{low} = 9.2 dBmV. The maximum permitted QPSK transmitted power is 54 – 1.2 = 52.8 dBmV, the minimum QPSK power is 9.2 – 1.2 = 8 dBmV, the maximum 64QAM power is 54 dBmV, and the minimum 64QAM power is 9.2 dBmV.

6.2.18.2 S-CDMA transmit power calculations

In S-CDMA mode, the power calculations depend on whether the maximum scheduled codes feature is enabled.

6.2.18.2.1 S-CDMA transmit power calculations when maximum scheduled codes are not enabled

In S-CDMA mode, when maximum scheduled codes are not enabled, the CM determines its target transmit power P_t as follows. Define:

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

P_{hi} = $\min[P_{max} - G_{const}]$ over all burst profiles used by the CM (see Table 6-6)

P_{low} = $\max[P_{min} - G_{const}]$ + 10 log (*number_active_codes/number_of_codes_ per_mini-slot)* where the maximum is over all burst profiles used by the CM (see Table 6-6)

The CM updates its reported power by the following steps:

1)	P_r =	$P_r + \Delta P$	//Add power level adjustment to reported power level
2)	P_r =	$\min[P_r, P_{hi}]$	//Clip at max. power limit
3)	P_r =	$\max[P_r, P_{low}]$	//Clip at min. power limit

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 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, P_t, in a frame is the sum of the individual transmit powers, $P_{t,i}$, of each code, where the sum is performed using absolute power quantities (non-dB domain).

In a spreader-off frame, the CM target transmit power is:

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

The transmitted power level 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 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 (e.g., 8 and 53 dBmV) may not be permitted. Also if, for example, only QPSK is used, the reported power may be greater than 53 dBmV, although the target transmit power will not exceed 53 dBmV.

If, for example, QPSK and 64QAM burst profiles are active, the number of active codes is 128 and the number of codes per mini-slot is 2, then P_{hi} = 53 dBmV and P_{low} = 27.24 dBmV. The maximum permitted QPSK transmitted power is 53 – 1.18 = 51.82 dBmV when all active codes are transmitted. The minimum QPSK power is 27.24 – 1.18 –10log(128) + 10log(2) = 8 dBmV 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 64QAM power is 53 dBmV when all active codes are transmitted, and the minimum 64QAM power is 27.24 –10log(128) + 10log(2) = 9.18 dBmV when one mini-slot is transmitted. The minimum QPSK power permitted while transmitting, for example, 2 mini-slots is 11 dBmV, and the minimum 64QAM power permitted while transmitting 2 mini-slots is 12.2 dBmV.

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.

6.2.18.2.2 S-CDMA transmit power calculations when maximum scheduled codes are enabled

In S-CDMA mode on channels on which maximum scheduled codes are enabled, the CM determines its target transmit power P_t as follows. Define:

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

P_{hi_S} = min[53 – G_{const}] over all spreader-on burst profiles used by the CM (see Table 6-6)

P_{low_S} = max[8 – G_{const}] + 10 log(*number_active_codes/number_of_codes_per_mini-slot*) where the maximum is over all burst profiles used by the CM (see Table 6-6)

P_{max_T} = Maximum target transmit power permitted for the CM in TDMA mode (see Table 6-6) 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 (see Table 6-6)

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

P_{sf} = CM power shortfall per clauses 8.3.5 and 8.3.26

P_{hr} = S-CDMA power headroom in dB. Equivalent to TLV-11 defined in Table 8-21 divided by 4

ΔP = power level adjustment in dB sent from CMTS to CM

The CM updates its power by the following steps:

1. $P_r = P_r + \Delta P$ //Add power level adjustment to reported power level
2. $P_r = \min[P_r, P_{hi\_T}]$ //Clip at max. TDMA power limit
3. $P_r = \max[P_r, P_{low\_S}]$ //Clip at min. S-CDMA power limit
4. $P_{on} = \min[P_r, P_{hi\_S}]$ //Clip at max. S-CDMA power limit

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

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

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

$$P_{sf} = P_r + \max[G_{const,i}] - 53$$

// Difference between spreader-off reported and max. spreader-on target powers

$$P_{sf} = \max[P_{sf}, 0]$$

// Set P_{sf} to 0 if P_t is less than 53 dBmV

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

$$P_{t,\,i} = P_{on} + G_{const,\,i} - 10\log(number\_active\_codes) + P_{hr}$$

i.e., the clipped reported power 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 P_t in a frame is the sum of the individual transmit powers, $P_{t,\ i}$, of each code, 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, the burst profile contains QPSK for IUCs 1,2,3 and 4 and 64QAM 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} = 27.24 dBmV, and P_{hi_T} = 58 dBmV. Assume the CM ranges at spreader-off target transmit power of 57 dBmV. The CM reports P_{sf} = 57 dBmV – 53 dBmV = 4 dB. The CMTS uses P_{sf} to set (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 to:

$$\begin{aligned} P_{t,i} &= P_{on} + G_{const,\ i} - 10 \log(\textit{number\_active\_codes}) + P_{hr} \\ &= 53 \text{ dBmV} + 0 \text{ dB} - 21 \text{ dB} + 6 \text{ dB} \qquad \text{// For a code with 64QAM modulation} \\ &= 38 \text{ dBmV} \end{aligned}$$

A parameter that may be used to illustrate the effect of increased power per code is the effective transmit power, P_{eff}, the power that would result hypothetically if all N_{act} active codes were transmitted. It is computed as:

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

$$= P_{on} + P_{hr} + 10 \log \frac{1}{N_{act}} \sum_{i=1}^{N_{act}} 10^{G_{const,i}/10}$$

where the last term is the average constellation gain.

For a reference case with all codes transmitted using 64QAM 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 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 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. 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:

$$\textit{effective\_SNR} = \text{measured\_SNR} - P_{sf} + P_{hr} + (G_{const,\ i} - G_{const,\ ranging})$$

where $G_{const, ranging}$ is the constellation gain of the ranging burst that resulted in the SNR measurement. In the MIB, *effective_SNR* corresponds to a reference case with 64QAM 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 the last station maintenance was 17 dB using QPSK modulation ($G_{const, ranging}$ = –1.2 dB) then the effective SNR referenced to 64QAM modulation is:

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

6.2.18.3 Transmit power step size

The step resolution in transmit power MUST be 1 dB or less. When a CM is commanded with finer resolution than it can implement, it MUST round to the nearest supported step size. If the commanded step is halfway between two supported step sizes, the CM MUST choose the smaller step. For example, with a supported step resolution of 1 dB, a command to step ±0.5 dB would result in no step, while a command to step ±0.75 dB would result in a ±1 dB step.

The step size accuracy MUST be within ±0.4 dB. For example, the actual power increase resulting from a command to increase the power level by 1 dB in a CM's next transmitted burst MUST be between 0.6 dB and 1.4 dB.

A relaxation in step size accuracy to ±1.4 dB is allowed for one gain change when changing the power throughout the full power control range in either direction (from low-end to high-end power and vice versa). The locations of these two gain changes with relaxed accuracy MUST be at least 2 dB apart, thus enabling the use of large step attenuators in the coverage of the full power control range (hysteresis effect).

6.2.19 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:

i) the modulation rate (six rates from 160 ksymb/s to 5.12 Msymb/s in octave steps);

ii) the centre frequency (Hz);

iii) the 1536-bit preamble superstring; and

iv) the S-CDMA channel parameters.

The channel parameters are further described in Table 8-18; these characteristics are shared by all users on a given channel. The burst profile attributes are listed in Table 6-7, and are further described in Table 8-19; 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 (see clause 8.3.4).

The CM MUST support all burst profiles commanded by the CMTS via the burst descriptors in the UCD (see clause 8.3.3), and subsequently assigned for transmission in a MAP (see clause 8.3.4).

Table 6-7 – Burst profile attributes

Burst profile attributes	Configuration settings
Modulation	QPSK, 8QAM, 16QAM, 32QAM, 64QAM, 128QAM (TCM only)
Differential encoding	On/off
TCM encoding	On/off
Preamble length	0-1536 bits (see clause 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)[a)]	0 to 255
Guard time	4 to 255 modulation intervals There is no guard time in S-CDMA channels
Last codeword length	Fixed, shortened
Scrambler on/off	On/off
Byte interleaver depth (I_r) [b)]	0 to floor (2048/N_r)[c)]
Byte interleave block size (B_r)[d)]	2 × N_r to 2048
Preamble type	QPSK0/QPSK1
S-CDMA spreader[e)]	On/off
S-CDMA codes per subframe[e)]	1 to 128
S-CDMA interleaver step[e)]	1 to (spreading intervals per frame – 1)

a) 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.

b) If depth = l, no interleaving; if depth = 0, dynamic mode.

c) N_r is the R-S codeword size k + 2T as defined in clause 6.2.6.1.

d) Used only in dynamic mode.

e) Used only for S-CDMA channels.

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 6-8.

Table 6-8 – User-unique burst parameters

User-unique parameter	Adjustment command	Resulting parameter value
Power level	8-bit two's complement, resolution = 0.25 dB	TDMA: +8 to +54 dBmV (32QAM, 64QAM) +8 to +55 dBmV (8QAM, 16QAM) +8 to +58 dBmV (QPSK) S-CDMA: +8 to +53 dBmV (all modulations) Resolution = 1 dB or better
Offset frequency	Range = 32 kHz, resolution = 1 Hz	Range = ±32 kHz; increment = 1 Hz; implement within ±10 kHz
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	Integer part: 0 to (2^{16} – 1), increments of 6.25 µs/64 Fractional part: Unsigned 8-bit fractional extension, units of 6.25 µs/(64 × 256) = 0.38146973 ns
Burst length (mini-slots) if variable on this channel (changes burst-to-burst)	N/A	1 to 255 mini-slots
Transmit equalizer coefficients	DOCSIS 2.0 mode: 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	Up to 64 coefficients; 4 bytes per coefficient: 2 real and 2 imaginary

The CM MUST implement the offset frequency adjustment to effect a change in upstream carrier frequency within 10 Hz of the commanded change.

6.2.19.1 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 of the CM from 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 (see clause 8.3.6), 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 one symbol duration (of the symbol rate in use for a given burst) and (other than a fixed bias) with accuracy within ±0.25 µs plus ±1/2 symbol owing to resolution. As an example, for the maximum symbol rate of 5120 ksymb/s, the corresponding symbol period would be 195 ns, the corresponding maximum resolution for the timing adjustment MUST be 195 ns, and the corresponding minimum accuracy MUST be ±348 ns. The accuracy of CM burst timing of ±0.25 µs plus ±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 MHz) = 6.25 µs/64 ~= 97.66 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 5120 ksymb/s, the corresponding maximum resolution for the timing correction would be 195 ns × (±0.01), or roughly ± 2 ns.

6.2.19.2 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; and
5) ranging offset.

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 µs between the last symbol centre of one burst and the first symbol centre 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 µs between the last symbol centre of one burst and the first symbol centre 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. The output power of the CM MUST be settled to within ±0.1 dB of its final output power level:

a) within 5 µs from the beginning of a change of 1 dB or less; and
b) within 10 µs 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, and 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 (see clause 11.3.2, "Changing upstream channel descriptor message parameters"). Symbol rate MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS, and 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 centre of one burst and the first symbol centre 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, and 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 centre 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 clauses 6.2.21.5 and 6.2.21.6 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, and 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** 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 centre of one burst and the first symbol centre 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. Ranging offset MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS, and 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, if that. 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).

6.2.19.3 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.

6.2.19.4 CM timing offsets when changing modulation rate

When making a modulation rate change, the CM MUST employ the following timing offsets when changing modulation rates. 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. These offsets are maintained into DOCSIS 2.0 but with the addition of including in the table the highest modulation rate. 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 centre of the first symbol in the burst, which is the reference point for burst timing as stated in clause 6.2.20. Specification of these offsets is needed so that CMs apply uniform adjustments to their ranging offsets and so that CMTSs can appropriately handle CMs that apply these offsets when making modulation rate changes.

Modulation rate	Timing offset (in units of 1/64 time ticks referenced to 5.12 MHz)
5.12 MHz	0 (reference)
2.56 MHz	0
1.28 MHz	24
0.64 MHz	72
0.32 MHz	168
0.16 MHz	360

As an 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 clause 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 0.25 s plus 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 necessarily to only these ways:

a) The CMTS may implement the internal physical delay offset, as specified in the above table.

b) The CMTS may implement an internal timing compensation based on the expected offset in the above table.

c) The CMTS may increase the guard time.

d) The CMTS may send an unsolicited RNG-RSP to each CM to adjust the delay offset. As discussed in clause 6.3.6, the CM is expected to be capable of adjusting its timing offset at any time with the accuracy specified within this clause.

6.2.20 Burst timing convention

Figure 6-29 illustrates the nominal burst timing for TDMA channels.

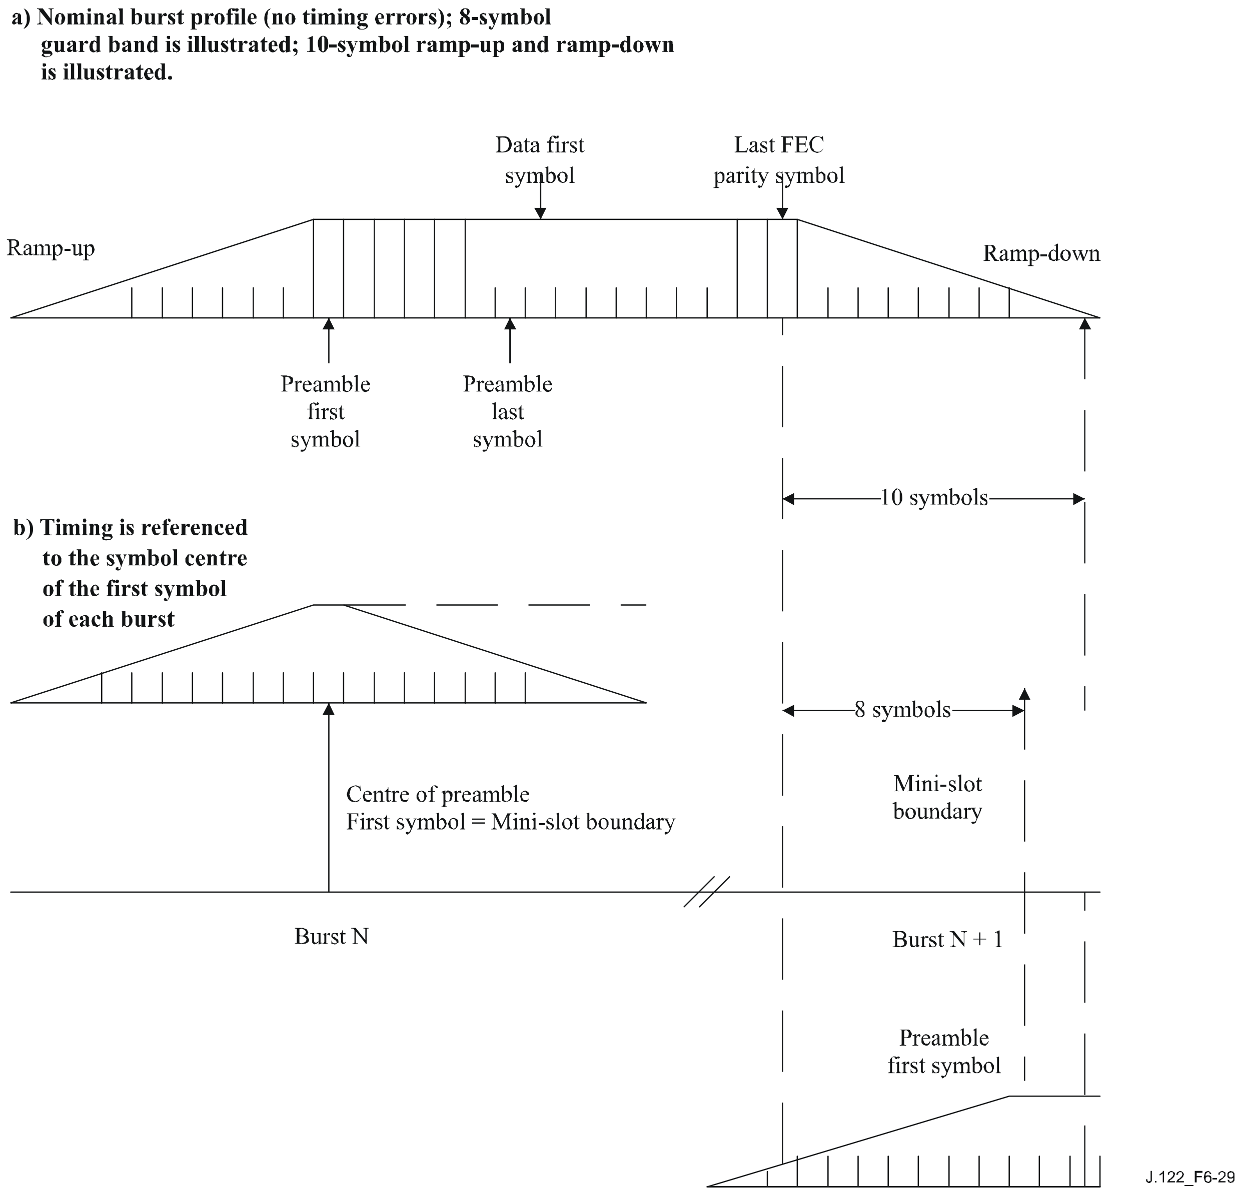

NOTE - Ramp-down of one burst can overlap ramp-up of following burst even with one transmitter assigned both bursts.

Figure 6-29 – Nominal TDMA burst timing

Figure 6-30 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; 8-symbol guard band shown.

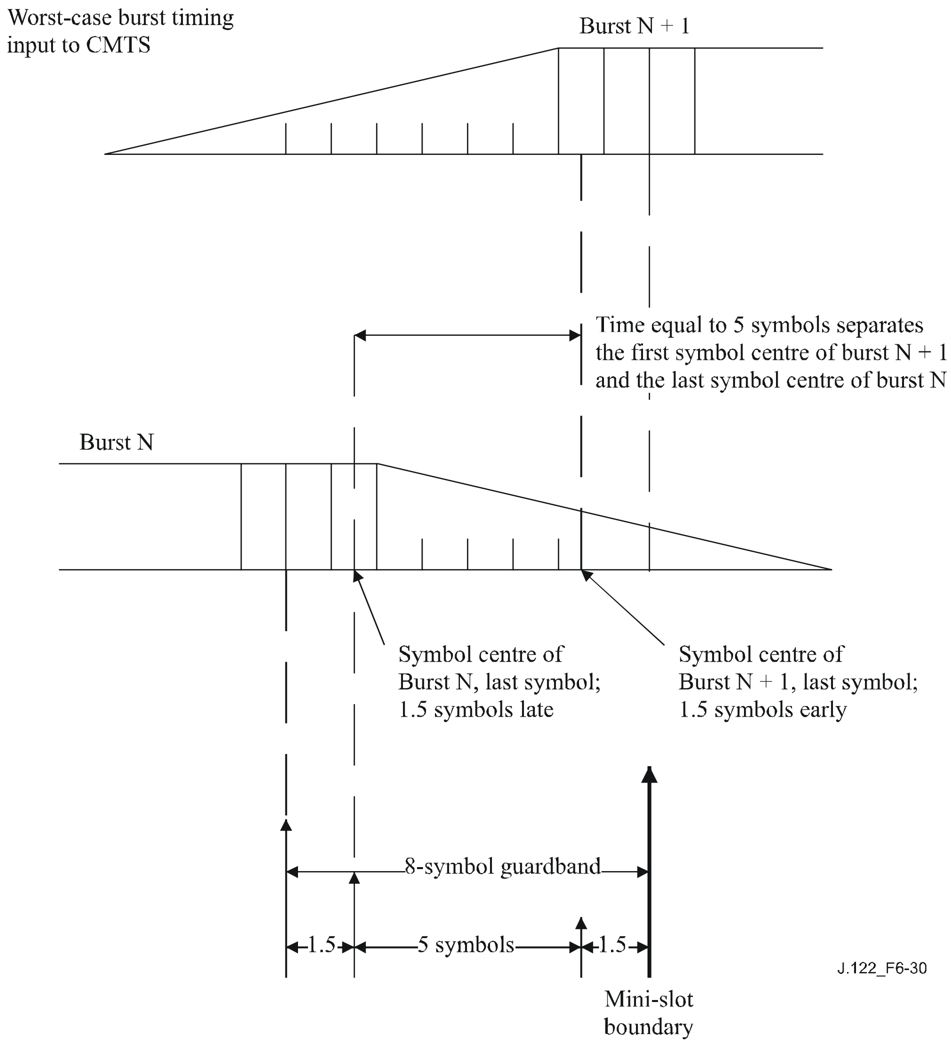

Figure 6-30 – Worst-case TDMA burst timing

At a symbol rate of R_s, symbols occur at a rate of one each $T_s = 1/R_s$ seconds. 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 $N \times T_s$, 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.

6.2.21 Fidelity requirements

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

6.2.21.1 Spurious emissions

The spurious emissions specifications are separated into two regions based on the transmit power. Region 1 is defined to have a power range of +14 dBmV to ($P_{max} - 3$), i.e., the central region. Region 2 is defined from +8 dBmV to +14 dBmV and ($P_{max} - 3$) to P_{max}, i.e., the low and high ends of the transmit power. P_{max} is defined in Table 6-8.

For S-CDMA mode, when a modem is transmitting fewer than 4 spreading codes, the region 2 specifications are used for all transmit power levels. Otherwise, for all other numbers of spreading codes (e.g., 4 to 128) or for TDMA mode, the spurious emissions specifications are used according to the power ranges defined for regions 1 and 2 above.

In addition, for S-CDMA, the spurious emission specifications for S-CDMA MUST be met for any *number_allocated_codes*, as defined in clause 6.2.19.

The noise and spurious power MUST NOT exceed the levels given in Tables 6-9, 6-10 and 6-11.

In Table 6-9, 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., 160 to 5120 kHz). All requirements expressed in dBc are relative to the actual transmit power that the CM emits.

Table 6-9 – Spurious emissions

Parameter	Transmitting burst	Between bursts
Inband	–40 dBc	The greater of –72 dBc or –59 dBmV
Adjacent band	See Table 6-10.	The greater of –72 dBc or –59 dBmV
3 or fewer carrier-related frequency bands (such as second harmonic, if <42 MHz)	Region 1: –50 dBc for transmitted modulation rate = 320 ksymb/s and above; –47 dBc for transmitted modulation rate = 160 ksymb/s Region 2: –47 dBc	The greater of –72 dBc or –59 dBmV
Bands within 5 to 42 MHz (excluding assigned channel, adjacent channels, and carrier-related channels)	See Table 6-11.	The greater of –72 dBc or –59 dBmV
CM integrated spurious emissions limits (all in 4 MHz, includes discretes)[a]		
42 to 54 MHz	Max(–40 dBc, –26 dBmV)	–26 dBmV
54 to 60 MHz	–35 dBmV	–40 dBmV
60 to 88 MHz	–40 dBmV	–40 dBmV
88 to 860 MHz	–45 dBmV	Max(–45 dBmV, –40 dB ref d/s[b])
CM discrete spurious emissions limits		
42 to 54 MHz	Max(–50 dBc, –36 dBmV)	–36 dBmV
54 to 88 MHz	–50 dBmV	–50 dBmV
88 to 860 MHz	–50 dBmV	–50 dBmV

a) 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.

b) "dB ref d/s" is relative to the received downstream signal level. Some spurious outputs are proportional to the receive signal level.

The measurement bandwidth for the three (or fewer) carrier-related frequency bands (below 42 MHz) is 160 kHz, with up to three 160 kHz bands, each with no more than the value given in Table 6-9, allowed to be excluded from the "bands within 5 to 42 MHz transmitting burst" specifications of Table 6-11. Carrier-related spurious emissions include all products whose frequency is a function of the carrier frequency of the upstream transmission such as, but not limited to, carrier harmonics.

The measurement bandwidth is also 160 kHz for the between bursts specifications of Table 6-9 below 42 MHz.

The transmitting burst specifications 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 specifications apply except during a used grant of mini-slots, 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 microseconds at the 5120 ksymb/s rate, or as short as 200 microseconds at the 160 ksymb/s rate.

6.2.21.1.1 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 6-10 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 that is 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 6-10 – Adjacent channel spurious emissions relative to the transmitted burst power level

Transmitted carrier modulation rate	Specification in the interval, region 1	Specification in the interval, region 2	Measurement interval and distance from carrier edge	Adjacent channel carrier modulation rate
All modulation rates	–47 dBc	–45 dBc	20 kHz to 180 kHz	160 kHz
	–47 dBc	–45 dBc	40 kHz to 360 kHz	320 kHz
	–46 dBc	–45 dBc	80 kHz to 720 kHz	640 kHz
	–45 dBc	–44 dBc	160 kHz to 1440 kHz	1280 kHz
	–44 dBc	–41 dBc	320 kHz to 2880 kHz	2560 kHz
	–42 dBc	–38 dBc	640 kHz to 5760 kHz	5120 kHz

6.2.21.1.2 Spurious emissions in 5 to 42 MHz

Spurious emissions, other than those in an adjacent channel or carrier-related emissions listed above, may occur in intervals (frequency bands) that could be occupied by other carriers of the same or different modulation rates. To accommodate these different modulation rates and associated bandwidths, the spurious emissions are measured in an interval equal to the bandwidth corresponding to the modulation rate of the carrier that could be transmitted in that interval. This interval is independent of the current transmitted modulation rate.

Table 6-11 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, 5 MHz or 42 MHz, is reached. Measurement intervals should not include the three or fewer carrier-related emission bands excluded above.

Table 6-11 – Spurious emissions in 5 to 42 MHz relative to the transmitted burst power level

Possible modulation rate in this interval	Specification in the interval, region 1	Specification in the interval, region 2	Initial measurement interval and distance from carrier edge
160 kHz	–54 dBc	–53 dBc	220 kHz to 380 kHz
320 kHz	–52 dBc	–50 dBc	240 kHz to 560 kHz
640 kHz	–50 dBc	–47 dBc	280 kHz to 920 kHz
1280 kHz	–48 dBc	–44 dBc	360 kHz to 1640 kHz
2560 kHz	–46 dBc	–41 dBc	520 kHz to 3080 kHz
5120 kHz	–44 dBc	–38 dBc	840 kHz to 5960 kHz

6.2.21.2 Spurious emissions during burst on/off transients

Each transmitter MUST control spurious emissions, prior to and during ramp-up and during and following ramp-down, 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 µs of constant slewing. This requirement applies when the CM is transmitting at +55 dBmV or more; 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 from +55 dBmV, down to a maximum change of 7 mV at 31 dBmV and below. This requirement does not apply to CM power-on and power-off transients.

6.2.21.3 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.

6.2.21.3.1 Definitions

Symbol MER (**MER_{symb}*****)***: MER_{symb} is defined as follows for TDMA or S-CDMA symbols. The transmitted RF waveform (after appropriate downconversion) 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_{symb}. 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_{symb}. MER_{symb} is defined at the F connector of the CM, except that when an echo channel is inserted, MER_{symb} is defined at the output of the echo channel. MER_{symb} is computed by the formula:

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

where:

E_{av} is the average constellation energy for equally likely symbols (see clause 6.2.13 and Figure 6-18)

N is the number of symbols averaged

e_j is the error vector from the jth received symbol to the ideal transmitted QAM symbol on the grid of Figure 6-18

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

MER of composite chips (**MER_{chip}*****)*:** 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 funnelling 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 $\overline{P_k}$ defined in clause 6.2.14.

MER_{chip} is defined as follows. The transmitted RF waveform (after appropriate downconversion) 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:

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

where:

p_j is the jth ideal transmitted composite chip

r_j is the jth received composite chip

N is the number of composite chips observed

6.2.21.3.2 Requirements

Unless otherwise stated, the MER MUST meet or exceed the following limits over the full transmit power range of Table 6-8 for each modulation, each modulation rate, and over the full carrier frequency range, and for S-CDMA, over any valid number of active and allocated codes. The 5-42 MHz carrier frequency range refers more precisely to [5 MHz + modulation rate × 1.25/2] to [42 MHz – 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:

MER_{symb} ≥ 30 dB over 15 to 30 MHz carrier frequency.

MER_{symb} ≥ 27 dB over 10 to 15 MHz and 30 to 35 MHz carrier frequency.

MER_{symb} ≥ 23 dB over 5 to 10 MHz and 35 to 42 MHz carrier frequency.

Case 1b: for modulation rate 5.12 MHz:

MER_{symb} ≥ 27 dB over 15 to 30 MHz carrier frequency.

MER_{symb} ≥ 24 dB over 10 to 15 MHz and 30 to 35 MHz carrier frequency.

MER_{symb} ≥ 20 dB over 5 to 10 MHz and 35 to 42 MHz carrier frequency.

Case 2: Flat channel, transmit equalization ON

Case 2a: For TDMA/QPSK, MER_{symb} ≥ 30 dB.

Case 2b: For S-CDMA and all TDMA modulations except QPSK, MER_{symb} ≥ 35 dB.

Case 2c: For S-CDMA, MER_{chip} ≥ 33 dB.

Case 3: Echo channel, transmit equalization ON

Case 3a: In the presence of a single echo selected from the channel micro-reflections defined in Table 4-2, 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 4-2 (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.

6.2.21.4 Filter distortion

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

6.2.21.4.1 Amplitude

The spectral mask MUST be the ideal square-root raised-cosine spectrum with alpha = 0.25, within the ranges given in Table 6-12.

Table 6-12 – Filter amplitude distortion

Frequency	Amplitude range	
	Low	High
$f_c - 5R_s/8$	–	–30 dB
$f_c - R_s/2$	–3.5 dB	–2.5 dB
$f_c - 3R_s/8$ to $f_c - R_s/4$	–0.5 dB	+0.3 dB
$f_c - R_s/4$ to $f_c + R_s/4$	–0.3 dB	+0.3 dB
$f_c + R_s/4$ to $f_c + 3R_s/8$	–0.5 dB	+0.3 dB
$f_c + R_s/2$	–3.5 dB	–2.5 dB
$f_c + 5R_s/8$	–	–30 dB

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

6.2.21.4.2 Phase

$f_c - 5R_s/8$ Hz to $f_c + 5R_s/8$ Hz: Group delay variation MUST NOT be greater than 100 ns.

6.2.21.5 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-bursted), 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.

6.2.21.6 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 up to five years from date of manufacture.

6.2.21.7 Modulation rate accuracy

For TDMA mode, the upstream modulator MUST provide an absolute 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.

For S-CDMA mode, the upstream modulator MUST lock the upstream chip rate to the downstream symbol rate, subject to the symbol timing jitter requirements of clause 6.2.21.8.

6.2.21.8 Modulation timing jitter

6.2.21.8.1 Symbol timing jitter for asynchronous operation

For TDMA mode, 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-second period. In other words, the difference between the maximum and the minimum symbol duration during the 2-second period shall be less than 0.02 of the nominal symbol duration for each of the five upstream symbol rates.

For TDMA mode, 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-second period. In other words, the difference between the maximum and the minimum cumulative phase error during the 0.1-second 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.

6.2.21.8.2 Chip timing jitter for synchronous operation

All jitter specifications assume a downstream input to the CM per clauses 6.3.5, 6.3.6, 6.3.7.2, 6.3.7.3, 6.3.9 and 6.3.10.

For S-CDMA mode, upstream chip clock timing error (with the mean error subtracted out) relative to the CMTS master clock MUST be less than 0.005 RMS of the chip period over a 35-second measurement interval. This applies:

1) to 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 chip 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 chip 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-bursted) upstream signal is transmitted at the commanded carrier frequency, modulation rate and level;
- the chip sequence at the spreader output 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 symbol clock as in normal synchronous operation.

6.2.22 Upstream demodulator input power characteristics

The maximum total input power to the upstream demodulator MUST NOT exceed 35 dBmV in the 5-42 MHz frequency range of operation.

The intended received power in each carrier MUST be within the values shown in Table 6-13.

Table 6-13 – Maximum range of commanded nominal receive power in each carrier

Modulation rate [kHz]	Maximum range [dBmV]
160	−16 to + 14
320	−13 to + 17
640	−10 to + 20
1280	–7 to + 23
2560	–4 to + 26
5120	–1 to + 29

The demodulator MUST operate within its defined performance specifications with received bursts within ±6 dB of the nominal commanded received power.

6.2.23 Upstream electrical output from the CM

The CM MUST output an RF modulated signal with the characteristics delineated in Table 6-14.

Table 6-14 – Electrical output from the CM

Parameter	Value
Frequency	5 to 42 MHz edge to edge
Level range (one channel)	TDMA: +8 to +54 dBmV (32QAM, 64QAM) +8 to +55 dBmV (8QAM, 16QAM) +8 to +58 dBmV (QPSK) S-CDMA: +8 to +53 dBmV (all modulations)
Modulation type	QPSK, 8QAM, 16QAM, 32QAM, 64QAM and 128QAM
Modulation rate (nominal)	TDMA: 160, 320, 640,1280, 2560 and 5120 kHz S-CDMA: 1280, 2560 and 5120 kHz
Bandwidth	TDMA: 200, 400, 800,1600, 3200 and 6400 kHz S-CDMA: 1600, 3200 and 6400 kHz
Output impedance	75 ohms
Output return loss	>6 dB (5-42 MHz)
Connector	F connector per [IEC 60169-24] (common with the input)

6.3 Downstream

This Recommendation applies only to a CMTS supporting exactly one QAM channel per RF output port. A CMTS supporting more than one QAM channel per RF output port would instead conform to the DOCSIS downstream radio frequency interface [DOCS9] specification in lieu of this downstream clause.

6.3.1 Downstream protocol

The downstream PMD sublayer MUST conform to Annex B of [ITU-T J.83], for low-delay video applications, with the exceptions called out in clause 6.3.2.

6.3.2 Scalable interleaving to support low latency

The downstream PMD sublayer MUST support a variable-depth interleaver with the characteristics defined in Table 6-15. The table contains a subset of interleaver modes found in Annex B of [ITU-T J.83].

Table 6-15 – Interleaver characteristics

I (number of taps)	J (increment)	Burst protection 64QAM/256QAM	Latency 64QAM/256QAM
8	16	5.9 µs/4.1 µs	0.22 ms/0.15 ms
16	8	12 µs/8.2 µs	0.48 ms/0.33 ms
32	4	24 µs/l 6 µs	0.98 ms/0.68 ms
64	2	47 µs/33 µs	2.0 ms/1.4 ms
128	1	95 µs/66 µs	4.0 ms/2.8 ms

The interleaver depth, which is coded in a 4-bit control word contained in the FEC frame synchronization trailer, always reflects the interleaving in the immediately-following frame. In addition, errors are allowed while the interleaver memory is flushed after a change in interleaving is indicated.

Refer to Annex B of [ITU-T J.83] for the control bit specifications required to specify which interleaving mode is used.

6.3.3 Downstream frequency plan

The downstream frequency plan should comply with harmonic related carrier (HRC), incremental related carrier (IRC) or standard (STD) North American frequency plans per [b-EIA 542]. However, operation below a centre frequency of 91 MHz is not required.

6.3.4 CMTS electrical output

The CMTS MUST output an RF modulated signal with the following characteristics defined in Table 6-16.

Table 6-16 – CMTS output

Parameter	Value
Centre frequency (f_c)	91 to 857 MHz ±30 kHz[a)]
Level	Adjustable over the range 50 to 61 dBmV
Modulation type	64QAM and 256QAM
Symbol rate (nominal) 64QAM 256QAM	 5.056941 Msymb/s 5.360537 Msymb/s
Nominal channel spacing	6 MHz
Frequency response 64QAM 256QAM	 –0.18 square root raised cosine shaping –0.12 square root raised cosine shaping
Total discrete spurious inband (fc ± 3 MHz) Inband spurious and noise (fc ± 3 MHz) Adjacent channel (fc ± 3.0 MHz) to (fc ± 3.75 MHz)	≤57 dBc ≤48 dBc; where channel spurious and noise includes all discrete spurious, noise, carrier leakage, clock lines, synthesizer products and other undesired transmitter products. Noise within ±50 kHz of the carrier is excluded. ≤58 dBc in 750 kHz

Table 6-16 – CMTS output

Parameter	Value
Adjacent channel (fc ± 3.75 MHz) to (fc ± 9 MHz) Next adjacent channel (fc ± 9 MHz), to (fc ± 15 MHz) Other channels (47 MHz to 1000 MHz)	≤62 dBc, in 5.25 MHz, excluding up to 3 spurs, each of which must be ≤60 dBc when measured in a 10-kHz band. Less than the greater of −65 dBc or −12 dBmV in 6 MHz, excluding up to three discrete spurs. The total power in the spurs must be ≤60 dBc when each is measured with 10-kHz bandwidth. <–12 dBmV in each 6-MHz channel, excluding up to three discrete spurs. The total power in the spurs must be ≤60 dBc when each is measured with 10-kHz bandwidth.
Phase noise Output impedance Output return loss	1 to 10 kHz: −33 dBc double-sided noise power 10 to 50 kHz: −51 dBc double-sided noise power 50 kHz to 3 MHz: −51 dBc double-sided noise power 75 ohms >14 dB within an output channel up to 750 MHz; >13 dB in an output channel above 750 MHz
Connector	F connector per [IEC 60169-24]
a) ±30 kHz includes an allowance of 25 kHz for the largest-frequency offset normally built into upconverters.	

6.3.5 Downstream electrical input to CM

The CM MUST be able to locate and accept RF modulated signals located within channels defined in [b-EIA 542] for harmonic related carrier (HRC), incremental related carrier (IRC) and standard (STD) North American frequency plans. Operation below a centre frequency of 91 MHz is not required. The signals will have the characteristics defined in Table 6-17.

Table 6-17 – Electrical input to CM

Parameter	Value
Centre frequency	91 to 857 MHz ±30 kHz
Level range (one channel)	−15 dBmV to +15 dBmV
Modulation type	64QAM and 256QAM
Symbol rate (nominal)	5.056941 Msymb/s (64QAM) and 5.360537 Msymb/s (256QAM)
Bandwidth	6 MHz (alpha = 0.18 square root raised cosine shaping for 64QAM and alpha = 0.12 square root raised cosine shaping for 256QAM)

Table 6-17 – Electrical input to CM

Parameter	Value
Total input power (40-900 MHz)	<30 dBmV
Input (load) impedance	75 ohms
Input return loss	>6 dB (88-860 MHz)
Connector	F connector per [IEC 60169-24] (common with the output)

6.3.6 CM BER performance

The bit-error-rate performance of a CM MUST be as described in this clause. The requirements apply to the I = 128, J = 1 mode of interleaving.

6.3.6.1 64QAM

6.3.6.1.1 64QAM 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) of 23.5 dB or greater.

6.3.6.1.2 64QAM image rejection performance

Performance as described in clause 6.3.6.1.1 MUST be met with an analogue or digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

6.3.6.1.3 64QAM adjacent channel performance

Performance as described in clause 6.3.6.1.1 MUST be met with a digital signal at 0 dBc in the adjacent channels.

Performance as described in clause 6.3.6.1.1 MUST be met with an analogue signal at +10 dBc in the adjacent channels.

Performance as described in clause 6.3.6.1.1, with an additional 0.2-dB allowance, MUST be met with a digital signal at +10 dBc in the adjacent channels.

6.3.6.2 256QAM

6.3.6.2.1 256QAM 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.

Input receive signal level E_s/N_o

−6 dBmV to +15 dBmV 30 dB or greater

Less than –6 dBmV down to −15dBmV 33 dB or greater

6.3.6.2.2 256QAM image rejection performance

Performance as described in clause 6.3.6.2.1 MUST be met with an analogue or a digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

6.3.6.2.3 256QAM adjacent channel performance

Performance as described in clause 6.3.6.2.1 MUST be met with an analogue or a digital signal at 0 dBc in the adjacent channels.

Performance as described in clause 6.3.6.2.1, with an additional 0.5-dB allowance, MUST be met with an analogue signal at +10 dBc in the adjacent channels.

Performance as described in clause 6.3.6.2.1, with an additional 1.0-dB allowance, MUST be met with a digital signal at +10 dBc in the adjacent channels.

6.3.7 CMTS timestamp jitter

The CMTS timestamp jitter MUST be less than 500 ns peak-to-peak at the output of the downstream transmission convergence sublayer. This jitter is relative to an ideal downstream transmission convergence sublayer that transfers the MPEG packet data to the downstream physical media dependent sublayer with a perfectly continuous and smooth clock at the MPEG packet data rate. Downstream physical media dependent sublayer processing MUST NOT be considered in timestamp generation and transfer to the downstream physical media dependent sublayer.

Thus, any two timestamps N1 and N2 (N2 > N1) which were transferred to the downstream physical media dependent sublayer, at times T1 and T2 respectively, must satisfy the following relationship:

$$\left|\frac{N2-N1}{f_{CMTS}}-(T2-T1)\right|<500\times10^{-9}$$

In the equation, the value of (N2 – N1) is assumed to account for the effect of rollover of the timebase counter, and T1 and T2 represent time in seconds. f_{CMTS} is the actual frequency of the CMTS master timebase and may include a fixed frequency offset from the nominal frequency of 10.24 MHz. This frequency offset is bounded by a requirement further below in this clause.

The jitter includes inaccuracy in timestamp value and the jitter in all clocks. The 500 ns allocated for jitter at the downstream transmission convergence sublayer output MUST be reduced by any jitter that is introduced by the downstream physical media dependent sublayer.

The CM is expected to meet the burst timing accuracy requirements in clause 6.2.19 when the time-stamps contain this worst-case jitter.

NOTE – Jitter is the error (i.e., measured) relative to the CMTS master clock (the CMTS master clock is the 10.24-MHz clock used for generating the timestamps).

6.3.7.1 CMTS master clock jitter for asynchronous operation

The CMTS 10.24-MHz master clock MUST have frequency accuracy of ≤ ±5 ppm, drift rate $\leq 10^{-8}$ per second, and edge jitter of ≤10 ns peak-to-peak (±5 ns) over a temperature range of 0 to 40 degrees C up to ten years from date of manufacture[5] (the drift rate and jitter requirements on the CMTS master clock implies that the duration of two adjacent segments of 10'240'000 cycles will be within 30 ns, due to 10 ns jitter on each segments' duration, and 10 ns due to frequency drift. Durations of other counter lengths also may be deduced: adjacent 1'024'000 segments, ≤21 ns; 1'024'000 length segments separated by one 10'240'000 cycle segment, ≤30 ns; adjacent 102'400'000 segments, ≤120 ns. The CMTS master clock MUST meet such test limits in 99% or more measurements).

[5] This Recommendation MAY also be met by synchronizing the CMTS master clock oscillator to an external frequency reference source. If this approach is used, the internal CMTS master clock MUST have frequency accuracy of ±20 ppm over a temperature range of 0 to 40 degrees C up to ten years from date of manufacture when no frequency reference source is connected. The drift rate and edge jitter MUST be as specified above.

6.3.7.2 CMTS master clock jitter for synchronous operation

In addition to the requirements in clause 6.3.7.1, the 10.24 MHz CMTS master clock MUST meet the following double sideband phase noise requirements over the specified frequency ranges:

< [–50 + 20*log(f_{MC}/10.24)] dBc (i.e., < 0.05 ns RMS)	10 to 100 Hz
< [–58 + 20*log(f_{MC}/10.24)] dBc (i.e., < 0.02 ns RMS)	100 to 1 kHz
< [–50 + 20*log(f_{MC}/10.24)] dBc (i.e., < 0.05 ns RMS)	1 to 10 kHz
< [–50 + 20*log(f_{MC}/10.24)] dBc (i.e., < 0.05 ns RMS)	10 kHz to f_{MC}/2

where f_{MC} is the frequency of the measured master clock in MHz. The value of f_{MC} MUST be either an integral multiple or divisor of 10.24 MHz. For example, if a 20.48-MHz oscillator is used as the master clock frequency source, and there is no explicit 10.24-MHz clock to test, the 20.48 MHz clock may be used with f_{MC} equal to 20.48 in the above expressions.

6.3.7.3 CMTS master clock frequency drift for synchronous operation

The frequency of the master clock MUST NOT drift more than 1e-8 per second.

6.3.8 CMTS clock generation

The CMTS has the following three options related to the synchronization of the CMTS master clock and the downstream symbol clock:

1) not locked;
2) downstream symbol clock locked to CMTS master clock;
3) CMTS master clock locked to downstream symbol clock.

For S-CDMA operation, the master clock and the downstream symbol clock MUST be locked using either option 2 or 3.

Let f_b' represent the rate of the downstream symbol clock which is locked to the CMTS master clock, and let f_m represent the rate of the CMTS master clock locked to the downstream symbol clock. Let f_b represent the nominal specified downstream symbol rate, and let f_m represent the nominal CMTS master clock rate (10.24 MHz).

With the downstream symbol clock locked to the CMTS master clock, the following equation MUST hold:

$$f_b' = f_m \times M/N$$

With the CMTS master clock locked to the downstream symbol clock, the following equation MUST hold:

$$f_m' = f_b \times N/M$$

M and N MUST be unsigned integer values each representable in 16 bits (these are specified in the channel TLV parameters of the UCD). When the downstream symbol clock and the CMTS master clock are not locked together, the values of M and N are not valid and are ignored by the CM.

The values of M and N MUST result in a value of f_b' or f_m', which is not more than ±1 ppm from its specified nominal value. Table 6-18 lists the downstream modes of operation, their associated nominal symbol rates, f_b, example values for M and N, the resulting synchronized clock rates, and their offsets from their nominal values.

Table 6-18 – Downstream symbol rates and example parameters for synchronization with the CMTS master clock

Downstream mode	Nominal specified symbol rate, f_b [MHz]	M/N	CMTS master clock rate, f_m' [MHz]	Downstream symbol rate, f_b' [MHz]	Offset from nominal
Annex B, 64QAM	5.056941	401/812	10.239990…	5.056945…	0.95 ppm
Annex B, 256QAM	5.360537	78/149	10.240000…	5.360536…	0.02 ppm

6.3.9 CMTS downstream symbol clock jitter for synchronous operation

The downstream symbol clock MUST meet the following double sideband phase noise requirements over the specified frequency ranges:

< [–53 + 20*log(f_{DS}/5.057)] dBc (i.e., < 0.07 ns RMS)	10 to 100 Hz
< [–53 + 20*log(f_{DS}/5.057)] dBc (i.e., < 0.07 ns RMS)	100 to 1 kHz
< [–53 + 20*log(f_{DS}/5.057)] dBc (i.e., < 0.07 ns RMS)	1 to 10 kHz
< [–36 + 20*log(f_{DS}/5.057)] dBc (i.e., < 0.5 ns RMS)	10 to 100 kHz
< [–30 + 20*log(f_{DS}/5.057)] dBc (i.e., < 1 ns RMS)	100 kHz to (f_{DS}/2)

where f_{DS} is the frequency of the measured clock in MHz. The value of f_{DS} MUST be an integral multiple or divisor of the downstream symbol clock. For example, an f_{DS} = 20.227764-MHz clock may be measured if there is no explicit 5.056941 MHz-clock available.

The CMTS MUST provide a test mode in which:

- the downstream QAM symbol sequence is replaced with an alternating binary sequence (1, –1, 1, –1, 1, –1, ...) at nominal amplitude, on both I and Q;
- the CMTS generates the downstream symbol clock from the 10.24-MHz reference clock as in normal synchronous operation.

If an explicit downstream symbol clock that is capable of meeting the above phase noise requirements is available (e.g., a smooth clock without clock domain jitter), this test mode is not required.

6.3.10 CMTS downstream symbol clock drift for synchronous operation

The frequency of the downstream symbol clock MUST NOT drift more than 1e-8 per second.

7 Downstream transmission convergence sublayer

7.1 Introduction

This clause applies to the first technology option referred to in clause 1.1. For the second and third options, refer to Annexes F and J, respectively.

In order to improve demodulation robustness, facilitate common receiving hardware for both video and data, and provide an opportunity for the possible future multiplexing of video and data over the PMD sublayer bitstream defined in clause 6, a sublayer is interposed between the downstream PMD sublayer and the data-over-cable MAC sublayer.

The downstream bitstream is defined as a continuous series of 188-byte MPEG [ITU-T H.222.0] packets. These packets consist of a 4-byte header followed by 184 bytes of payload. The header identifies the payload as belonging to the data-over-cable MAC. Other values of the header may indicate other payloads. The mixture of MAC payloads and those of other services is optional and is controlled by the CMTS.

Figure 7-1 illustrates the interleaving of data-over-cable (DOC) MAC bytes with other digital information (digital video in the example shown).

Header = DOC	DOC MAC payload
Header = video	Digital video payload
Header = video	Digital video payload
Header = DOC	DOC MAC payload
Header = video	Digital video payload
Header = DOC	DOC MAC payload
Header = video	Digital video payload
Header = video	Digital video payload
Header = video	Digital video payload

Figure 7-1 – Example of interleaving MPEG packets in downstream

7.2 MPEG packet format

The format of an MPEG packet carrying DOCS data is shown in Figure 7-2. The packet consists of a 4-byte MPEG header, a pointer_field (not present in all packets) and the DOCS payload.

MPEG header (4 bytes)	pointer_field (1 byte)	DOCS payload (183 or 184 bytes)

Figure 7-2 – Format of an MPEG packet

7.3 MPEG header for DOCS data-over-cable

The format of the MPEG transport stream header is defined in clause 2.4 of [ITU-T H.222.0]. The particular field values that distinguish data-over-cable MAC streams are defined in Table 7-1. Field names are from [ITU-T H.222.0].

Table 7-1 – MPEG header format for DOCS data-over-cable packets

Field	Length (bits)	Description
sync_byte	8	0 × 47; MPEG packet sync byte
transport_error_indicator	1	Indicates an error has occurred in the reception of the packet. This bit is reset to zero by the sender, and is set to one whenever an error occurs in transmission of the packet.
payload unit_start_indicator	1	A value of one indicates the presence of a pointer_field as the first byte of the payload (fifth byte of the packet)
transport_priority	1	Reserved; set to zero
PID (see Note)	13	DOCS data-over-cable well-known PID (0xlFFE)
transport_scrambling_control	2	Reserved; set to '00'

Table 7-1 – MPEG header format for DOCS data-over-cable packets

Field	Length (bits)	Description
adaptation_field_control	2	'0l'; use of the adaptation_field is NOT ALLOWED on the DOCS PID
continuity_counter	4	Cyclic counter within this PID
NOTE – The MPEG header consists of 4 bytes that begin the 188-byte MPEG packet. The format of the header for use on a DOCS data-over-cable PID is restricted to that shown in this table. The header format conforms to the MPEG standard, but its use is restricted in this Recommendation to NOT ALLOW inclusion of an adaptation_field in the MPEG packets.		

7.4 MPEG payload for DOCS data-over-cable

The MPEG payload portion of the MPEG packet will carry the DOCS MAC frames. The first byte of the MPEG payload will be a 'pointer_field' if the payload_unit_start_indicator (PUSI) of the MPEG header is set.

7.4.1 stuff_byte

This Recommendation defines a stuff_byte pattern having a value (0xFF) that is used within the DOCS payload to fill any gaps between the DOCS MAC frames. This value is chosen as an unused value for the first byte of the DOCS MAC frame. The 'FC' byte of the MAC header will be defined to never contain this value (FC_TYPE = '11' indicates a MAC-specific frame, and FC_PARM = '11111' is not currently used and, according to this Recommendation, is defined as an illegal value for FC_PARM).

7.4.2 pointer_field

The pointer_field is present as the fifth byte of the MPEG packet (first byte following the MPEG header) whenever the PUSI is set to one in the MPEG header. The interpretation of the pointer_field is as follows.

The pointer_field contains the number of bytes in this packet that immediately follow the pointer_field that the CM decoder must skip past before looking for the beginning of a DOCS MAC frame. A pointer field MUST be present if it is possible to begin a data-over-cable MAC frame in the packet, and MUST point to either:

1) the beginning of the first MAC frame to start in the packet; or

2) to any stuff_byte preceding the MAC frame.

7.5 Interaction with the MAC sublayer

MAC frames may begin anywhere within an MPEG packet, MAC frames may span MPEG packets, and several MAC frames may exist within an MPEG packet.

The following figures show the format of the MPEG packets that carry DOCS MAC frames. In all cases, the PUSI flag indicates the presence of the pointer_field as the first byte of the MPEG payload.

Figure 7-3 shows a MAC frame that is positioned immediately after the pointer_field byte. In this case, pointer_field is zero, and the DOCS decoder will begin searching for a valid FC byte at the byte immediately following the pointer_field.

MPEG header (PUSI = 1)	pointer_field (= 0)	MAC frame (up to 183 bytes)	stuff_byte(s) (0 or more)

Figure 7-3 – Packet format where a MAC frame immediately follows the pointer_field

Figure 7-4 shows the more general case where a MAC frame is preceded by the tail of a previous MAC frame and a sequence of stuffing bytes. In this case, the pointer_field still identifies the first byte after the tail of frame #1 (a stuff_byte) as the position where the decoder should begin searching for a legal MAC sublayer FC value. This format allows the multiplexing operation in the CMTS to immediately insert a MAC frame that is available for transmission if that frame arrives after the MPEG header and pointer_field have been transmitted.

MPEG header (PUSI = 1)	pointer_field (= M)	Tail of MAC frame #1 (M bytes)	stuff_byte(s) (0 or more)	Start of MAC frame #2

Figure 7-4 – Packet format with MAC frame preceded by stuffing bytes

In order to facilitate multiplexing of the MPEG packet stream carrying DOCS data with other MPEG-encoded data, the CMTS SHOULD NOT transmit MPEG packets with the DOCS PID which contain only stuff_bytes in the payload area. MPEG null packets SHOULD be transmitted instead. Note that there are timing relationships implicit in the DOCS MAC sublayer which must also be preserved by any MPEG multiplexing operation.

Figure 7-5 shows that multiple MAC frames may be contained within the MPEG packet. The MAC frames may be concatenated one after the other or be separated by an optional sequence of stuffing bytes.

MPEG header (PUSI = 1)	pointer_field (= 0)	MAC frame #1	MAC frame #2	stuff_byte(s) (0 or more)	MAC frame #3

Figure 7-5 – Packet format showing multiple MAC frames in a single packet

Figure 7-6 shows the case where a MAC frame spans multiple MPEG packets. In this case, the pointer_field of the succeeding frame points to the byte following the last byte of the tail of the first frame.

MPEG header (PUSI = 1)	pointer_field (= 0)	stuff_byte(s) (0 or more)	Start of MAC frame #1 (up to 183 bytes)	
MPEG header (PUSI = 0)	Continuation of MAC frame #1 (184 bytes)			
MPEG header (PUSI = 1)	pointer_field (= M)	Tail of MAC frame #1 (M bytes)	stuff_byte(s) (0 or more)	Start of MAC frame #2 (M bytes)

Figure 7-6 – Packet format where a MAC frame spans multiple packets

The transmission convergence sublayer must operate closely with the MAC sublayer in providing an accurate timestamp to be inserted into the time synchronization message (refer to clauses 8.3.2 and 9.3).

7.6 Interaction with the Physical Layer

The MPEG-2 packet stream MUST be encoded according to Annex B of [ITU-T J.83], including MPEG-2 transport framing using a parity checksum as described in Annex B of [ITU-T J.83].

7.7 MPEG header synchronization and recovery

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.

The format of MAC frames is described in detail in clause 8.

8 Media access control specification

8.1 Introduction

8.1.1 Overview

This clause describes version 2.0 of the DOCS MAC protocol. Some of the MAC protocol highlights include:

- bandwidth allocation controlled by CMTS;
- a stream of mini-slots in the upstream;
- dynamic mix of contention- and reservation-based upstream transmit opportunities;
- bandwidth efficiency through support of variable-length packets;
- extensions provided for future support of ATM or other data PDU;
- Quality of service including:
  - support for bandwidth and latency guarantees;
  - packet classification;
  - dynamic service establishment.
- extensions provided for security at the data link layer;
- support for a wide range of data rates.

8.1.2 Definitions

8.1.2.1 MAC-sublayer domain

A MAC-sublayer domain is a collection of upstream and downstream channels for which a single MAC allocation and management protocol operates. Its attachments include one CMTS and some number of CMs. The CMTS MUST service all of the upstream and downstream channels; each CM accesses one logical upstream and one downstream channel at a time. The CMTS MUST police and discard any packets received that have a source MAC address that is not a unicast MAC address. The upstream channels may be any combination of DOCS 1.x or 2.0 formats. A single upstream channel MAY transport DOCS 1.x and 2.0 bursts.

8.1.2.2 MAC service access point

A MAC service access point (MSAP) is an attachment to a MAC-sublayer domain (refer to clause 5.2).

8.1.2.3 Service flows

The concept of service flows is central to the operation of the MAC protocol. Service flows provide a mechanism for upstream and downstream quality of service management. In particular, they are integral to bandwidth allocation.

A service flow ID defines a particular unidirectional mapping between a CM and the CMTS. Active upstream service flow IDs also have associated service IDs or SIDs. Upstream bandwidth is allocated to SIDs, and hence to CMs, by the CMTS. Service IDs provide the mechanism by which upstream quality of service is implemented.

The CMTS MAY assign one or more service flow IDs (SFIDs) to each CM, corresponding to the service flows required by the CM. This mapping can be negotiated between the CMTS and the CM during CM registration or via dynamic service establishment (refer to clause 11.4).

In a basic CM implementation, two service flows (one upstream, one downstream) could be used, for example, to offer best-effort IP service. However, the service flow concept allows for more complex CMs to be developed with support for multiple service classes while supporting interoperability with more basic modems. With these more complex modems, it is possible that certain service flows will be configured in such a way that they cannot carry all types of traffic. That is, they may have a maximum packet size limitation or be restricted to small fixed size unsolicited grants. Furthermore it might not be appropriate to send other kinds of data on service flows that are being used for constant bit rate (CBR)-type applications.

Even in these complex modems, it is necessary to be able to send certain upstream packets needed for MAC management, SNMP management, key management, etc. For the network to function properly, all CMs MUST support at least one upstream and one downstream service flow. These service flows MUST always be provisioned to allow the CM to request and to send the largest possible unconcatenated MAC frame (refer to clause 8.2.2). These service flows are referred to as the upstream and downstream primary service flows. The SID assigned to the upstream primary service flow is referred to as the primary SID.

The primary SID MUST always be assigned to the first provisioned upstream service flow during the registration process (which may or may not be the same temporary SID used for the registration process). The primary service flows MUST be immediately activated at registration time. The primary SID MUST always be used for periodic ranging after registration. The primary service flows MAY be used for traffic. All unicast service flows MUST use the security association defined for the primary service flow (refer to [DOCS8]).

All service flow IDs are unique within a single MAC-sublayer domain. The mapping of a unicast service identifier to an active/admitted service flow MUST be unique within a single MAC-sublayer domain. The length of the service flow ID is 32 bits. The length of the service ID is 14 bits (although the service ID is sometimes carried in the low-order bits of a 16-bit field).

Unicast flows on different logical upstreams that are attached to a single MAC-sublayer domain MAY be assigned the same SID, as long as the SFIDs are unique.

8.1.2.4 Upstream intervals, mini-slots and 6.25-microsecond increments

The upstream transmission time-line is divided into intervals by the upstream bandwidth allocation mechanism. Each interval is an integral number of mini-slots. A "mini-slot" is the unit of granularity for upstream transmission opportunities. There is no implication that any PDU can actually be transmitted in a single mini-slot. Each interval is labelled with a usage code which defines both the type of traffic that can be transmitted during that interval and the physical-layer modulation encoding. The usage code values are defined in Table 8-20 and allowed use is defined in clause 8.3. The binding of these values to physical-layer parameters is defined in Table 8-18.

8.1.2.4.1 TDMA mode

For DOCS 1.x channels, a mini-slot is a power-of-two multiple of 6.25-µs increments, i.e., 2, 4, 8, 16, 32, 64 or 128 times 6.25 µs. For DOCS 2.0 TDMA, a mini-slot is a power-of-two multiple of 6.25-µs increments: 1, 2, 4, 8, 16, 32, 64 or 128 times 6.25 µs. The relationship between mini-slots, bytes, and time ticks is described further in clause 9.3.4.

8.1.2.4.2 S-CDMA mode

For DOCS 2.0 S-CDMA channels, a mini-slot is not restricted to be a power-of-two multiple of 6.25-µs increments. Instead a mini-slot is a unit of capacity that is dependent on the modulation rate, number of spreading codes, and number of spreading intervals configured for the upstream channel[6]. While the channel may be configured such that the time duration of a mini-slot is a power-of-two multiple of 6.25-µs increments, there is no special significance to 6.25-µs time ticks for S-CDMA channels. The relationship between mini-slots and S-CDMA framing is described further in clause 6.2.11. The relationship between mini-slots, bytes and time ticks is described further in clause 9.3.4.

8.1.2.5 MAC frame

A MAC layer frame is a unit of data exchanged between two (or more) entities at the data link layer. A MAC frame consists of a MAC header (beginning with a frame control byte; see Figure 8-3), and may incorporate a variable-length data PDU. The variable-length PDU includes a pair of 48-bit addresses, data and a CRC. In special cases, the MAC header may encapsulate multiple MAC frames (see clause 8.2.5.5) into a single MAC frame. The MAC layer definition of a frame is different from any physical layer or MPEG layer definition of a frame.

8.1.2.6 Logical upstream channels

The MAC layer deals with logical upstreams. A logical upstream is identified with an upstream channel ID which is unique within the MSAP. A logical upstream consists of a contiguous stream of mini-slots which are described and allocated by the UCD and MAP messages associated with a channel ID. A CM can only register to operate on a single logical upstream channel.

There are 4 distinct types of logical upstream:

- Type 1: DOCS 1.x upstreams that support no DOCS 2.0 TDMA features;
- Type 2: Mixed upstreams that support DOCS 1.x and DOCS 2.0 TDMA bursts;
- Type 3A: DOCS 2.0 TDMA-only upstreams that cannot support DOCS 1.x CMs;
- Type 3S: S-CDMA upstreams that support only CMs operating in S-CDMA mode.

All valid logical upstreams fall into one of these 4 categories.

In DOCS 2.0 it is possible for multiple logical upstreams to share the same spectrum. When this occurs, the logical upstreams sharing the same spectrum are time domain multiplexed and only one is active at any time. The only exception to this rule is that it is possible for the broadcast initial maintenance regions to be simultaneous. When a logical upstream channel is inactive, its mini-slots are allocated to the NULL SID by its associated MAP messages. Having multiple logical upstreams that share the same spectrum is the only way to have modems operating in S-CDMA mode share the same upstream spectrum with modems not operating in S-CDMA mode.

The CMTS MUST support all four logical upstream channel types individually, and MUST also support the following three combinations of two logical upstream channels sharing the same upstream spectrum:

[6] This relationship holds true on an S-CDMA channel even if the burst parameters for a particular IUC have the spreader disabled.

- one DOCS 1.x-only logical channel plus one S-CDMA logical channel with the same modulation rate on both logical channels;
- one mixed DOCS 1.x and A-TDMA logical channel plus one S-CDMA logical channel with the same modulation rate on both logical channels;
- one A-TDMA-only logical channel plus one S-CDMA logical channel with the same modulation rate on both logical channels.

The CMTS MAY support other combinations of logical channels sharing the same upstream spectrum, including combinations of logical channels with different modulation rates.

8.1.2.7 DOCS 2.0-only logical upstreams

DOCS 2.0-only logical upstreams have operational parameters in their associated UCD messages that prevent the operation of DOCS 1.x CMs. See clause 8.3.3 for a detailed description of which parameter values make a channel a DOCS 2.0-only upstream. The UCD messages for DOCS 2.0-only logical upstreams use a different MAC management message type (see clause 8.3.1) than do UCD messages for channels that can support 1.x CMs. This prevents 1.x CMs from attempting to use DOCS 2.0-only upstreams or from being confused by UCD messages for those channels. A logical upstream is a DOCS 2.0-only upstream if and only if it is an S-CDMA upstream or a DOCS 2.0 TDMA-only upstream.

8.1.3 Future use

A number of fields are defined as being "for future use" or reserved in the various MAC frames described in this Recommendation. These fields MUST NOT be interpreted or used in any manner by this version (2.0) of the MAC protocol.

8.2 MAC frame formats

8.2.1 Generic MAC frame format

A MAC frame is the basic unit of transfer between MAC sublayers at the CMTS and the cable modem. The same basic structure is used in both the upstream and downstream directions. MAC frames are variable in length. The term "frame" is used in this context to indicate a unit of information that is passed between MAC sublayer peers. This is not to be confused with the term "framing" that indicates some fixed timing relationship.

There are three distinct regions to consider, as shown in Figure 8-1. Preceding the MAC frame is either PMD sublayer overhead (upstream) or an MPEG transmission convergence header (downstream). The first part of the MAC frame is the MAC header. The MAC header uniquely identifies the contents of the MAC frame. Following the header is the optional data PDU region. The format of the data PDU and whether it is even present is described in the MAC header.

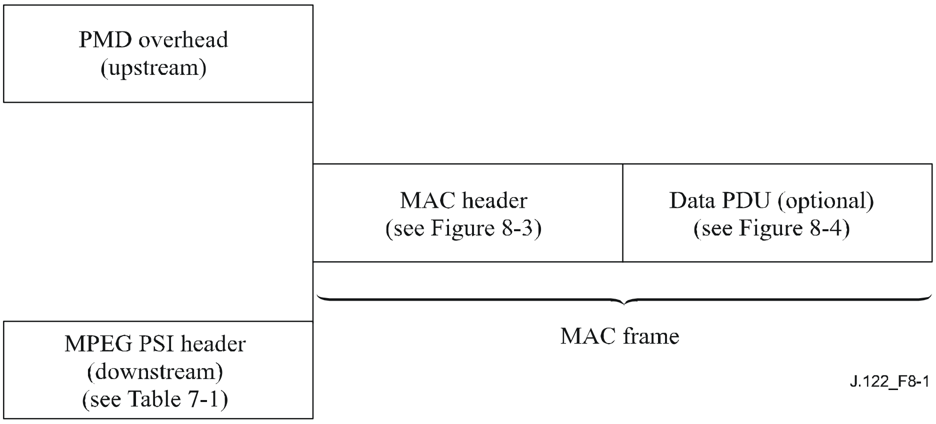

Figure 8-1 – Generic MAC frame format

8.2.1.1 PMD overhead

In the upstream direction, the PHY layer indicates the start of the MAC frame to the MAC sublayer. From the MAC sublayer's perspective, it only needs to know the total amount of overhead so it can account for it in the bandwidth allocation process. More information on this may be found in the PMD sublayer clause (clause 6).

The FEC overhead is spread throughout the MAC frame and is assumed to be transparent to the MAC data stream. The MAC sublayer does need to be able to account for the overhead when doing bandwidth allocation. More information on this may be found in the upstream bandwidth allocation clause (refer to clause 9.1).

8.2.1.2 MAC frame transport

The transport of MAC frames by the PMD sublayer for upstream channels is shown in Figure 8-2.

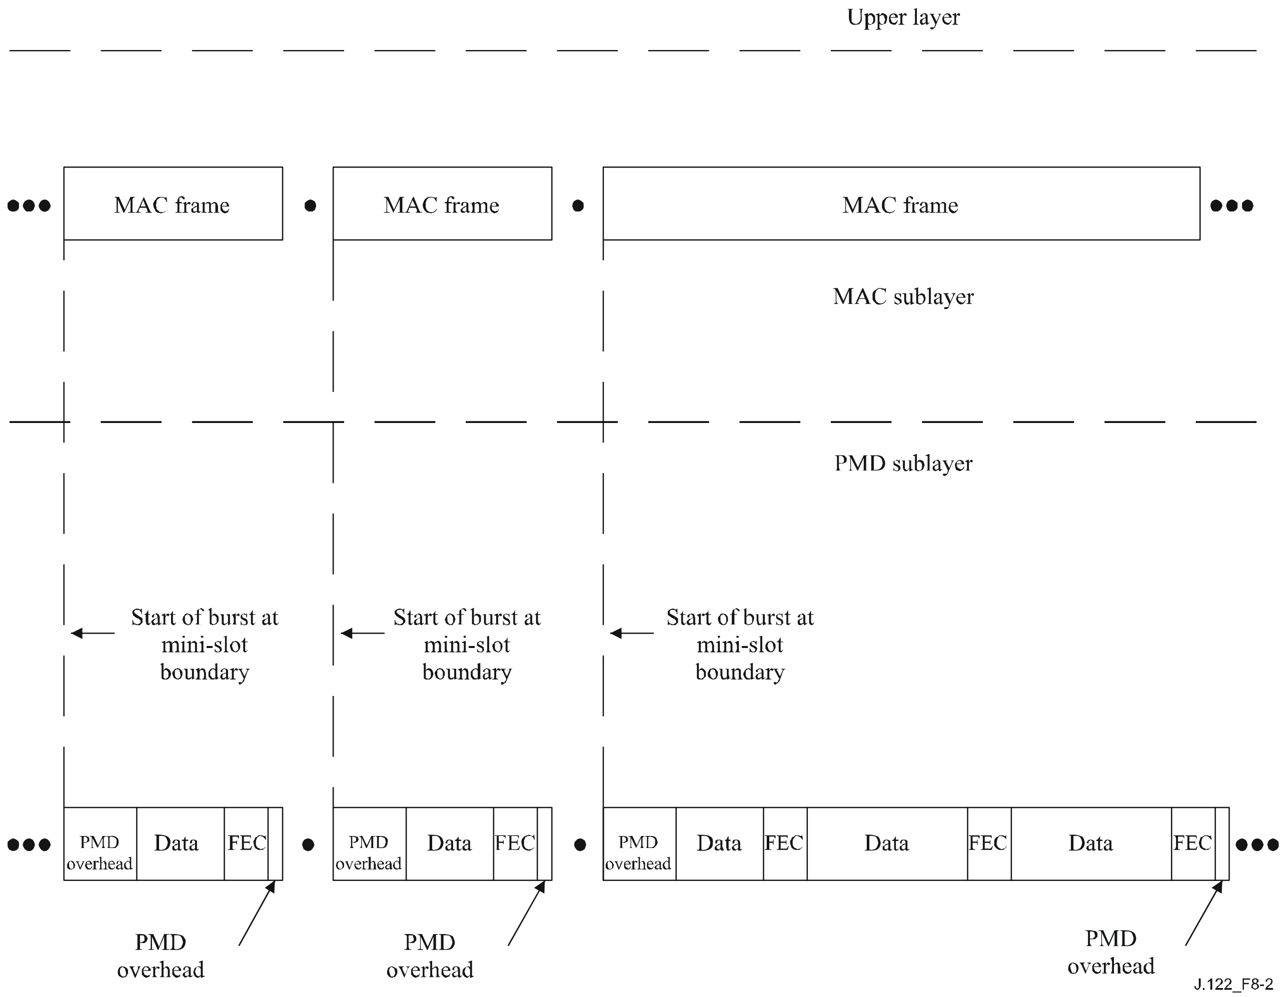

Figure 8-2 – Upstream MAC/PMD convergence

The layering of MAC frames over MPEG in the downstream channel is described in clause 7.

Note that the CMTS PHY ensures that the CMTS MAC receives upstream MAC frames in the same order the CM mapped the MAC frames onto mini-slots. That is to say that if MAC frame X begins in mini-slot n and MAC frame Y begins in mini-slot n+m, then the CMTS MAC will receive X before it receives Y. This is true even when, as is possible with S-CDMA, mini-slots n and n+m are actually simultaneously transmitted within the PHY layer.

8.2.1.3 Ordering of bits and octets

Within an octet, the least-significant bit is the first transmitted on the wire. This follows the convention used by Ethernet and [ISO/IEC 8802-3]. This is often called bit-little-endian order[7].

Within the MAC layer, when numeric quantities are represented by more than one octet (i.e., 16-bit and 32-bit values), the octet containing the most-significant bits is the first transmitted on the wire. This is sometimes called byte-big-endian order.

[7] This applies to the upstream channel only. For the downstream channel, the MPEG transmission convergence sublayer presents an octet-wide interface to the MAC, so the MAC sublayer does not define bit order.

This clause follows the textual convention that when bit-fields are presented in tables, the most-significant bits are topmost in the table. For example, in Table 8-2, FC_TYPE occupies the two most-significant bits and EHDR_ON occupies the least-significant bit.

8.2.1.3.1 Representing negative numbers

Signed integer values MUST be transmitted and received in two's complement format.

8.2.1.3.2 Type-Length-Value fields

Many MAC messages incorporate type-length-value (TLV) fields. TLV fields are unordered lists of TLV-tuples. Some TLVs are nested (see Annex C). All TLV length fields, except for EH_LEN (see clause 8.2.6), MUST be greater than zero. Unless otherwise specified, type is one byte and length is one byte.

Using this encoding, new parameters may be added that some devices cannot interpret. A CM or CMTS which does not recognize a parameter type MUST skip over this parameter and MUST NOT treat the event as an error condition.

8.2.1.4 MAC header format

The MAC header format MUST be as shown in Figure 8-3.

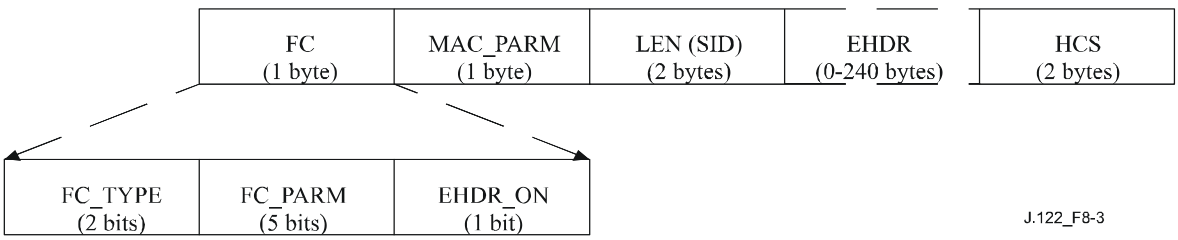

Figure 8-3 – MAC header format

All MAC headers MUST have the general format as shown in Table 8-1. The frame control (FC) field is the first byte and uniquely identifies the rest of the contents within the MAC header. The FC field is followed by 3 bytes of MAC control; an OPTIONAL extended header (EHDR) field; plus a header check sequence (HCS) to ensure the integrity of the MAC header.

Table 8-1 – Generic MAC header format

MAC header field	Usage	Size
FC	Frame control: Identifies type of MAC header	8 bits
MAC_PARM	Parameter field whose use is dependent on FC: if EHDR_ON=l; used for EHDR field length (ELEN) else if for concatenated frames (see Table 8-10) used for MAC frame count else (for requests only) indicates the number of mini-slots requested	8 bits
LEN (SID)	The length of the MAC frame. The length is defined to be the sum of the number of bytes in the extended header (if present) and the number of bytes following the HCS field (for a REQ header, this field is the service ID instead).	16 bits
EHDR	Extended MAC header (where present; variable size)	0-240 bytes
HCS	MAC header check sequence	2 bytes
	Length of a MAC header	6 bytes + EHDR

The HCS field is a 16-bit CRC that ensures the integrity of the MAC header, even in a collision environment. The HCS field coverage MUST include the entire MAC header, starting with the FC field and including any EHDR field that may be present. The HCS is calculated using ITU-T CRC ($x^{16} + x^{12} + x^5 + 1$) as defined in [ITU-T X.25].

The FC field is broken down into the FC_TYPE sub-field, FC_PARM sub-field and an EHDR_ON indication flag. The format of the FC field MUST be as shown in Table 8-2.

Table 8-2 – FC field format

FC field	Usage	Size
FC_TYPE	MAC frame control type field: 00: Packet PDU MAC header 01: ATM PDU MAC header 10: Reserved PDU MAC header 11: MAC-specific header	2 bits
FC_PARM	Parameter bits, use dependent on FC_TYPE	5 bits
EHDR_ON	When = 1, indicates that EHDR field is present. (Length of EHDR (ELEN) determined by MAC_PARM field)	1 bit

The FC_TYPE sub-field is the two MSBs of the FC field. These bits MUST always be interpreted in the same manner to indicate one of four possible MAC frame formats. These types include: MAC header with packet PDU; MAC header with ATM cells; MAC header reserved for future PDU types; or a MAC header used for specific MAC control purposes. These types are spelled out in more detail in the remainder of this clause.

The five bits following the FC_TYPE sub-field is the FC_PARM sub-field. The use of these bits is dependent on the type of MAC header. The LSB of the FC field is the EHDR_ON indicator. If this bit is set, then an extended header (EHDR) is present. The EHDR provides a mechanism to allow the MAC header to be extensible in an inter-operable manner.

The transmission convergence sublayer stuff-byte pattern is defined to be a value of 0xFF. This precludes the use of FC byte values which have FC_TYPE = '11' and FC_PARM = '11111'.

The MAC_PARM field of the MAC header serves several purposes depending on the FC field. If the EHDR_ON indicator is set, then the MAC_PARM field MUST be used as the extended header length (ELEN). The EHDR field may vary from 0 to 240 bytes. If this is a concatenation MAC header, then the MAC_PARM field represents the number of MAC frames (CNT) in the concatenation (see clause 8.2.5.5). If this is a request MAC header (REQ) (see clause 8.2.5.3), then the MAC_PARM field represents the amount of bandwidth being requested. In all other cases, the MAC_PARM field is reserved for future use.

The third field has two possible uses. In most cases, it indicates the length (LEN) of this MAC frame. In one special case, the request MAC header, it is used to indicate the cable modem's service ID since no PDU follows the MAC header.

The extended header (EHDR) field provides extensions to the MAC frame format. It is used to implement data link security as well as frame fragmentation, and can be extended to add support for additional functions in future releases. Initial implementations SHOULD pass this field to the processor. This will allow future software upgrades to take advantage of this capability (refer to clause 8.2.6 for details).

8.2.1.5 Data PDU

The MAC header may be followed by a data PDU. The type and format of the data PDU is defined in the frame control field of the MAC header. The FC field explicitly defines a packet data PDU, an ATM data PDU, a MAC-specific frame and a reserved code point (used as an escape mechanism for future extensions). All CMs MUST use the length in the MAC header to skip over any reserved data.

8.2.2 Packet-based MAC frames

8.2.2.1 Variable-length packets

The MAC sublayer MUST support a variable-length Ethernet/[ISO/IEC 8802-3]-type packet data PDU. With the exception of packets that have been subject to payload header suppression, the packet PDU MUST be passed across the network in its entirety, including its original CRC. In the case where payload header suppression has been applied to the packet PDU, all bytes except those suppressed MUST be passed across the network, and the CRC covers only those bytes actually transmitted (refer to clause 8.2.6.3.1). A unique packet MAC header is appended to the beginning. The frame format without an extended header MUST be as shown in Figure 8-4 and Table 8-3.

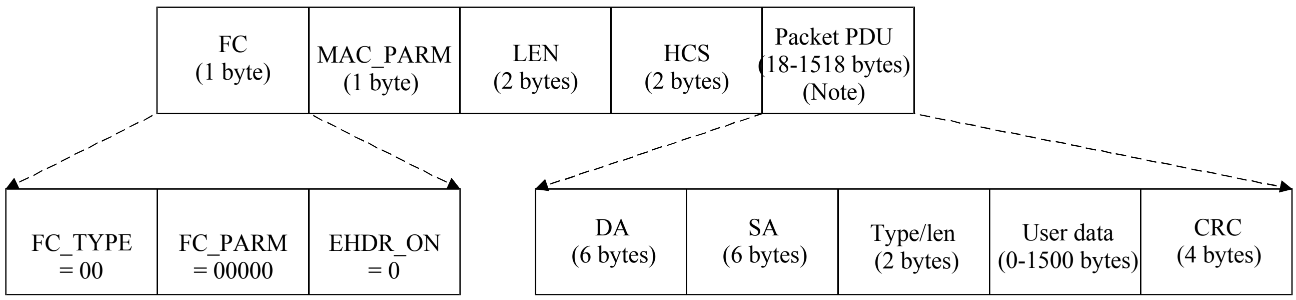

NOTE – Frame size is limited to 1518 bytes in the absence of VLAN tagging. Cooperating devices which implement IEE 802.1Q VLAN tagging MAY use a frame size up to 1522 bytes.

J.122_F8-4

Figure 8-4 – Ethernet/802.3 packet PDU format

Table 8-3 – Packet PDU format

Field	Usage	Size
FC	FC_TYPE = 00; packet MAU header FC_PARM[4:0] = 00000; other values reserved for future use and ignored EHDR_ON = 0 if there is no extended header, 1 if there is an EHDR	8 bits
MAC_PARM	MAC_PARM = x; MUST be set to zero if there is no EHDR; Otherwise set to length of EHDR	8 bits
LEN	LEN = n + x; length of packet PDU in bytes + length of EHDR	16 bits
EHDR	Extended MAC header, if present	x (0-240) bytes
HCS	MAC header check sequence	16 bits
Packet data, packet PDU	DA – 48-bit destination address SA – 48-bit source address Type/LEN – 16-bit Ethernet type or [ISO/IEC 8802-3] length field User data (variable length, 0-1500 bytes) CRC – 32-bit CRC over packet PDU (as defined in Ethernet/[ISO/IEC 8802-3])	n bytes
	Length of packet MAC frame	6 + x + n bytes

Under certain circumstances (see Appendix VI) it may be necessary to transmit a packet PDU MAC frame without an actual PDU. This is done so that the extended header can be used to carry certain information about the state of the service flow. This could also happen as a result of PHS (see clause 8.2.6.3.1). Such a frame will have the length field in MAC header set to the length of the extended header and will have no packet data, and therefore no CRC. This can only happen with frames transmitted on the upstream as frames transmitted on the downstream always have at least the DA and SA fields of the packet PDU.

8.2.3 ATM cell MAC frames

The FC_TYPE 0x01 is reserved for future definition of ATM cell MAC frames. This FC_TYPE field in the MAC header indicates that an ATM PDU is present. This PDU MUST be silently discarded by MAC implementations of this version of the Recommendation. Implementations MUST use the length field to skip over the ATM PDU.

8.2.4 Reserved PDU MAC frames

The MAC sublayer provides a reserved FC code point to allow for support of future (to be defined) PDU formats. The FC field of the MAC header indicates that a reserved PDU is present. This PDU MUST be silently discarded by MAC implementations of this version of the Recommendation. Implementations compliant to this version of the Recommendation MUST use the length field to skip over the reserved PDU.

The format of the reserved PDU without an extended header MUST be as shown in Figure 8-5 and Table 8-4.

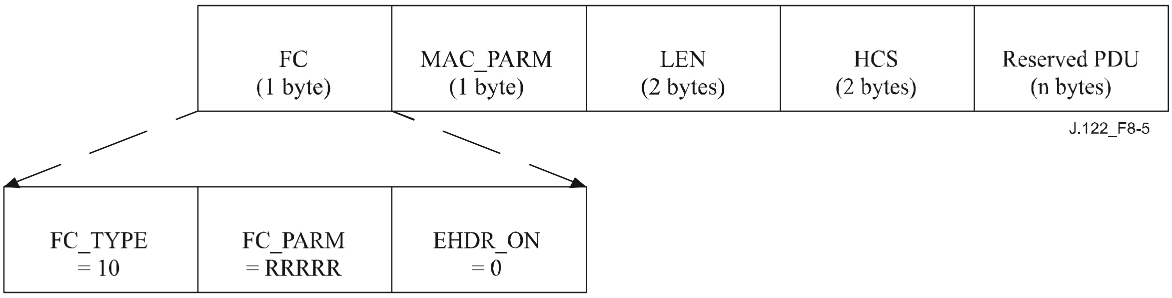

Figure 8-5 – Reserved PDU format

Table 8-4 – Reserved PDU format

Field	Usage	Size
FC	FC_TYPE = 10; reserved PDU MAC header FC_PARM[4:0]; reserved for future use EHDR_ON = 0 if there is no extended header, 1 if there is an EHDR	8 bits
MAC_PARM	MAC_PARM = x; MUST be set to zero if there is no EHDR; Otherwise set to length of EHDR	8 bits
LEN	LEN = n + x; length of reserved PDU + length of EHDR in bytes	16 bits
EHDR	Extended MAC header, if present	x (0-240) bytes
HCS	MAC header check sequence	16 bits
User data	Reserved data PDU	n bytes
	Length of reserved PDU MAC frame	6 + x + n bytes

8.2.5 MAC-specific headers

There are several MAC headers which are used for very specific functions. These functions include support for downstream timing and upstream ranging/power adjust, requesting bandwidth, fragmentation and concatenating multiple MAC frames.

Table 8-5 describes FC_PARM usage within the MAC-specific header.

Table 8-5 – MAC-specific headers and frames

FC_PARM	Header/Frame type
00000	Timing header
00001	MAC management header
00010	Request frame
00011	Fragmentation header
11100	Concatenation header

8.2.5.1 Timing header

A specific MAC header is identified to help support the timing and adjustments required. In the downstream, this MAC header MUST be used to transport the global timing reference to which all cable modems synchronize. In the upstream, this MAC header MUST be used as part of the ranging message needed for a cable modem's timing and power adjustments. The timing MAC header is followed by a packet data PDU. The format MUST be as shown in Figure 8-6 and Table 8-6.

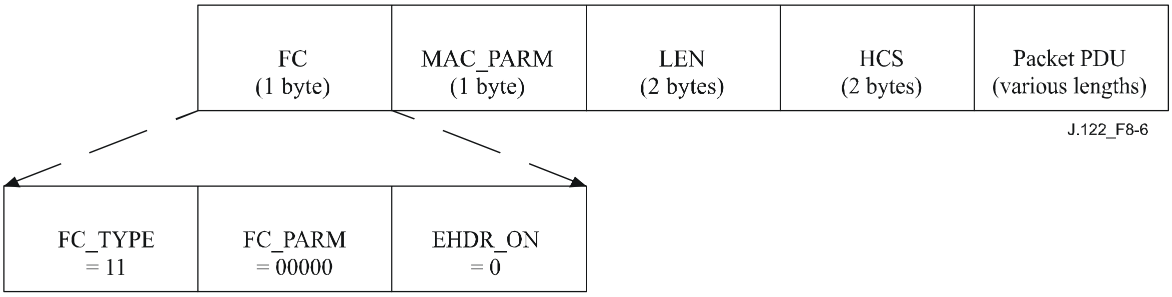

Figure 8-6 – Timing MAC header

Table 8-6 – Timing MAC header format

Field	Usage	Size
FC	FC_TYPE = 11; MAC-specific header FC_PARM[4:0] = 00000; timing MAC header EHDR_ON = 0; extended header prohibited for SYNC and RNG-REQ	8 bits
MAC_PARM	Reserved for future use	8 bits
LEN	LEN = n; length of packet PDU in bytes	16 bits
EHDR	Extended MAC header (not present)	0 bytes
HCS	MAC header check sequence	2 bytes
Packet data	MAC management message: SYNC message (downstream only) RNG-REQ (upstream only)	n bytes
	Length of timing message MAC frame	6 + n bytes

8.2.5.2 MAC management header

A specific MAC header is identified to help support the MAC management messages required. This MAC header MUST be used to transport all MAC management messages (refer to clause 8.3). The format MUST be as shown Figure 8-7 and Table 8-7.

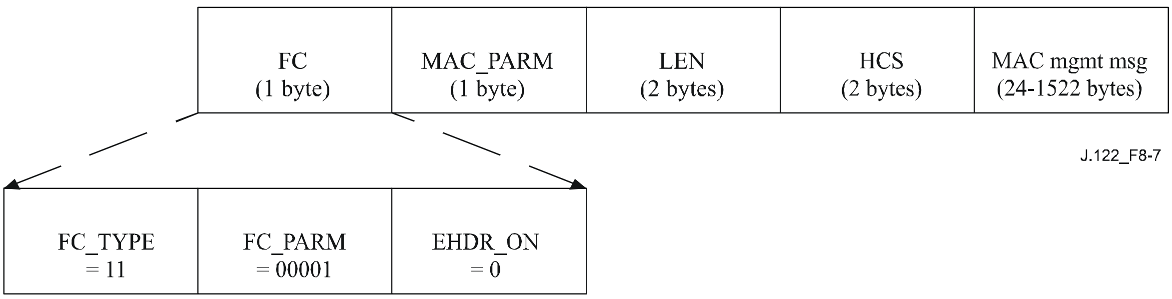

Figure 8-7 – MAC management header

Table 8-7 – MAC management header format

Field	Usage	Size
FC	FC_TYPE = 11; MAC-specific header FC_PARM[4:0] = 00001; MAC management header EHDR_ON = 0 if there is no extended header, 1 if there is an EHDR	8 bits
MAC_PARM	MAC_PARM = x; MUST be set to zero if there is no EHDR; Otherwise set to length of EHDR	8 bits
LEN	LEN = n + x; length of MAC management message + length of EHDR in bytes	16 bits
EHDR	Extended MAC header, if present	x (0-240) bytes
HCS	MAC header check sequence	16 bits
Packet data	MAC management message	n bytes
	Length of packet MAC frame	6 + x + n bytes

8.2.5.3 Request frame

The request frame is the basic mechanism that a cable modem uses to request bandwidth. As such, it is only applicable in the upstream. There MUST be no data PDUs following the request frame. The general format of the request MUST be as shown in Figure 8-8 and Table 8-8.

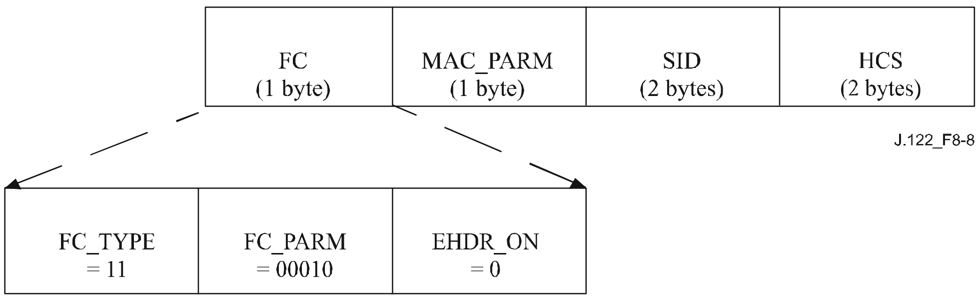

Figure 8-8 – Request frame format

Table 8-8 – Request frame (REQ) format

Field	Usage	Size
FC	FC_TYPE = 11; MAC-specific header FC_PARM[4:0] = 00010; MAC header only; no data PDU following EHDR_ON = 0; no EHDR allowed	8 bits
MAC_PARM	REQ; total number of mini-slots requested	8 bits
SID	Service ID used for requesting bandwidth. For valid SID ranges, see clause 9.1.2.	16 bits
EHDR	Extended MAC header not allowed	0 bytes
HCS	MAC header check sequence	2 bytes
	Length of a REQ MAC header	6 bytes

Because the request frame does not have a data PDU following it, the LEN field is not needed. The LEN field MUST be replaced with an SID. The SID MUST uniquely identify a particular service flow within a given CM.

The bandwidth request, REQ, MUST be specified in mini-slots. The REQ field MUST indicate the current total amount of bandwidth requested for this service queue, including appropriate allowance for the PHY overhead.

8.2.5.4 Fragmentation header

The fragmentation MAC header provides the basic mechanism to split a larger MAC PDU into smaller pieces that are transmitted individually and then reassembled at the CMTS. As such, it is only applicable in the upstream. The general format of the fragmentation MAC header MUST be as shown in Figure 8-9.

A compliant CM MUST support fragmentation. A compliant CMTS MUST support fragmentation. To decrease the burden on the CMTS and to reduce unnecessary overhead, fragmentation headers MUST NOT be used on unfragmented frames.

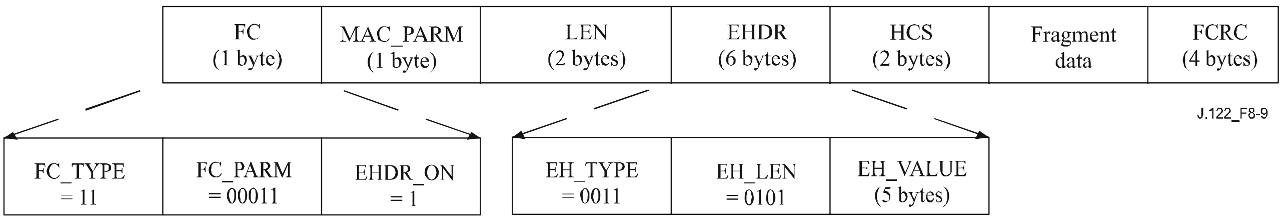

Figure 8-9 – Fragmentation MAC header format

Table 8-9 – Fragmentation MAC frame (FRAG) format

Field	Usage	Size
FC	FC_TYPE = 11; MAC-specific header FC_PARM [4:0] = 00011; fragmentation MAC header EHDR_ON = 1; fragmentation EHDR follows	8 bits
MAC_PARM	ELEN = 6 bytes; length of fragmentation EHDR	8 bits
LEN	LEN = length of fragment payload + EHDR length + FCRC length	16 bits
EHDR	Refer to clause 8.2.6.2	6 bytes
HCS	MAC header check sequence	2 bytes
Fragment data	Fragment payload; portion of total MAC PDU being sent	n bytes
FCRC	CRC – 32-bit CRC over fragment data payload (as defined in Ethernet/ [ISO/IEC 8802-3])	4 bytes
	Length of a MAC fragment frame	16 + n bytes

8.2.5.5 Concatenation header

A specific MAC header is defined to allow multiple MAC frames to be concatenated. This allows a single MAC "burst" to be transferred across the network. The PHY overhead[8] and the concatenation MAC header only occur once. Concatenation of multiple MAC frames MUST be as shown in Figure 8-10. Concatenation of multiple MAC frames is the only method by which the CM can transmit more than one MAC frame in a single transmit opportunity.

[8] This includes the preamble, guard time and possibly zero-fill bytes in the last codeword. The FEC overhead recurs for each codeword.

A compliant CM MUST support concatenation. A compliant CMTS MUST support concatenation. Concatenation only applies to upstream traffic. Concatenation MUST NOT be used on downstream traffic.

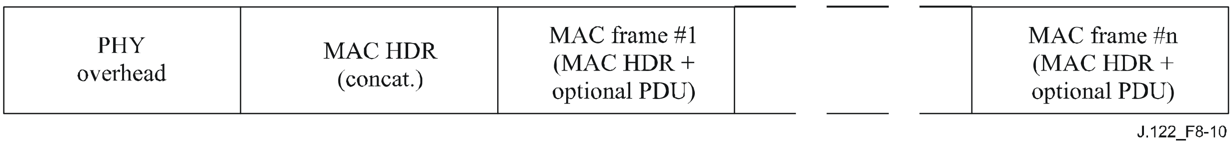

Figure 8-10 – Concatenation of multiple MAC frames

Only one concatenation MAC header MUST be present per MAC "burst." Nested concatenation MUST NOT be allowed. Immediately following the concatenation MAC header MUST be the MAC header of the first MAC frame. Information within the MAC header indicates the length of the first MAC frame and provides a means to find the start of the next MAC frame. Each MAC frame within a concatenation MUST be unique and MAY be of any type. This means that packet and MAC-specific frames MAY be mixed together. However, all frames in a concatenation MUST be assigned to the same service flow. The CMTS MUST support concatenations containing multiple frame types, including both packet and MAC-specific frames.

The embedded MAC frames MAY be addressed to different destinations and MUST be delivered as if they were transmitted individually.

The format of the concatenation MAC header MUST be as shown in Figure 8-11 and Table 8-10.

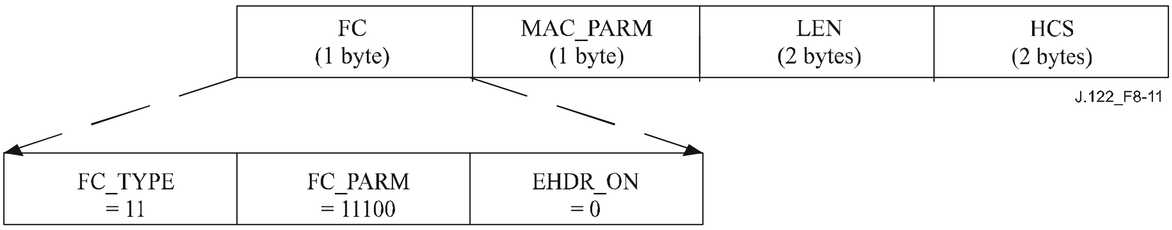

Figure 8-11 – Concatenation MAC header format

Table 8-10 – Concatenated MAC frame format

Field	Usage	Size
FC	FC_TYPE = 11; MAC-specific header FC_PARM[4:0] = 11100; concatenation MAC header EHDR_ON = 0; no EHDR with concatenation header	8 bits
MAC_PARM	CNT, number of MAC frames in this concatenation CNT = 0 indicates unspecified number of MAC frames	8 bits
LEN	LEN = x +...+ y; length of all following MAC frames in bytes	16 bits
EHDR	Extended MAC header MUST NOT be used	0 bytes
HCS	MAC header check sequence	2 bytes
MAC frame 1	First MAC frame: MAC header plus OPTIONAL data PDU	x bytes
MAC frame n	Last MAC frame: MAC header plus OPTIONAL data PDU	y bytes
	Length of concatenated MAC frame	6 + LEN bytes

The MAC_PARM field in the concatenation MAC header provides a count of MAC frames as opposed to EHDR length or REQ amount as used in other MAC headers. If the field is non-zero, then it MUST indicate the total count of MAC frames (CNT) in this concatenation burst.

8.2.6 Extended MAC headers

Every MAC header, except the timing, concatenation MAC header and request frame, has the capability of defining an extended header field (EHDR). The presence of an EHDR field MUST be indicated by the EHDR_ON flag in the FC field being set. Whenever this bit is set, then the MAC_PARM field MUST be used as the EHDR length (ELEN). The minimum defined EHDR is 1 byte. The maximum EHDR length is 240 bytes.

A compliant CMTS and CM MUST support extended headers.

The format of a generic MAC header with an extended header included MUST be as shown in Figure 8-12 and Table 8-11.

NOTE – Extended headers MUST NOT be used in a concatenation MAC header, but MAY be included as part of the MAC headers within the concatenation.

Extended headers MUST NOT be used in request frames and timing MAC headers.

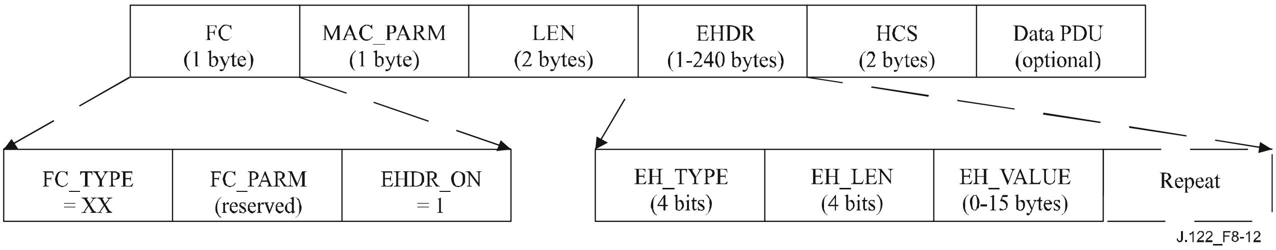

Figure 8-12 – Extended MAC format

Table 8-11 – Example extended header format

Field	Usage	Size
FC	FC_TYPE = XX; applies to all MAC headers FC_PARM[4:0] = XXXXX; dependent on FC_TYPE EHDR_ON = 1; EHDR present this example	8 bits
MAC_PARM	ELEN = x; length of EHDR in bytes	8 bits
LEN	LEN = x + y; length of EHDR plus OPTIONAL data PDU in bytes	16 bits
EHDR	Extended MAC header present this example	x bytes
HCS	MAC header check sequence	2 bytes
PDU	OPTIONAL data PDU	y bytes
	Length of MAC frame with EHDR	6 + x + y bytes

Since the EHDR increases the length of the MAC frame, the LEN field MUST be increased to include both the length of the data PDU and the length of the EHDR.

The EHDR field consists of one or more EH elements. Each EH element is variably sized. The first byte of the EH element MUST contain a type and a length field. Every CM MUST use this length to skip over any unknown EH elements. The format of an EH element MUST be as shown in Table 8-12.

Table 8-12 – EH element format

EH element fields	Usage	Size
EH_TYPE	EH element type field	4 bits
EH_LEN	Length of EH_VALUE	4 bits
EH_VALUE	EH element data	0-15 bytes

The types of EH element defined in Table 8-13 MUST be supported. Reserved and extended types are undefined at this point and MUST be ignored.

Table 8-13 – Extended header types

EH_TYPE	EH_LEN	EH_VALUE
0	0	Null configuration setting; may be used to pad the extended header. The EH_LEN MUST be zero, but the configuration setting may be repeated.
1	3	Request: mini-slots requested (1 byte); SID (2 bytes) [CM → CMTS]
2	2	Acknowledgment requested; SID (2 bytes) [CM → CMTS]
3 (= BP-UP)	4	Upstream privacy EH element [DOCS8]
	5	Upstream privacy with fragmentation[a)] EH element [DOCS8] (See clause 8.2.7)
4 (= BP-DOWN)	4	Downstream privacy EH element [DOCS8]
5	1	Service flow EH element; payload header suppression header downstream.
6	1	Service flow EH element; payload header suppression header upstream
	2	Service flow EH element; payload header suppression header upstream (1 byte), unsolicited grant synchronization header (1 byte)
7-9		Reserved
10-14		Reserved [CM ↔ CM]
15	xx	Extended EH element: EHX_TYPE (1 byte), EHX_LEN (1 byte), EH_VALUE (length determined by EHX_LEN)
a) An upstream privacy with fragmentation EH element MUST only occur within a fragmentation MAC-specific header (refer to clause 8.2.5.4).		

The first ten EH element types are intended for one-way transfer between the cable modem and the CMTS. The next five EH element types are for end-to-end usage within a MAC-sublayer domain. Thus, the information attached to EHDR elements 10-14 on the upstream MUST also be attached when the information is forwarded within a MAC-sublayer domain. The final EH element type is an escape mechanism that allows for more types and longer values, and MUST be as shown in Table 8-13.

8.2.6.1 Piggyback requests

Several extended headers can be used to request bandwidth for subsequent transmissions. These requests are generically referred to as "piggyback requests". They are extremely valuable for performance because they are not subject to contention as request frames generally are (refer to clause 9.4).

Requests for additional bandwidth can be included in request, upstream privacy and upstream privacy with fragmentation extended header elements.

8.2.6.2 Fragmentation extended header

Fragmented packets use a combination of the fragmentation MAC header and a modified version of the upstream privacy extended header. Clause 8.2.5.4 describes the fragmentation MAC header. The upstream privacy extended header with fragmentation, also known as the fragmentation extended header, MUST be as shown in Table 8-14.

Table 8-14 – Fragmentation extended header format

EH Element Fields	Usage	Size
EH_TYPE	Upstream privacy EH element = 3	4 bits
EH_LEN	Length of EH_VALUE = 5	4 bits
EH_VALUE	Key_seq; same as in BP_UP	4 bits
	Ver = 1; version number for this EHDR	4 bits
	BPI_ENABLE If BPI_ENABLE = 0, BPI disabled If BPI_ENABLE = l, BPI enabled	1 bit
	Toggle bit; same as in BP-UP[a)]	1 bit
	SID; service ID associated with this fragment	14 bits
	REQ; number of mini-slots for a piggyback request	8 bits
	Reserved; must be set to zero	2 bits
	First_Frag; set to one for first fragment only	1 bit
	Last_Frag; set to one for last fragment only	1 bit
	Frag_seq; fragment sequence count, incremented for each fragment.	4 bits
a) Refer to [DOCS8].		

8.2.6.3 Service flow extended header

The service flow EH element is used to enhance service flow operations. It may consist of one or two bytes in the EH_VALUE field. The payload header suppression header is the only byte in a one-byte field or the first byte in a two-byte field. The unsolicited grant synchronization header is the second byte in a two-byte field.

8.2.6.3.1 Payload header suppression header

In payload header suppression (PHS), a repetitive portion of the payload headers following the HCS is suppressed by the sending entity and restored by the receiving entity. In the upstream, the sending entity is the CM and the receiving entity is the CMTS. In the downstream, the sending entity is the CMTS and the receiving entity is the CM.

For small payloads, payload header suppression provides increased bandwidth efficiency without having to use compression. Payload header suppression may be separately provisioned in the upstream and downstream, and is referenced with an extended header element.

A compliant CM MUST support payload header suppression[9]. A compliant CMTS MUST support payload header suppression.

The payload header suppression extended header sub-element has the format shown in Table 8-15:

9 This is not intended to imply that the CM must be capable of determining when to invoke payload header suppression. Payload header suppression support is only required for the explicitly signalled case.

Table 8-15 – Payload header suppression EHDR sub-element format

<table>
<tr><th>EH element fields</th><th colspan="2">Usage</th><th>Size</th></tr>
<tr><td>EH_TYPE</td><td colspan="2">Service flow EH_TYPE = 5 for downstream and EH_TYPE = 6 for upstream</td><td>4 bits</td></tr>
<tr><td>EH_LEN</td><td colspan="2">Length of EH_VALUE = 1</td><td>4 bits</td></tr>
<tr><td rowspan="2">EH_VALUE</td><td>0</td><td>Indicates no payload header suppression on current packet</td><td rowspan="2">8 bits</td></tr>
<tr><td>1-255</td><td>Payload header suppression index (PHSI)</td></tr>
</table>

The payload header suppression index is unique per SID in the upstream and unique per CM in the downstream. Payload header suppression is disabled if this extended header element is omitted or, if included, with the PHSI value set to 0. The payload header suppression index (PHSI) references the suppressed byte string known as a payload header suppression field (PHSF).

NOTE – While PHS signalling allows for up to 255 payload header suppression rules per service flow, the exact number of PHS rules supported per service flow is implementation-dependent. Similarly, PHS signalling allows for PHS sizes of up to 255 bytes, however, the maximum PHS size supported is implementation-dependent. For interoperability, the minimum PHS size that MUST be supported is 64 bytes for any PHS rule supported. As with any other parameter requested in a dynamic service request, a PHS-related DSx request can be denied because of a lack of resources.

The upstream suppression field MUST begin with the first byte following the MAC header checksum. The downstream suppression field MUST begin with the thirteenth byte following the MAC header checksum. This allows the Ethernet SA and DA to be available for filtering by the CM.

The operation of baseline privacy (refer to [DOCS8]) is not affected by the use of PHS. When fragmentation is inactive, baseline privacy begins encryption and decryption with the thirteenth byte following the MAC header checksum.

Unless the entire packet PDU is suppressed, the packet PDU CRC is always transmitted, and MUST be calculated only on the bytes transmitted. The bytes that are suppressed MUST NOT be included in the CRC calculation.

8.2.6.3.2 Unsolicited grant synchronization header

The unsolicited grant synchronization header may be used to pass status information regarding service flow scheduling between the CM and CMTS. It is currently only defined for use in the upstream with unsolicited grant and unsolicited grant with activity detection scheduling services (refer to clause 10.2.3).

This extended header is similar to the payload suppression EHDR except that the EH_LEN is 2, and the EH_VALUE has one additional byte which includes information related to unsolicited grant synchronization. For all other service flow scheduling types, the field SHOULD NOT be included in the extended header element generated by the CM. The CMTS MAY ignore this field.

Table 8-16 – Unsolicited grant synchronization EHDR sub-element format

<table>
<tr><th>EH element fields</th><th colspan="2">Usage</th><th>Size</th></tr>
<tr><td>EH_TYPE</td><td colspan="2">Service flow EH_TYPE = 6</td><td>4 bits</td></tr>
<tr><td>EH_LEN</td><td colspan="2">Length of EH_VALUE = 2</td><td>4 bits</td></tr>
<tr><td rowspan="4">EH_VALUE</td><td>0</td><td>Indicates no payload header suppression on current packet.</td><td rowspan="2">8 bits
(always present)</td></tr>
<tr><td>1-255</td><td>Payload header suppression index (PHSI)</td></tr>
<tr><td colspan="2">Queue indicator</td><td>1 bit</td></tr>
<tr><td colspan="2">Active grants</td><td>7 bits</td></tr>
</table>

8.2.7 Fragmented MAC frames

When enabled, fragmentation is initiated any time the grant length is less than the requested length. This normally occurs because the CMTS chooses to grant less than the requested bandwidth.

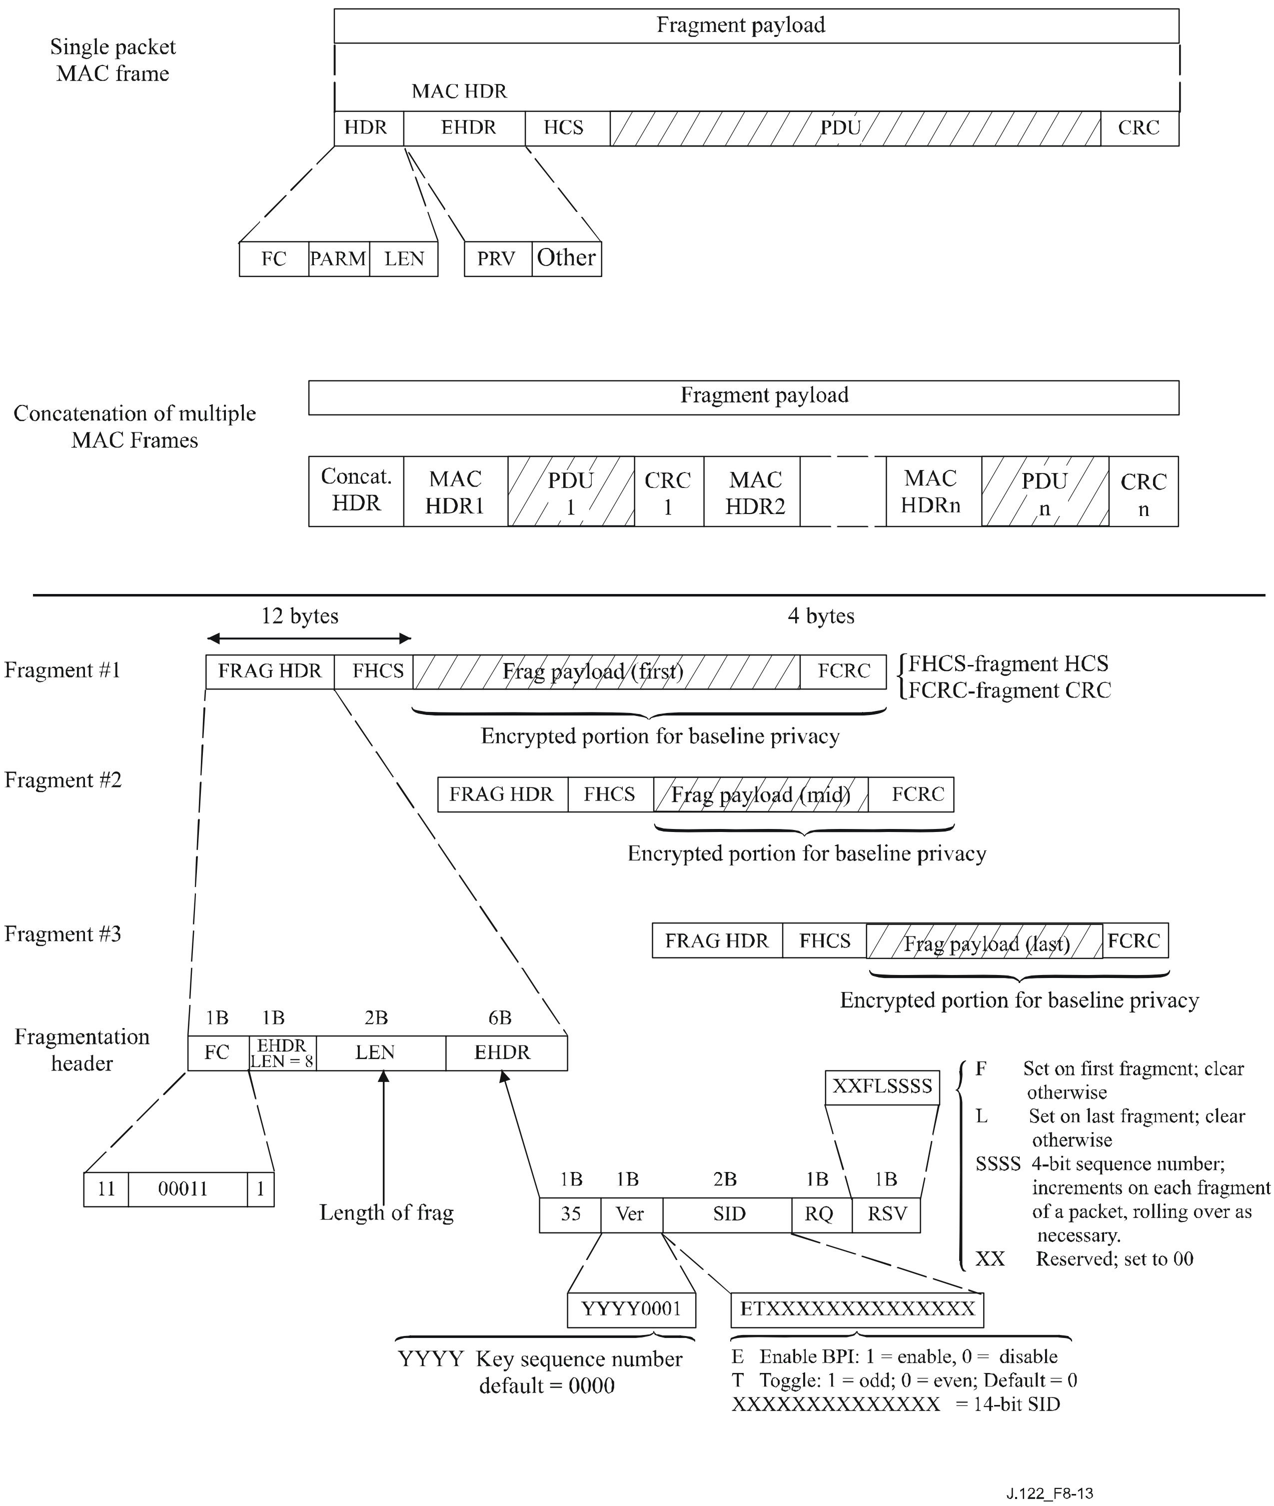

Figure 8-13 – Fragmentation details

The CM MAC calculates how many bytes of the original frame, including overhead for a fragmentation header and CRC, can be sent in the received grant. The CM MAC generates a fragmentation header for each fragment. Fragmented frames use the MAC message type (FC = 11). The FC parameter field is set to (00011), in order to uniquely identify the fragmentation header from other MAC message types. A four bit sequence field is used in the last byte of the extended header field to aid in reassembly and to detect dropped or missing fragments. The CM arbitrarily selects a sequence number for the first fragment of a frame[10]. Once the sequence number is selected for the first fragment, the CM MUST increment the sequence number by one for each fragment transmitted for that frame. There are two flags associated with the sequence number, F and L, where F is set to indicate the first fragment and L is set to indicate the last fragment. Both are cleared for middle fragments. The CMTS stores the sequence number of the first fragment (F bit set) of each frame. The CMTS MUST verify that the fragment sequence field increments (by one) for each fragment of the frame.

The REQ field in the fragmentation header is used by the fragmentation protocol for first and middle fragments (refer to clause 10.3). For the last fragment, the REQ field is interpreted as a request for bandwidth for a subsequent frame.

Fragmentation headers are fixed size and MUST contain only a fragmentation extended header element. The extended header consists of a privacy EH element extended by one byte to make the fragment overhead an even 16 bytes. A privacy EH element is used whether the original packet header contained a privacy EH element or not. If privacy is in use, the following fields, version, enable bit and SID, in the fragment EH element are the same with those of BP EH element inside the original MAC frame. If privacy is not in use, the privacy EH element is used but the enable bit is cleared. The SID used in the fragment EH element MUST match the SID used in the partial grant that initiated the fragmentation. A separate CRC MUST be calculated for each fragment (note that each MAC frame payload will also contain the CRC for that packet). A packet CRC of a reassembled packet MAY be checked by the CMTS even though an FCRC covers each fragment.

The CMTS MUST make certain that any fragmentary grant it makes is large enough to hold at least 17 bytes of MAC layer data. This is to ensure that the grant is large enough to accommodate fragmentation overhead plus at least 1 byte of actual data. The CMTS may want to enforce an even higher limit as small fragments are extremely inefficient.

When fragmentation is active, baseline privacy encryption and decryption begin with the first byte following the MAC header checksum.

8.2.7.1 Considerations for concatenated packets and fragmentation

MAC management messages and data PDUs can occur in the same concatenated frame. Without fragmentation, the MAC management messages within a concatenated frame would be unencrypted. However, with fragmentation enabled on the concatenated frame, the entire concatenated frame is encrypted based on the privacy extended header element. This allows baseline privacy to encrypt each fragment without examining its contents. Clearly, this only applies when baseline privacy is enabled.

To ensure encryption synchronization, if fragmentation, concatenation and baseline privacy are all enabled, a CM MUST NOT concatenate BPKM MAC management messages. This ensures that BPKM MAC management messages are always sent unencrypted.

8.2.8 Error-handling

The cable network is a potentially harsh environment that can cause several different error conditions to occur. This clause, together with clause 11.5, describes the procedures that are required when an exception occurs at the MAC framing level.

The most obvious type of error occurs when the HCS on the MAC header fails. This can be a result of noise on the network or possibly by collisions in the upstream channel. Framing recovery on the downstream channel is performed by the MPEG transmission convergence sublayer. In the upstream channel, framing is recovered on each transmitted burst, such that framing on one burst is

[10] 'Frame' always refers to either frames with a single packet PDU or concatenated frame.

independent of framing on prior bursts. Hence, framing errors within a burst are handled by simply ignoring that burst; i.e., errors are unrecoverable until the next burst.

A second exception, which applies only to the upstream, occurs when the length field is corrupted and the MAC thinks the frame is longer or shorter than it actually is. Synchronization will recover at the next valid upstream data interval.

For every MAC transmission, The HCS MUST be verified. When a bad HCS is detected, the MAC header and any payload MUST be dropped.

For packet PDU transmissions, a bad CRC may be detected. Since the CRC only covers the data PDU and the HCS covers the MAC Header; the MAC header is still considered valid. Thus, the packet PDU MUST be dropped, but any pertinent information in the MAC header (e.g., bandwidth request information) MAY be used.

8.2.8.1 Error recovery during fragmentation

There are some special error handling considerations for fragmentation. Each fragment has its own fragmentation header complete with an HCS and its own FCRC. There may be other MAC headers and CRCs within the fragmented payload. However, only the HCS of the fragment header and the FCRC are used for error detection during fragment reassembly.

If the HCS for a fragment fails, the CMTS MUST discard that fragment. If the HCS passes but the FCRC fails, the CMTS MUST discard that fragment, but MAY process any requests in the fragment header. The CMTS SHOULD process such a request if it is performing fragmentation in piggyback mode (refer to clause 10.3.2.2). This allows the remainder of the frame to be transmitted as quickly as possible.

If a CMTS is performing fragmentation in multiple grant mode (refer to clause 10.3.2.1) it SHOULD complete all the grants necessary to fulfil the CM's original request even if a fragment is lost or discarded. This allows the remainder of the frame to be transmitted as quickly as possible.

If any fragment of a non-concatenated MAC frame is lost or discarded, the CMTS MUST discard the rest of that frame. If a fragment of a concatenated MAC frame is lost or discarded, the CMTS MAY forward any frames within the concatenation that have been received correctly or it MAY discard all the frames in the concatenation.

A CMTS MUST terminate fragment reassembly if any of the following occurs for any fragment on a given SID:

- The CMTS receives a fragment with the L-bit set.
- The CMTS receives an upstream fragment, other than the first one, with the F-bit set.
- The CMTS receives a packet PDU frame with no fragmentation header.
- The CMTS deletes the SID for any reason.

In addition, the CMTS MAY terminate fragment reassembly based on implementation-dependent criteria such as a reassembly timer. When a CMTS terminates fragment reassembly, it MUST dispose of (either by discarding or forwarding) the reassembled frame(s).

8.2.8.2 Error codes and messages

Annex D of [DOCS5] lists CM and CMTS error codes and messages. When reporting error conditions, these codes MUST be used as indicated in [DOCS5] and MAY be used for reporting errors via vendor-specific interfaces. If the error codes are used, the error messages MAY be replaced by other descriptive messages.

8.3 MAC management messages

8.3.1 MAC management message header

MAC management messages MUST be encapsulated in an LLC unnumbered information frame per [ISO/IEC 8802-2], which in turn is encapsulated within the cable network MAC framing, as shown in Figure 8-14. Figure 8-14 shows the MAC header and the MAC management message header fields which are common across all MAC management messages.

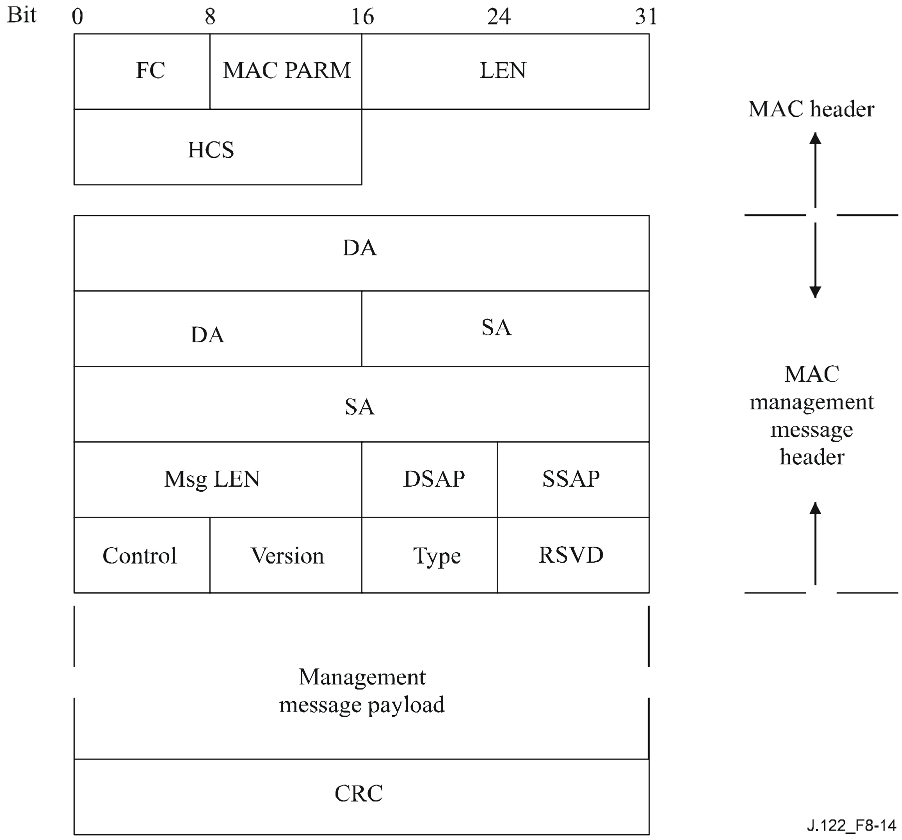

Figure 8-14 – MAC header and MAC management message header fields

The fields MUST be as defined below:

- **FC, MAC PARM, LEN, HCS**: Common MAC frame header (refer to clause 8.2.1.4 for details). All messages use a MAC-specific header.
- **Destination address (DA)**: MAC management frames will be addressed to a specific CM unicast address or to the DOCS management multicast address. These DOCS MAC management addresses are described in Annex A.
- **Source address (SA)**: The MAC address of the source CM or CMTS system.
- **Msg LEN**: Length of the MAC message from DSAP to the end of the payload.
- **DSAP**: The LLC null destination SAP (00) as defined by [ISO/IEC 8802-2].
- **SSAP**: The LLC null source SAP (00) as defined by [ISO/IEC 8802-2].
- **Control**: Unnumbered information frame (03) as defined by [ISO/IEC 8802-2].
- **Version and Type**: Each 1 octet. Refer to Table 8-17. Messages with a version number of 1 are understood by all CMs and CMTSs compliant with all versions of the DOCS Recommendation. Messages with a version number of 2 are understood by DOCS 1.1 and 2.0 equipment, and messages with a version number of 3 are understood by DOCS 2.0

equipment. DOCS 2.0-compliant CMs and CMTSs MUST silently discard any message with a version number greater than 3.

Table 8-17 – MAC management message types

Type value	Version	Message name	Message description
1	1	SYNC	Timing synchronization
2 or 29	1 or 3	UCD	Upstream channel descriptor A UCD for a DOCS 2.0-only channel uses a type of 29 and a version of 3. All other UCDs use a type of 2 and a version of 1 (see clause 8.3.3).
3	1	MAP	Upstream bandwidth allocation
4	1	RNG-REQ	Ranging request
5	1	RNG-RSP	Ranging response
6	1	REG-REQ	Registration request
7	1	REG-RSP	Registration response
8	1	UCC-REQ	Upstream channel change request
9	1	UCC-RSP	Upstream channel change response
10	1	TRI-TCD	Telephony channel descriptor – [DOCS6]
11	1	TRI-TSI	Termination system information – [DOCS6]
12	1	BPKM-REQ	Privacy key management request [DOCS8]
13	1	BPKM-RSP	Privacy key management response [DOCS8]
14	2	REG-ACK	Registration acknowledge
15	2	DSA-REQ	Dynamic service addition request
16	2	DSA-RSP	Dynamic service addition response
17	2	DSA-ACK	Dynamic service addition acknowledge
18	2	DSC-REQ	Dynamic service change request
19	2	DSC-RSP	Dynamic service change response
20	2	DSC-ACK	Dynamic service change acknowledge
21	2	DSD-REQ	Dynamic service deletion request
22	2	DSD-RSP	Dynamic service deletion response
23	2	DCC-REQ	Dynamic channel change request
24	2	DCC-RSP	Dynamic channel change response
25	2	DCC-ACK	Dynamic channel change acknowledge
26	2	DCI-REQ	Device class identification request
27	2	DCI-RSP	Device class identification response
28	2	UP-DIS	Upstream transmitter disable
29	3		(See entry for UCD)
30	3	INIT-RNG-REQ	Initial ranging request
31	3	TST-REQ	Test request message
32	3	DCD	Downstream channel descriptor (DSG)
33-255			Reserved for future use

- **RSVD**: 1 octet. This field is used to align the message payload on a 32-bit boundary. Set to 0 for this version for all messages other than the RNG-REQ and INIT-RNG-REQ. See clauses 8.3.5 and 8.3.26 for setting this value for these two messages.
- **Management message payload**: Variable length. As defined for each specific management message.
- **CRC**: Covers message including header fields (DA, SA, ...). Polynomial defined by [ISO/IEC 8802-3].

A compliant CMTS or CM MUST support the MAC management message types listed in Table 8-17, except messages specific to telephony return devices which MAY be supported.

8.3.2 Time synchronization (SYNC)

Time synchronization (SYNC) MUST be transmitted by CMTS at a periodic interval to establish MAC sublayer timing. This message MUST use an FC field with FC_TYPE = MAC-specific header and FC_PARM = timing MAC header. This MUST be followed by a packet PDU in the format shown in Figure 8-15.

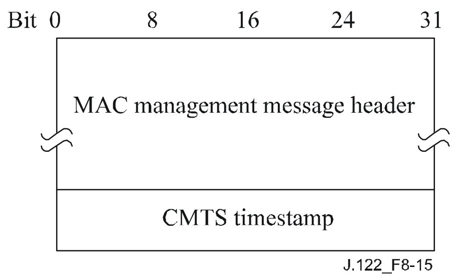

Figure 8-15 – Format of packet PDU following the timing header

The parameters shall be as defined below:

- **CMTS timestamp**: The count state of an incrementing 32-bit binary counter clocked with the CMTS 10.24-MHz master clock.

  The CMTS timestamp represents the count state at the instant that the first byte (or a fixed time offset from the first byte) of the time synchronization MAC management message is transferred from the downstream transmission convergence sublayer to the downstream physical media dependent sublayer as described in clause 6.3.7. The CMTS MUST NOT allow a SYNC message to cross an MPEG packet boundary[11].

8.3.3 Upstream channel descriptor (UCD)

An upstream channel descriptor MUST be transmitted by the CMTS at a periodic interval to define the characteristics of an logical upstream channel (see Figure 8-16). A separate message MUST be transmitted for each logical upstream that is currently available for use.

[11] Since the SYNC message applies to all upstream channels within this MAC domain, units were chosen to be independent of the modulation rate of any particular upstream channel. See clause 9.3.4 for time-unit relationships.

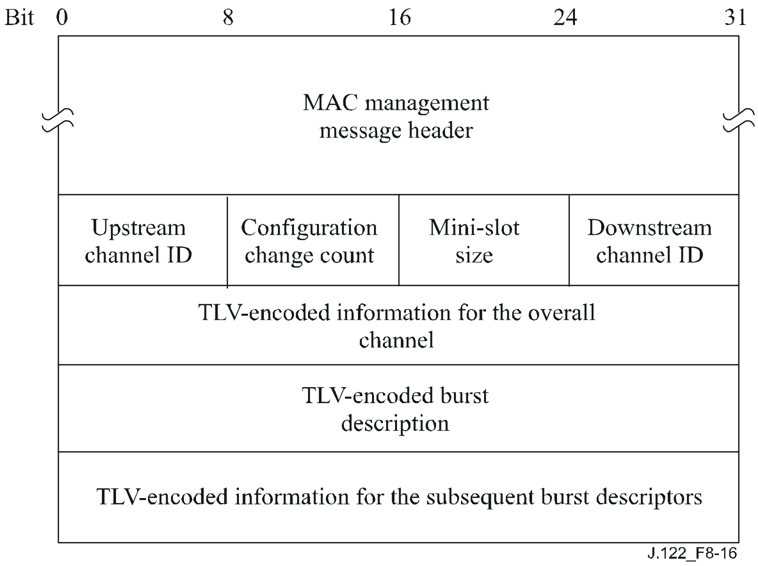

Figure 8-16 – Upstream channel descriptor

The MAC management header for this message has 2 possible values for the type field and for the version field. For DOCS 2.0-only upstreams the CMTS MUST use a value of 29 for the type field and MUST use a value of 3 for the version field. For all other logical upstreams, the CMTS MUST use a value of 2 for the type field and a value of 1 for the version field. A CMTS MUST NOT use type 5 TLVs to encode IUCs 1-6 in a UCD with a message type of 2. A CMTS MUST use type 5 TLVs to encode all burst profiles in a UCD with a message type of 29.

A type 2 UCD MUST contain type 4 burst descriptors for IUCs 3 and 4, a type 4 burst descriptor for requests and a type 4 burst descriptor for data. The burst descriptors for requests and data MUST use 1.x TDMA. To describe a 1.x/2.0 mixed logical channel, the UCD MUST additionally contain a type 5 burst descriptor for 2.0 TDMA data opportunities.

A type 29 UCD MUST contain a type 5 burst descriptor for IUC 3, a type 5 burst descriptor for requests, and a type 5 burst descriptor for data.

A CMTS MUST NOT include burst descriptors for IUCs 5 or 6 in a UCD message for a DOCS 2.0-only upstream.

For interoperability, a CMTS SHOULD provide:

- burst descriptors for IUCs 1, 5 and 6 in a type 2 UCD describing a 1.x-only channel;
- burst descriptors for IUCs 1, 5, 6, 9 and 10 in a type 2 UCD describing a 1.x/2.0 mixed logical channel;
- burst descriptors for IUCs 1, 9 and 10 in a type 29 UCD.

The CMTS MUST treat an upstream as a DOCS 2.0-only upstream if any of the following is true about the channel wide parameters: S-CDMA mode is enabled, the mini-slot size is 1 time tick, or the value of the modulation rate parameter is 32. The CMTS MUST treat an upstream as a DOCS 2.0-only upstream if any of the following is true about any of IUCs 1-4: A modulation type other than QPSK or 16QAM is used, the FEC error correction (T) parameter is greater than 10, any portion of the extended preamble is used, any attribute from Table 8-19, "upstream physical-layer burst attributes," with a type greater than 11 is present in the descriptor. Note that none of these conditions can ever be true for IUCs 5 or 6.

To provide for flexibility, the message parameters following the channel ID MUST be encoded in a type/length/value (TLV) form in which the type and length fields are each 1 octet long.

A CMTS MUST generate UCDs in the format shown in Figure 8-16, including all of the following parameters:

- **Configuration change count**: Incremented by one (modulo the field size) by the CMTS whenever any of the values of this channel descriptor change, excluding the S-CDMA snapshot TLV[12]. If the value of this count in a subsequent UCD remains the same, the CM can quickly decide that the channel operating parameters have not changed, and may be able to disregard the remainder of the message. This value is also referenced from the MAP.
- **Mini-slot size**: The size T of the mini-slot for this upstream channel in units of the timebase tick of 6.25 µs. For channels that can support DOCS 1.x CMs, the allowable values are $T = 2^M$, M = 1,...,7. That is, T = 2, 4, 8, 16, 32, 64 or 128.

  For DOCS 2.0-only channels, the relationship between M and T remains the same, but the allowable values are M = 0,1,…,7, with T = 1, 2, 4, 8, 16, 32, 64 or 128. If the value of T is 1 then the channel MUST be treated as a DOCS 2.0-only channel. On S-CDMA channels, this parameter will not have any effect.
- **Upstream channel ID**: The identifier of the upstream channel to which this message refers. This identifier is arbitrarily chosen by the CMTS and is only unique within the MAC-sublayer domain.

  NOTE – Upstream channel ID = 0 is reserved to indicate telephony return.
- **Downstream channel ID**: The identifier of the downstream channel on which this message has been transmitted. This identifier is arbitrarily chosen by the CMTS at startup, and is only unique within the MAC-sublayer domain.

All other parameters are coded as TLV tuples. The type values used MUST be those defined in Table 8-18, for channel parameters, and Table 8-19, for upstream physical layer burst attributes. Burst descriptors (type 4 and/or type 5) MUST appear in the UCD message after all other channel-wide parameters.

Table 8-18 – Channel TLV parameters

Name	Type (1 byte)	Length (1 byte)	Value (variable length)
Modulation rate	1	1	Multiples of base rate of 160 kHz (value is 1, 2, 4, 8,16, or 32). A value of 32 means that this is a DOCS 2.0-only upstream. If S-CDMA mode is enabled, then this parameter MUST have a value of 8, 16 or 32.
Frequency	2	4	Upstream centre frequency (Hz)

[12] Refer to clause 6.2.11.2 for a description of the timestamp snapshot. The periodic update of the snapshot association does not represent a change in the operating parameters of the channel, hence the UCD configuration change count will not be incremented.

Table 8-18 – Channel TLV parameters

Name	Type (1 byte)	Length (1 byte)	Value (variable length)
Preamble pattern	3	1-128	The value field defines the first portion of the preamble superstring. If there is no extended preamble pattern parameter, then this parameter defines the entire preamble superstring. All burst-specific preamble values are chosen as bit-substrings of the preamble superstring. The first byte of the value field contains the first 8 bits of the superstring, with the first bit of the preamble superstring in the MSB position of the first value field byte, the eighth bit of the preamble superstring in the LSB position of the first value field byte; the second byte in the value field contains the second eight bits of the superstring, with the ninth bit of the superstring in the MSB of the second byte and sixteenth bit of the preamble superstring in the LSB of the second byte, and so forth.
Burst descriptor (DOCS 1.x)	4	n	May appear more than once; described below.
Burst descriptor (DOCS 2.0)	5	n	May appear more than once; described below.
Extended preamble pattern	6	1-64	512-bit preamble superstring extension. The value field is concatenated to the end of the value field of the preamble pattern to complete the preamble superstring. This parameter MUST NOT be included unless the length of the preamble pattern parameter is 128 bytes. Therefore, the MSB of the first byte of the value field of this parameter always follows the LSB of the 128th byte of the value field of the preamble pattern parameter in the preamble superstring.
S-CDMA mode enable[a)]	7	1	1 = on; 2 = off. If parameter is on, the upstream will operate in S-CDMA mode. Otherwise, it operates in TDMA mode. If this parameter is set to on, this is a DOCS 2.0-only upstream.
S-CDMA spreading intervals per frame	8	1	Number of consecutive spreading intervals mapped onto a two-dimensional frame (value is 1 through 32). This TLV MUST be present if S-CDMA mode is enabled, and MUST NOT be present if it is not.
S-CDMA codes per mini-slot	9	1	Number of consecutive codes mapped into a two-dimensional mini-slot (value is 2 through 32). This TLV MUST be present if S-CDMA mode is enabled, and MUST NOT be present if it is not.
S-CDMA number of active codes	10	1	Number of codes available to carry data payload (value is 64 through 128). This value MUST be a multiple of codes per mini-slot (TLV type 9). This TLV MUST be present if S-CDMA Mode is enabled, and MUST NOT be present if it is not.
S-CDMA code-hopping seed	11	2	15-bit seed to initialize code-hopping sequence. The value is left-justified in the 2-byte field. Set seed = 0 to disable code hopping. This TLV MUST be present if S-CDMA Mode is enabled, and MUST NOT be present if it is not.

Table 8-18 – Channel TLV parameters

Name	Type (1 byte)	Length (1 byte)	Value (variable length)
S-CDMA US ratio numerator "M"	12	2	The numerator (M) of the M/N ratio relating the downstream symbol clock to the upstream modulation clock. This TLV MUST be present if S-CDMA mode is enabled, and MUST NOT be present if it is not.
S-CDMA US ratio denominator "N"	13	2	The denominator (N) of the M/N ratio relating the downstream symbol clock to the upstream modulation clock. This TLV MUST be present if S-CDMA mode is enabled, and MUST NOT be present if it is not.
S-CDMA timestamp snapshot[b)]	14	9	Snapshot of the timestamp, mini-slot and S-CDMA frame taken at an S-CDMA frame boundary at the CMTS. A new value MUST be sampled and sent with each UCD message. Refer to clause 6.2.11.2, "mini-slot numbering". This TLV MUST be present if S-CDMA mode is enabled, and MUST NOT be present if it is not.
Maintain power spectral density	15	1	1 = on; 2 = off. If this value is on and the modulation rate is different from the previous UCD, the CM MUST change its transmit power level to keep the power spectral density as close as possible to what it was prior to the modulation rate change. If this value is off or this parameter is omitted then the CM maintains the same power level that it was using prior to the modulation rate change. In any case the effect of this parameter only lasts until the CM receives a power adjustment in a RNG-RSP.
Ranging required	16	1	0 = no ranging required; 1 = unicast initial ranging required; 2 = broadcast initial ranging required. If this value is non-zero and the UCD change count does not match the UCD currently in effect, the CM MUST perform ranging as specified by this TLV before using any other transmit opportunities with the new UCD parameters. If ranging is required, and the CM is already registered, then it MUST maintain its SIDs and not re-register. If this value is 0 or this TLV is omitted, no ranging is required.
S-CDMA maximum scheduled codes enabled	17	1	1 = maximum scheduled codes is enabled 2 = maximum scheduled codes is disabled. CMs that implement the S-CDMA maximum scheduled codes MUST set the RSVD field in the ranging requests as described in clauses 8.3.5 and 8.3.26. If S-CDMA mode is disabled on this channel, this TLV MUST NOT be present.
Ranging hold-off priority field	18	4	Bit field with values representing device classes, as defined in clause C.1.3.1.16 that should temporarily inhibit initial ranging. The CMTS MAY include this TLV in the UCD message. The CM MUST observe this TLV, as described in clause 11.2.2.

Table 8-18 – Channel TLV parameters

Name	Type (1 byte)	Length (1 byte)	Value (variable length)
Channel class ID	19	4	Bit field with values representing device classes as defined in clause C.1.3.1.16 that are allowed to use the channel. The CMTS MAY include this TLV in the UCD message. The CM MUST observe this TLV, as described in clause 11.2.2.

[a] CM MUST assume S-CDMA mode is off if TLV is not present.

[b] Refer to clause 6.2.11.2, "mini-slot numbering", for a description of the timestamp snapshot. A change solely in this parameter for a particular UCD does not represent a change in overall channel operating parameters, hence the UCD channel change count will not be incremented.

Burst descriptors are composed of an upstream interval usage code, followed by TLV encodings that define, for each type of upstream usage interval, the physical-layer characteristics that are to be used during that interval. The upstream interval usage codes are defined in the MAP message (see clause 8.3.4 and Table 8-20). The format of the burst descriptors is shown in Figure 8-17.

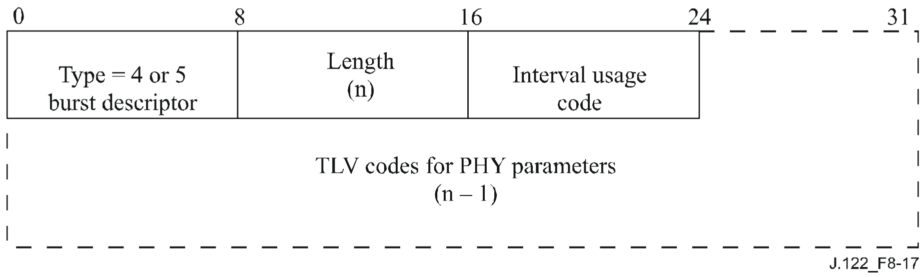

Figure 8-17 – Top-level encoding for burst descriptors

In Figure 8-17:

– **Type**: 4 for burst descriptors intended for DOCS 1.x and/or DOCS 2.0 modems. 5 is for burst descriptors intended for DOCS 2.0 modems only.

– **Length**: The number of bytes in the overall object, including the IUC and the embedded TLV items.

– **IUC**: Interval usage code defined in Table 8-20. The IUC is coded on the 4 least-significant bits. The four most-significant bits are unused (=0).

– **TLV items**: TLV parameters described in Table 8-18, "channel TLV parameters".

Two different type values are used to describe burst descriptors. Type 4 burst descriptors MUST be understood by all modems and MUST only be used to describe IUCs 1 through 6 from Table 8-20. Type 5 burst descriptors MUST be understood by DOCS 2.0 modems. A type 5 burst descriptor MUST be used to describe any IUC if any of the following is true: a modulation type other than QPSK or 16QAM is used, the FEC error correction (T) attribute is greater than 10, any portion of the extended preamble is used, or any attribute from Table 8-19, "upstream physical-layer burst attributes", with a type greater than 11 is present in the descriptor. Type 5 burst descriptors MUST NOT be used to describe IUC 5 or IUC 6. Type 5 burst descriptors MUST be used to describe IUCs 9-11.

A burst descriptor MUST be included for each interval usage code that is to be used in the allocation MAP. The interval usage code MUST be one of the values from Table 8-20.

Within each burst descriptor is an unordered list of physical-layer attributes, encoded as TLV values. These attributes are shown in Table 8-19. The CMTS MUST ensure that the set of burst attributes for all the burst descriptors in the UCD allow any CM on the upstream to be able to request enough mini-slots to be able to transmit a maximum size PDU (see clause 8.2.2, "packet-based MAC Frames").

Table 8-19 – Upstream physical-layer burst attributes

Name	Type (1 byte)	Length (1 byte)	Value (variable length)
Modulation type	1	1	1 = QPSK 2 = 16QAM 3 = 8QAM 4 = 32QAM 5 = 64QAM 6 = 128QAM (S-CDMA only) Values greater than 2 MUST NOT be used in a descriptor encoded in a type 4 TLV.
Differential encoding	2	1	1 = on, 2 = off (see clause 6.2.13, "symbol mapping".)
Preamble length	3	2	Up to 1536 bits for burst descriptors encoded in a type 5 TLV. Up to 1024 bits for descriptors encoded in a type 4 TLV. If this descriptor is encoded in a type 4 TLV then the substring of the preamble superstring defined by this parameter and the preamble value offset MUST NOT include any bits from the extended preamble pattern. The value must be an integral number of symbols (see clause 6.2.9, "preamble prepend").
Preamble value offset	4	2	Identifies the bits to be used in the preamble. This is specified as a starting offset into the preamble superstring (see Table 8-18). That is, a value of zero means that the first bit of the preamble for this burst type is the value of the first bit of the preamble superstring. A value of 100 means that the preamble is to use the 101st and succeeding bits from the preamble superstring. This value must be a multiple of the symbol size. The first bit of the preamble is the first bit into the symbol mapper (see Figures 6-2 and 6-3), and is I_1 in the first symbol of the burst (see clause 6.2.13).
FEC error correction (T)	5	1	0-16 for descriptors encoded in a type 5 TLV. 0-10 for descriptors encoded in a type 4 TLV. (0 implies no FEC. The number of codeword parity bytes is 2 × T)
FEC codeword information bytes (k)	6	1	Fixed: 16 to 253 (assuming FEC on) Shortened: 16 to 253 (assuming FEC on) (Not used if no FEC, T = 0)
Scrambler seed	7	2	The 15-bit seed value left justified in the 2-byte field. Bit 15 is the MSB of the first byte and the LSB of the second byte is not used (not used if scrambler is off).

Table 8-19 – Upstream physical-layer burst attributes

Name	Type (1 byte)	Length (1 byte)	Value (variable length)
Maximum burst size	8	1	The maximum number of mini-slots that can be transmitted during this burst type. Absence of this configuration setting implies that the burst size is limited elsewhere. When the interval type is short data grant (IUC 5) or advanced PHY short data grant (IUC 9), this value MUST be present and greater than zero (see clause 9.1.2.5). If the CMTS needs to limit the maximum length of concatenated frames it SHOULD use this configuration setting to do so.
Guard time size	9	1	For TDMA channels, the number of modulation intervals 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. In type 4 burst descriptors, the CMTS MUST choose the parameters such that the number of bytes that fit into any valid number of mini-slots will not change if the guard time is increased by 1. For S-CDMA channels, there is no guard time, and hence the CM MUST ignore this value. This TLV MUST NOT be present for S-CDMA channels.
Last codeword length	10	1	1 = fixed; 2 = shortened
Scrambler on/off	11	1	1 = on; 2 = off
R-S interleaver depth (I_r)	12	1	Reed-Solomon block interleaving depth. A depth of 0 indicates dynamic mode; a depth of 1 indicates R-S interleaving disabled (see clause 6.2.6) (0 through $\lfloor 2048/(K + 2T) \rfloor$). This TLV MUST be present for burst descriptors encoded in type 5 TLVs on DOCS 2.0 TDMA channels. This TLV MUST NOT be present for S-CDMA channels or in descriptors encoded in a type 4 TLV.
R-S interleaver block size (B_r)	13	2	Reed-Solomon block interleaving size in dynamic mode (18 through 2048). This TLV MUST be present in burst descriptors encoded in type 5 TLVs for DOCS 2.0 TDMA channels. This TLV MUST NOT be present on S-CDMA channels or in descriptors encoded in a type 4 TLV.
Preamble type	14	1	1 = QPSK0 2 = QPSK1 (Refer to Figure 6-18 and to clause 6.2.9) This TLV MUST NOT be present in descriptors encoded in a type 4 TLV.
S-CDMA spreader on/off	15	1	1 = on; 2 = off. This TLV MUST be present for S-CDMA channels. This TLV MUST NOT be present for non-S-CDMA channels or in descriptors encoded in a type 4 TLV.
S-CDMA codes per subframe	16	1	Number of codes per sub-frame used in the S-CDMA framer (1 through 128). This TLV MUST be present for S-CDMA channels. This TLV MUST NOT be present for non-S-CDMA channels or in descriptors encoded in a type 4 TLV.

Table 8-19 – Upstream physical-layer burst attributes

Name	Type (1 byte)	Length (1 byte)	Value (variable length)
S-CDMA framer interleaving step size	17	1	Size of interleaving steps used in S-CDMA framer (1 through 32). This TLV MUST be present for S-CDMA channels. This TLV MUST NOT be present for non-S-CDMA channels or in descriptors encoded in a type 4 TLV.
TCM encoding	18	1	1 = on; 2 = off. This TLV MUST be present for S-CDMA channels. This TLV MUST NOT be present for non-S-CDMA channels or in descriptors encoded in a type 4 TLV.

8.3.3.1 Example of UCD encoded TLV data

An example of UCD encoded TLV data is given in Figure 8-18.

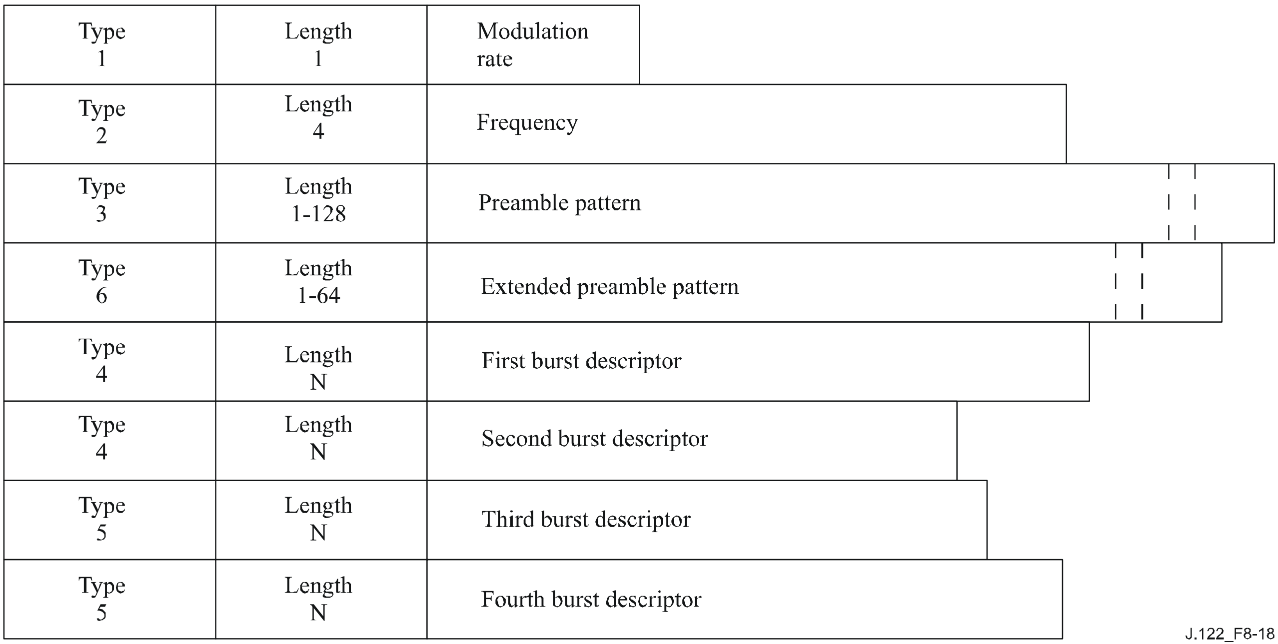

Figure 8-18 – Example of UCD encoded TLV data

8.3.4 Upstream bandwidth allocation map (MAP)

A CMTS MUST generate MAPs in the format shown in Figure 8-19.

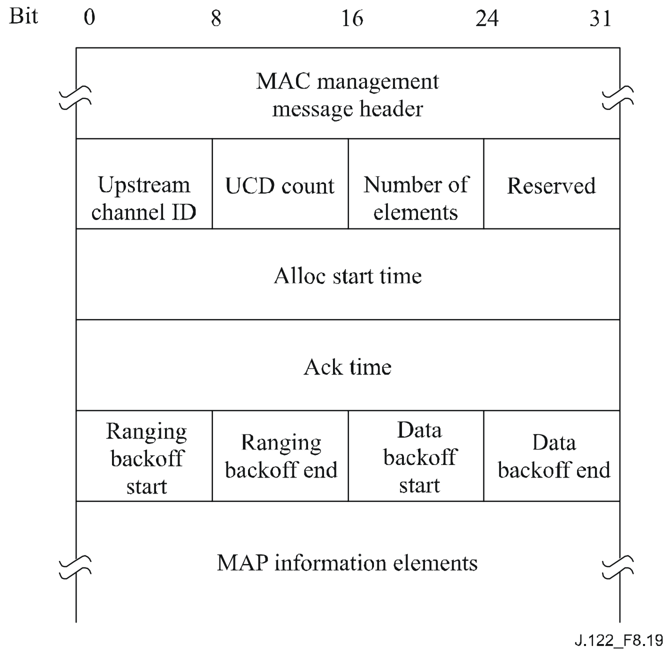

Figure 8-19 – MAP format

The parameters MUST be as follows:

- **Upstream channel ID**: The identifier of the upstream channel to which this message refers.
- **UCD count**: Matches the value of the configuration change count of the UCD which describes the burst parameters which apply to this map. See clause 11.3.2.
- **Number of elements**: Number of information elements in the map.
- **Reserved**: Reserved field for alignment.
- **Alloc start time**: Effective start time from CMTS initialization (in mini-slots) for assignments within this map.
- **Ack time**: Latest time, from CMTS initialization, (mini-slots) processed in upstream. This time is used by the CMs for collision detection purposes. See clause 9.4.
- **Ranging backoff start**: Initial back-off window for initial ranging contention, expressed as a power of two. Values range 0-15 (the highest order bits must be unused and set to 0).
- **Ranging backoff end**: Final back-off window for initial ranging contention, expressed as a power of two. Values range 0-15 (the highest order bits must be unused and set to 0).
- **Data backoff start**: Initial back-off window for contention data and requests, expressed as a power of two. Values range 0-15 (the highest order bits must be unused and set to 0).
- **Data backoff end**: Final back-off window for contention data and requests, expressed as a power of two. Values range 0-15 (the highest order bits must be unused and set to 0).
- **MAP information elements**: MUST be in the format defined in Figure 8-20 and Table 8-20. Values for IUCs are defined in Table 8-20 and are described in detail in clause 9.1.2.

NOTE – Refer to clause 9.1.1 for the relationship between alloc start/ack time and the timebase.

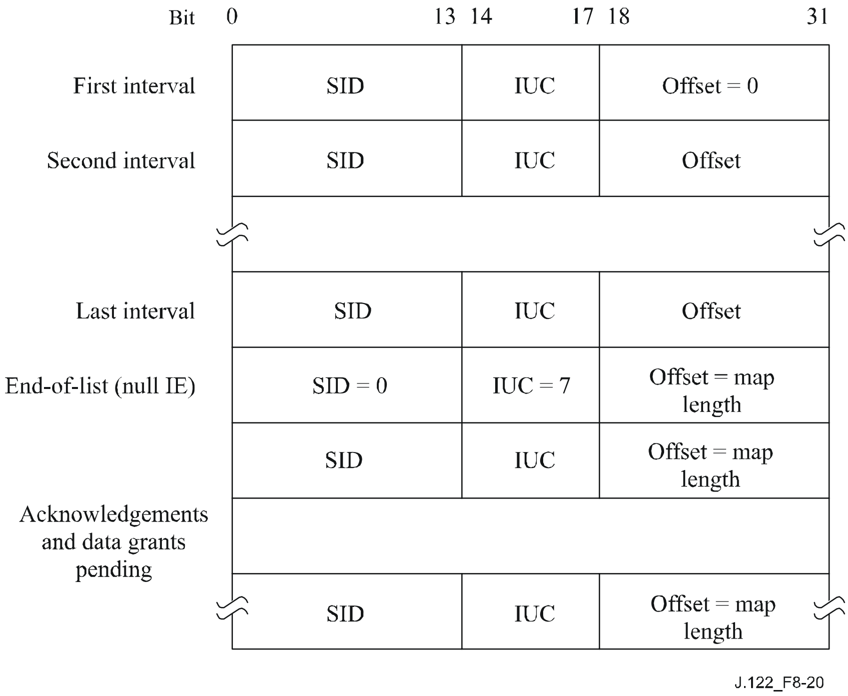

Figure 8-20 – MAP information element structure

Table 8-20 – Allocation MAP information elements (IE)

IE name[a)]	Interval usage code (IUC) (4 bits)	SID (14 bits)	Mini-slot offset (14 bits)
Request	1	Any	Starting offset of REQ region
REQ/data (refer to Annex A, "well-known addresses", for multicast definition)	2	Multicast	Starting offset of IMMEDIATE data region (well-known multicasts define start intervals)
Initial maintenance[b)]	3	Broadcast or unicast	Starting offset of MAINT region (used in initial or periodic ranging)
Station maintenance	4	Unicast[c)]	Starting offset of MAINT region (used in periodic ranging)
Short data grant[d)]	5	Unicast	Starting offset of data grant assignment; If inferred length = 0, then it is a data grant pending
Long data grant	6	Unicast	Starting offset of data grant assignment; If inferred length = 0, then it is a data grant pending
Null IE	7	Zero	Ending offset of the previous grant. Used to bound the length of the last actual interval allocation.
Data ack	8	Unicast	CMTS sets to map length
Advanced PHY short[e)] data grant	9	Unicast	Starting offset of data grant assignment; If inferred length = 0, then it is a data grant pending

Table 8-20 – Allocation MAP information elements (IE)

IE name[a)]	Interval usage code (IUC) (4 bits)	SID (14 bits)	Mini-slot offset (14 bits)
Advanced PHY long data grant	10	Unicast	Starting offset of data grant assignment; If inferred length = 0, then it is a data grant pending
Advanced PHY unsolicited grant	11	Unicast	Starting offset of data grant assignment
Reserved	12-14	Any	Reserved
Expansion	15	Expanded IUC	# of additional 32-bit words in this IE

a) Each IE is a 32-bit quantity, of which the most significant 14 bits represent the SID, the middle 4 bits the IUC, and the low-order 14 bits the mini-slot offset.

b) The CMTS MUST NOT use a unicast SID with an initial maintenance IUC on any upstream that is not a DOCS 2.0-only upstream.

c) The SID used in the station maintenance IE MUST be a temporary SID, or the primary SID that was assigned in the REG-RSP message to a CM.

d) The distinction between long and short data grants is related to the amount of data that can be transmitted in the grant. A short data grant interval may use FEC parameters that are appropriate to short packets while a long data grant may be able to take advantage of greater FEC coding efficiency.

e) The advanced PHY types are provided for channels carrying a combination of DOCS 1.x and DOCS 2.0 bursts and also for channels carrying DOCS 2.0 bursts only.

8.3.5 Ranging request (RNG-REQ)

A ranging request MUST be transmitted by a CM at initialization on an upstream other than a DOCS 2.0-only upstream, and periodically on request from CMTS to determine network delay and request power adjustment. On a DOCS 2.0-only upstream, the CM transmits a INIT-RNG-REQ (see clause 8.3.26) message at initialization instead, but uses the RNG-REQ for all unicast maintenance opportunities provided by the CMTS. This message MUST use an FC_TYPE = MAC-specific header and FC_PARM = timing MAC header. This MUST be followed by a packet PDU in the format shown in Figure 8-21.

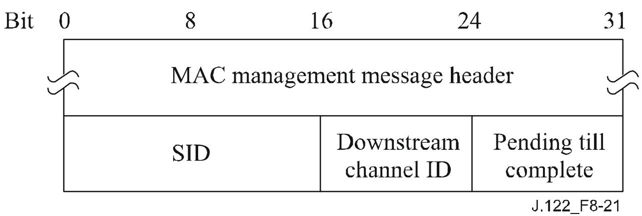

Figure 8-21 – Packet PDU following the timing header

Parameters MUST be as follows:

- **SID**: For RNG-REQ messages transmitted in broadcast initial maintenance intervals:
  – Initialization SID if modem is attempting to join the network.

– Initialization SID if modem has not yet registered and is changing upstream, downstream, or both downstream and upstream channels as directed by a downloaded parameter file.

– Primary SID (previously assigned in REG-RSP) if modem is registered and is changing upstream channels, or if the CM is redoing initial ranging as a result of a DCC, UCC or UCD change (see clauses 8.3.3 and 11.3.2).

For RNG-REQ messages transmitted in unicast initial maintenance or station maintenance intervals:

– Temporary SID if modem has not yet registered.

– Primary SID (previously assigned in REG-RSP) if modem is registered or is redoing initial ranging as a result of DCC, UCC or UCD change.

This is a 16-bit field of which the lower 14 bits define the SID, with bits 14 and 15 defined to be 0.

- **Downstream channel ID**: The identifier of the downstream channel on which the CM received the UCD which described this upstream. This is an 8-bit field.
- **Pending till complete**: If zero, then all previous ranging response attributes have been applied prior to transmitting this request. If non-zero, then this is time estimated to be needed to complete assimilation of ranging parameters. Note that only equalization can be deferred. Units are in unsigned centiseconds (10 ms).

A CM MUST set the RSVD field of the MAC management message header to report support of the S-CDMA maximum scheduled codes if and only if the CMTS indicated that it supports the S-CDMA maximum scheduled codes from TLV-17 of the UCD for the upstream channel on which the CM is ranging. In this case, the CM MUST report the maximum ratio of number of active codes to maximum scheduled codes that the CM can support; this ratio is indicated by a bit mask in the reserved field as shown below. For example, if the number of active codes on the channel is 128 and the CM supports a minimum of 64 scheduled codes (the minimum number of allowed active codes), the CM would report a ratio of 2. The CMTS will use this value in calculating an appropriate value for maximum scheduled codes to assign to the CM. The CM SHOULD support a maximum ratio of 32.

When the CM reports support for the S-CDMA maximum scheduled codes, the CM MUST also report its current transmit power shortfall (in dB). The CM power shortfall is the difference between the current target transmit power of the ranging request and the maximum S-CDMA spreader-on transmit power of 53 dBmV. The CM MUST report a power shortfall of 0 if the current target transmit power of the ranging request is less than or equal to 53 dBmV. This value will be used by the CMTS for calculating appropriate values for S-CDMA maximum scheduled codes and S-CDMA power headroom for the CM.

The format of the RSVD field is:

Bit 7: 1= S-CDMA maximum scheduled codes supported	Bits 6 to 5: CM maximum ratio of: <Anchor1> 00 = 20 1 = 8 10 = 16 11 = 32	Bit 4 to 0: CM power shortfall (1/4 dB)

8.3.6 Ranging response (RNG-RSP)

A ranging response MUST be transmitted by a CMTS in response to received RNG-REQ or INIT-RNG-REQ. The state machines describing the ranging procedure appear in clause 11.2.4. In that procedure it may be noted that, from the point of view of the CM, reception of a ranging response is stateless. In particular, the CM MUST be prepared to receive a ranging response at any time, not just following a ranging request.

To provide for flexibility, the message parameters following the upstream channel ID MUST be encoded in a type/length/value (TLV) form.

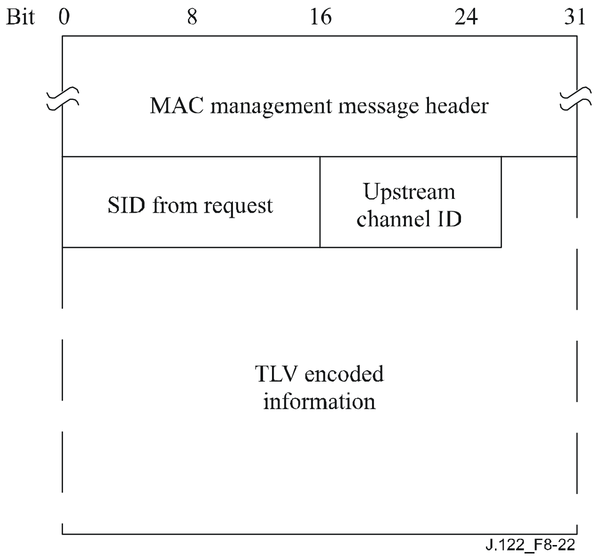

Figure 8-22 – Ranging response

A CMTS MUST generate ranging responses in the form shown in Figure 8-22, including all of the following parameters:

- **SID**: If the modem is being instructed by this response to move to a different channel, this is initialization SID. Otherwise, this is the SID from the corresponding RNG-REQ to which this response refers, except that if the corresponding RNG-REQ was an initial ranging request specifying an initialization SID, then this is the assigned temporary SID.
- **Upstream channel ID**: The identifier of the upstream channel on which the CMTS received the RNG-REQ or INIT-RNG-REQ to which this response refers. On the first ranging response received by the CM during initial ranging, this channel ID may be different from the channel ID the CM used to transmit the range request (see Appendix III). Thus, the CM MUST use this channel ID for the rest of its transactions, not the channel ID from which it initiated the range request.

All other parameters are coded as TLV tuples:

- **Ranging status**: Used to indicate whether upstream messages are received within acceptable limits by CMTS.
- **Timing adjust information**: The amount by which to change the ranging offset of the burst transmission so that bursts arrive at the expected mini-slot time at the CMTS. The units are (1/10.24 MHz) = 97.65625 ns. A negative value implies the ranging offset is to be decreased, resulting in later times of transmission at the CM (see Table 6-8 and clause 6.2.19.1, "ranging offset").
- **Power adjust information**: Specifies the relative change in transmission power level that the CM is to make in order that transmissions arrive at the CMTS at the desired power.

- **Frequency adjust information**: Specifies the relative change in transmission frequency that the CM is to make in order to better match the CMTS (this is fine-frequency adjustment within a channel, not re-assignment to a different channel).
- **CM transmitter equalization information**: This provides the equalization coefficients for the pre-equalizer.
- **Downstream frequency override**: An optional parameter. The downstream frequency with which the modem should redo initial ranging (see clause 8.3.6.3).
- **Upstream channel ID override**: An optional parameter. The identifier of the upstream channel with which the modem should redo initial ranging (see clause 8.3.6.3).
- **Fine timing adjust extension**: Higher resolution timing adjust offset to be appended to timing adjust, integer part. The units are (1/(256×10.24 MHz)) = 0.3814697265625 ns. This parameter provides finer granularity timing offset information for transmission in S-CDMA mode. This TLV MUST be present for S-CDMA channels (see clause 6.2.20).
- **S-CDMA maximum scheduled codes**: The value that the CMTS uses to limit the number of codes scheduled to a CM in an S-CDMA frame. CMs that implement the S-CDMA maximum scheduled codes use this value to limit the maximum size of a concatenated burst in an S-CDMA frame.
- **S-CDMA power headroom**: CMs that implement the S-CDMA maximum scheduled codes MUST use this value to control transmit power as per clause 6.2.18.2.

8.3.6.1 Encodings

The type values used MUST be those defined in Table 8-21 and Figure 8-23. These are unique within the ranging response message but not across the entire MAC message set. The type and length fields MUST each be 1 octet in length.

Table 8-21 – Ranging response message encodings

Name	Type (1 byte)	Length (1 byte)	Value (variable length)
Timing adjust, integer part	1	4	TX timing offset adjustment (signed 32-bit, units of (6.25 µs/64))
Power level adjust	2	1	TX power offset adjustment (signed 8-bit, 1/4-dB units)
Offset frequency adjust	3	2	TX frequency offset adjustment (signed 16-bit, Hz units)
Transmit equalization adjust	4	n	TX equalization data to be convolved with current values (refer to clause 6.2.15). See below for details about representation. This TLV MUST NOT be included in a RNG-RSP that includes a type 9 TLV.
Ranging status	5	1	1 = continue, 2 = abort, 3 = success
Downstream frequency override	6	4	Centre frequency of new downstream channel in Hz
Upstream channel ID override	7	1	Identifier of the new upstream channel
Timing adjust, fractional part	8	1	TX timing fine offset adjustment. Eight-bit unsigned value specifying the fine timing adjustment in units of 1/(245×10.24 MHz)

Table 8-21 – Ranging response message encodings

Name	Type (1 byte)	Length (1 byte)	Value (variable length)
Transmit equalization set	9	n	TX equalization data to be loaded in place of current values (refer to clause 6.2.10). See below for details about representation. This TLV MUST NOT be included in a RNG-RSP to a DOCS 1.x CM, and MUST not be included in a RNG-RSP that includes a type 4 TLV.
S-CDMA maximum scheduled codes	10	1	A CMTS MAY send this TLV only if a CM indicated in the RNG-REQ or INIT-RNG-REQ that it supports the S-CDMA maximum scheduled codes. A value of 0 means no code limit. Other possible values range from 4 to number_active_codes inclusive. Maximum scheduled codes must be an integer multiple of codes_per_mini-slots. If S-CDMA mode is disabled, this TLV MUST NOT be present. Absence of this TLV indicates that maximum scheduled codes is inactive for this CM, which MUST then use the S-CDMA number of active codes.
S-CDMA power headroom	11	1	A CMTS MUST send this TLV to a CM in conjunction with TLV-10. If S-CDMA mode is disabled, this TLV MUST NOT be present. The units are dB. The range of this TLV is from 0 to 4*10log(<Anchor2>) NOTE – A value of 0 for TLV-10 restricts the range to 0 for TLV-11.
Reserved	12-255	n	Reserved for future use

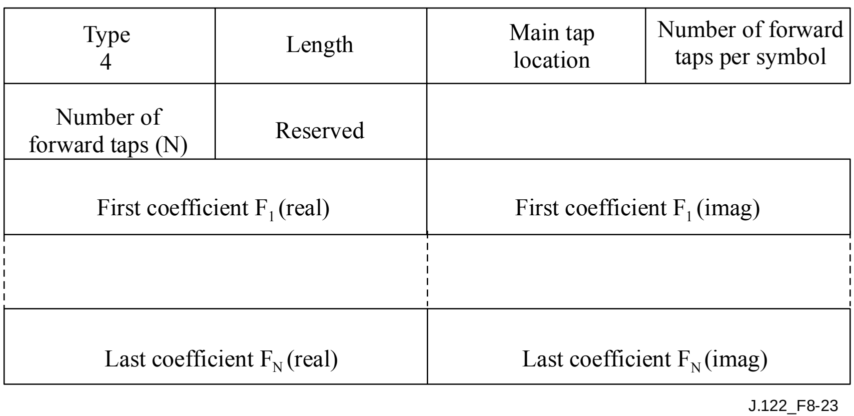

Figure 8-23 – Generalized decision feedback equalization coefficients

The number of taps per modulation interval T MUST be either 1, 2 or 4. The main tap location refers to the position of the zero delay tap, between 1 and N. For a T-spaced equalizer, the number of taps per modulation interval field MUST be set to "1". The total number of taps MAY range up to 64. Each tap consists of a real and imaginary coefficient entry in the table.

If more than 255 bytes are needed to represent equalization information, then several type 4 or 9 elements MAY be used. Data MUST be treated as if byte-concatenated, that is, the first byte after the length field of the second type 4 or 9 element is treated as if it immediately followed the last byte of the first type 4 or 9 element.

Figure 6-28 depicts the operation of the equalizer.

8.3.6.2 Example of TLV data

An example of TLV data is given in Figure 8-24.

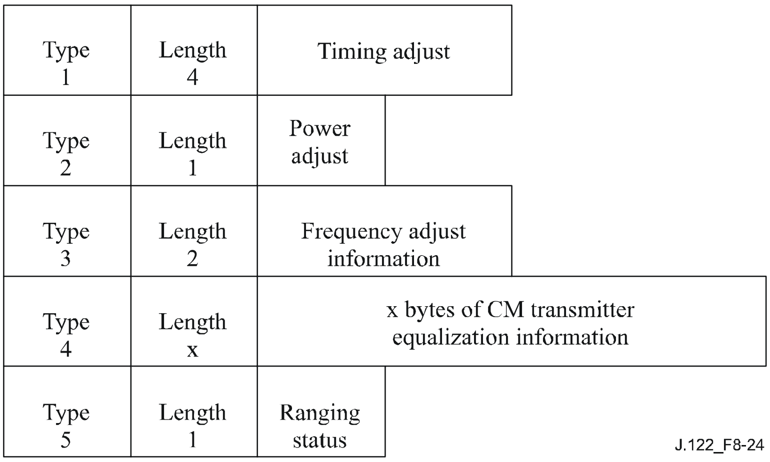

Figure 8-24 – Example of TLV data

8.3.6.3 Overriding channels prior to registration

The RNG-RSP message allows the CMTS to instruct the modem to move to a new downstream and/or upstream channel and to repeat initial ranging. However, the CMTS may do this only in response to an initial ranging request from a modem that is attempting to join the network, or in response to any of the unicast ranging requests that take place immediately after this initial ranging and up to the point where the modem successfully completes periodic ranging. If a downstream frequency override is specified in the RNG-RSP, the modem MUST re-initialize its MAC (see clause 11.2) using initial ranging with the specified downstream centre frequency as the first scanned channel. For the upstream channel, the modem selects its channel based on received UCD messages as per clause 11.2.2.

If an upstream channel ID override is specified in the RNG-RSP, the modem MUST re-initialize its MAC (see clause 11.2) using initial ranging with the upstream channel specified in the RNG-RSP for its first attempt and the same downstream frequency on which the RNG-RSP was received.

If both downstream frequency and upstream channel ID overrides are present in the RNG-RSP, the modem MUST re-initialize its MAC (refer to clause 11.2) using initial ranging with the specified downstream frequency and upstream channel ID for its first attempt.

Note that when a modem with an assigned temporary SID is instructed to move to a new downstream and/or upstream channel and to redo initial ranging, the modem MUST consider the temporary SID to be de-assigned. The modem MUST redo initial ranging using the initialization SID.

Configuration file settings for upstream channel ID and downstream frequency are optional, but if specified in the configuration file they take precedence over the ranging response parameters. Once ranging is complete, only clause C.1.1.2, UCC-REQ and DCC-REQ mechanisms are available for moving the modem to a new upstream channel, and only clauses C.1.1.1, C.1.1.21 and DCC-REQ mechanisms are available for moving the modem to a new downstream channel.

8.3.7 Registration request (REG-REQ)

A registration request MUST be transmitted by a CM at initialization after receipt of a CM parameter file, except as outlined in clauses 11.2.8 and 11.2.9.

To provide for flexibility, the message parameters following the SID MUST be encoded in a type/length/value form.

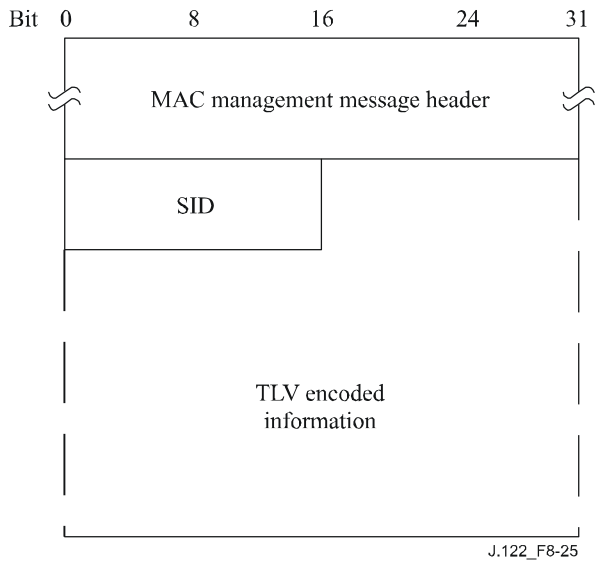

Figure 8-25 – Registration request

A CM MUST generate registration requests in the form shown in Figure 8-25, including the following parameters:

- **SID**: Temporary SID for this CM.

All other parameters are coded as TLV tuples as defined in Annex C.

Registration requests can contain many different TLV parameters, some of which are set by the CM according to its configuration file and some of which are generated by the CM itself. If found in the configuration file, the following configuration settings MUST be included in the registration request.

Configuration file settings:

- All configuration settings included in the CMTS MIC calculation as specified in clause D.3.1.
- Enable 2.0 mode.
- Downstream channel list.
- CMTS MIC configuration setting.

NOTE – The CM MUST forward DOCSIS extension field configuration settings to the CMTS in the same order in which they were received in the configuration file to allow the message integrity check to be performed.

The following registration parameter MUST be included in the registration request.

- Vendor-specific parameter: Vendor ID configuration setting (vendor ID of CM).

The modem capabilities encodings registration parameter MUST also be included in the registration request[13].

The following registration parameters MAY also be included in the registration request:

- Modem IP address.
- Vendor-specific capabilities.

  The vendor-specific capabilities field is for vendor-specific information not included in the configuration file.

The following configuration settings MUST NOT be forwarded to the CMTS in the registration request:

- Software upgrade filename.
- Software upgrade TFTP server IP address.
- SNMP write-access control.
- SNMP MIB object.
- SNMPv3 kickstart value.
- CPE Ethernet MAC address.
- HMAC digest.
- End configuration setting.
- Pad configuration setting.
- Telephone settings option.
- SNMPv3 notification receiver.

8.3.8 Registration response (REG-RSP)

A registration response MUST be transmitted by the CMTS in response to a received REG-REQ.

To provide for flexibility, the message parameters following the response field MUST be encoded in a TLV format.

[13] The CM MUST specify all of its modem capabilities in its registration request subject to the restrictions in clause C.1.3.1. The CMTS MUST NOT assume any modem capability which is defined but not explicitly indicated in the CM's registration request.

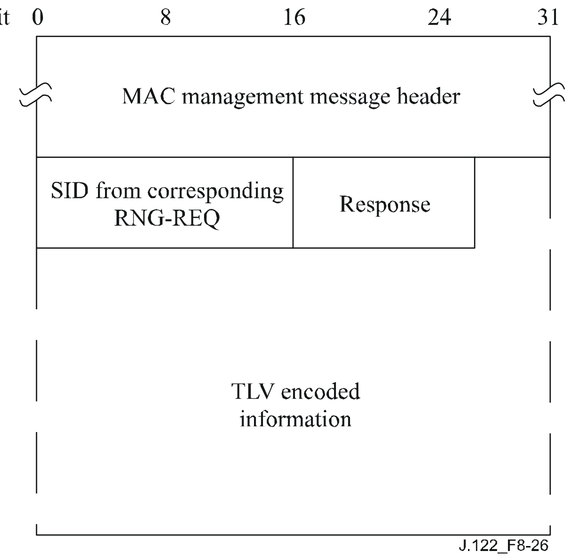

Figure 8-26 – Registration response format

A CMTS MUST generate registration responses in the form shown in Figure 8-26, including both of the following parameters:

- **SID from corresponding REG-REQ**: SID from corresponding REG-REQ to which this response refers (this acts as a transaction identifier).
- **Response**: For REG-RSP to a modem registering as a 1.0 modem (i.e., REG-REQ contains DOCS 1.0 class of service encodings).

  0 = Okay

  1 = Authentication failure

  2 = Class of service failure

For REG-RSP to a modem registering as a 1.1 or 2.0 modem (i.e., REG-REQ contains service flow encodings), this field MUST contain one of the confirmation codes in clauses C.4 and C.4.2.

NOTE 1 – Failures apply to the entire registration request. Even if only a single requested service flow or DOCS 1.0 service class is invalid or undeliverable, the entire registration is failed.

If the REG-REQ was successful, and contained service flow parameters, classifier parameters or payload header suppression parameters, the REG-RSP MUST contain, for each of these:

- **Classifier parameters**: All of the classifier parameters from the corresponding REG-REQ, plus the classifier identifier assigned by the CMTS.
- **Service flow parameters**: All the service flow parameters from the REG-REQ, plus the service flow ID assigned by the CMTS. Every service flow that contained a service class name that was admitted/activated[14] MUST be expanded into the full set of TLVs defining the service flow. Every upstream service flow that was admitted/activated MUST have a service identifier assigned by the CMTS. A service flow that was only provisioned will include only those QoS parameters that appeared in the REG-REQ, plus the assigned service flow ID.
- **Payload header suppression parameters**: All the payload header suppression parameters from the REG-REQ, plus the payload header suppression index assigned by the CMTS.

[14] The ActiveQoSParamSet or AdmittedQoSParamSet is non-null.

If the REG-REQ failed due to service flow parameters, classifier parameters or payload header suppression parameters, and the response is not one of the major error codes in clause C.4.2, the REG-RSP MUST contain at least one of the following:

- **Classifier error set**: A classifier error set and identifying classifier reference and service flow reference MUST be included for at least one failed classifier in the corresponding REG-REQ. Every classifier error set MUST include at least one specific failed classifier parameter of the corresponding classifier.
- **Service flow error set**: A service flow error set and identifying service flow reference MUST be included for at least one failed service flow in the corresponding REG-REQ. Every service flow error set MUST include at least one specific failed QoS parameter of the corresponding service flow.
- **Payload header suppression error set**: A PHS error set and identifying service flow reference and classifier reference pair MUST be included for at least one failed PHS rule in the corresponding REG-REQ. Every PHS error set MUST include at least one specific failed PHS parameter of the corresponding failed PHS rule.

Service class name expansion always occurs at admission time. Thus, if a registration request contains a service flow reference and a service class name for deferred admission/activation, the registration response MUST NOT include any additional QoS parameters except the service flow identifier (refer to clause 10.1.3).

If the corresponding registration request contains DOCS 1.0 service class TLVs (refer to clause C.1.1.4), the registration response MUST contain the following TLV tuples:

- **DOCS 1.0 service class data**: Returned when response = okay. This is a service ID/service class tuple for each class of service granted.

  NOTE 2 – Service class IDs MUST be those requested in the corresponding REG-REQ.
- **Service not available**: Returned when response = class of service failure. If a service class cannot be supported, this configuration setting is returned in place of the service class data.

All other parameters are coded TLV tuples.

- **Modem capabilities**: The CMTS response to the capabilities of the modem (if present in the registration request).
- **Vendor-specific data**: As defined in Annex C:
  – vendor ID configuration setting (vendor ID of the CMTS);
  – vendor-specific extensions.

8.3.8.1 Encodings

The type values used MUST be those shown below. These are unique within the registration response message but not across the entire MAC message set. The type and length fields MUST each be one octet.

8.3.8.1.1 Modem capabilities

This field defines the CMTS response to the modem capabilities field in the registration request. The CMTS MUST respond to the modem capability to indicate whether they may be used. If the CMTS does not recognize a modem capability, it MUST return the TLV with the value zero ("off") in the registration response.

Only capabilities set to "on" in the REG-REQ may be set "on" in the REG-RSP as this is the handshake indicating that they have been successfully negotiated. Capabilities set to "off" in the REG-REQ MUST also be set to "off" in the REG-RSP.

Encodings are as defined for the registration request.

8.3.8.1.2 DOCS 1.0 service class data

A DOCS 1.0 service-class-data parameter MUST be present in the registration response for each DOCS 1.0 class-of-service parameter (refer to clause C.1.1.4) in the registration request.

This encoding defines the parameters associated with a requested class of service. It is somewhat complex in that it is composed from a number of encapsulated type/length/value fields. The encapsulated fields define the particular class-of-service parameters for the class of service in question. Note that the type fields defined are only valid within the encapsulated service class data configuration setting string. A single service class data configuration setting MUST be used to define the parameters for a single service class. Multiple class definitions MUST use multiple service class data configuration setting sets.

Each received DOCS 1.0 class-of-service parameter must have a unique class ID in the range 1 to 16. If no class ID is present for any single DOCS 1.0 class-of-service TLV in the REG-REQ, the CMTS MUST send a REG-RSP with a class-of-service failure response and no DOCS 1.0 class-of-service TLVs.

Type	Length	Value
1	n	Encoded service class data

Class ID:

The value of the field MUST specify the identifier for the class of service to which the encapsulated string applies. This MUST be a class which was requested in the associated REG-REQ, if present.

Type	Length	Value
1.1	1	From REG-REQ

Valid range: The class ID MUST be in the range 1 to 16.

Service ID: The value of the field MUST specify the SID associated with this service class.

Type	Length	Value
1.2	2	SID

8.3.9 Registration acknowledge (REG-ACK)

A registration acknowledge MUST be transmitted by the CM in response to a REG-RSP from the CMTS with a confirmation code of okay (0)[15]. It confirms acceptance by the CM of the QoS parameters of the flow as reported by the CMTS in its REG-RSP. The format of a REG-ACK MUST be as shown in Figure 8-27.

[15] The registration acknowledge is a DOCS 1.1/2.0 message. Refer to Annex G for details of registration interoperability issues.

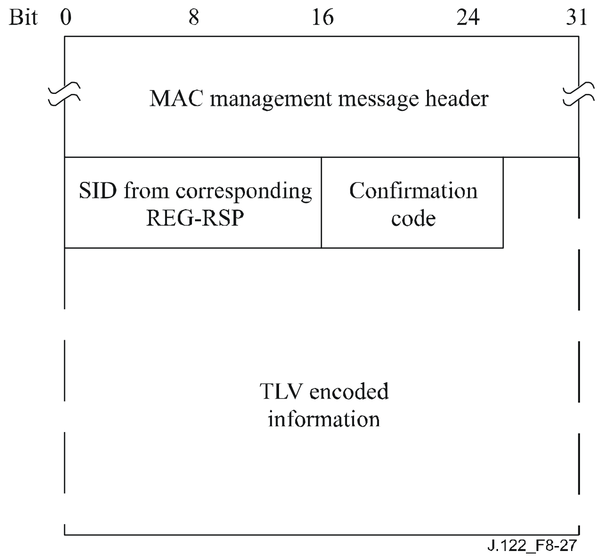

Figure 8-27 – Registration acknowledgment

The parameters MUST be as follows:

- **SID from corresponding REG-RSP**: SID from corresponding REG-RSP to which this acknowledgment refers (this acts as a transaction identifier).
- **Confirmation code**: The appropriate confirmation code (refer to clause C.4) for the entire corresponding registration response.

The CM is required to send all provisioned classifiers, service flows and payload header suppression rules to the CMTS in the REG-REQ (see clause 8.3.7). The CMTS will return them with identifiers, expanding service class names if present, in the REG-RSP (see clause 8.3.8). Since the CM may be unable to support one or more of these provisioned items, the REG-ACK includes error sets for all failures related to these provisioned items.

If there were any failures of provisioned items, the REG-ACK MUST include the error sets corresponding to those failures. The error set identification is provided by using service flow ID and classifier ID from corresponding REG-RSP. If a classifier ID or SFID was omitted in the REG-RSP, the CM MUST use the appropriate reference (classifier reference, SF reference) in the REG-ACK.

- **Classifier error set**: A classifier error set and identifying classifier reference/identifier and service flow reference/identifier pair MUST be included for at least one failed classifier in the corresponding REG-RSP. Every classifier error set MUST include at least one specific failed classifier parameter of the corresponding classifier. This parameter MUST be omitted if the entire REG-REQ/RSP is successful.
- **Service flow error set**: A service flow error set of the REG-ACK message encodes specifics of failed service flows in the REG-RSP message. A service flow error set and identifying service flow reference/identifier MUST be included for at least one failed QoS parameter of at least one failed service flow in the corresponding REG-RSP message. This parameter MUST be omitted if the entire REG-REQ/RSP is successful.
- **Payload header suppression error set**: A PHS error set and identifying service flow reference/identifier and classifier reference/identifier pair MUST be included for at least one failed PHS rule in the corresponding REG-RSP. Every PHS error set MUST include at least one specific failed PHS of the failed PHS rule. This parameter MUST be omitted if the entire REG-REQ/RSP is successful.

Per service flow acknowledgment is necessary not just for synchronization between the CM and CMTS, but also to support use of the service class name (refer to clause 10.1.3). Since the CM may not know all of the service flow parameters associated with a service class name when making the registration request, it may be necessary for the CM to NAK a registration response if it has insufficient resources to support this service flow.

8.3.10 Upstream channel change request (UCC-REQ)

An upstream channel change request MAY be transmitted by a CMTS to cause a CM to change the upstream channel on which it is transmitting. The format of an UCC-REQ message is shown in Figure 8-28.

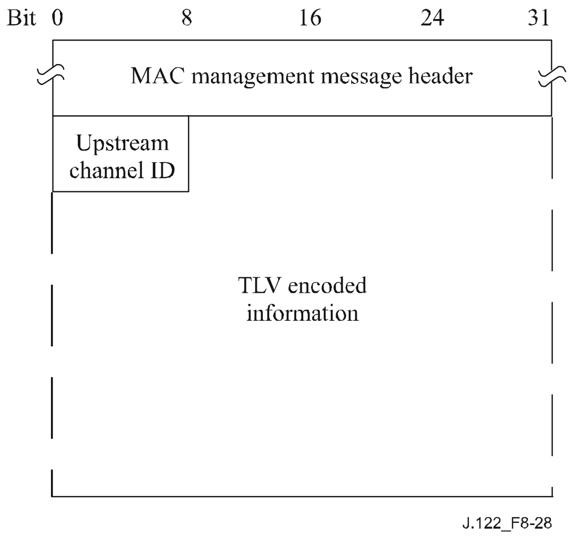

Figure 8-28 – Upstream channel change request

Parameters MUST be as follows:

- **Upstream channel ID**: The identifier of the upstream channel to which the CM is to switch for upstream transmissions. This is an 8-bit field.

Upon receipt of a UCC-REQ message, the CM MUST perform ranging with broadcast initial maintenance.

8.3.11 Upstream channel change response (UCC-RSP)

An upstream channel change response MUST be transmitted by a CM in response to a received upstream channel change request message to indicate that it has received and is complying with the UCC-REQ. The format of an UCC-RSP message is shown in Figure 8-29.

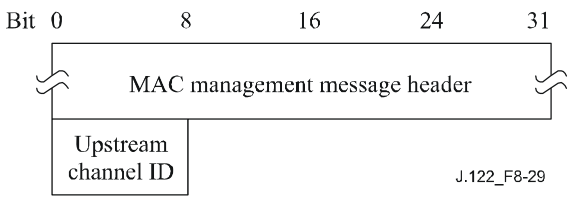

Figure 8-29 – Upstream channel change response

Before it begins to switch to a new upstream channel, a CM MUST transmit a UCC-RSP on its existing upstream channel. A CM MAY ignore an UCC-REQ message while it is in the process of performing a channel change. When a CM receives a UCC-REQ message requesting that it switch to an upstream channel that it is already using, the CM MUST respond with a UCC-RSP message on that channel indicating that it is already using the correct channel.

After switching to a new upstream channel, a CM MUST re-range using broadcast initial ranging, and then MUST proceed without re-performing registration. The full procedure for changing channels is described in clause 11.3.3.

Parameters MUST be as follows:

- **Upstream channel ID**: The identifier of the upstream channel to which the CM is to switch for upstream transmissions. This MUST be the same channel ID specified in the UCC-REQ message. This MUST be an 8-bit field.

8.3.12 Dynamic service addition request (DSA-REQ)

A dynamic service addition request MAY be sent by a CM or CMTS to create a new service flow.

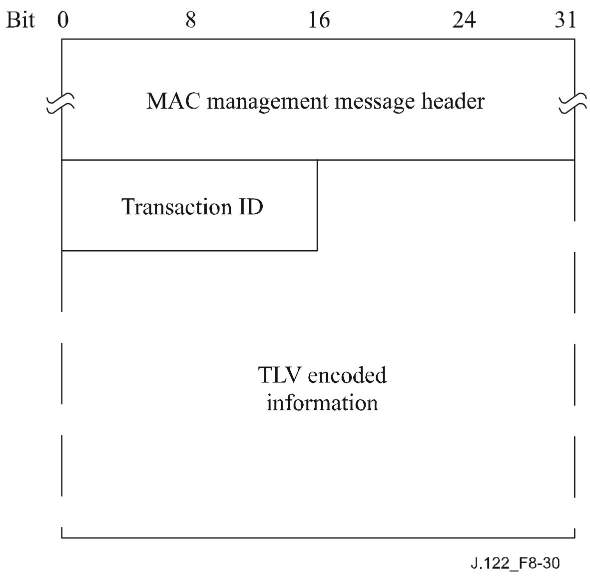

Figure 8-30 – Dynamic service addition request

A CM or CMTS MUST generate DSA-REQ messages in the form shown in Figure 8-30, including the following parameter:

- **Transaction ID**: Unique identifier for this transaction assigned by the sender.

All other parameters are coded as TLV tuples as defined in Annex C.

A DSA-REQ message MUST NOT contain parameters for more than one service flow in each direction, i.e., a DSA-REQ message MUST contain parameters for either a single upstream service flow, or for a single downstream service flow, or for one upstream and one downstream service flow.

The DSA-REQ message MUST contain:

- **Service flow parameters**: Specification of the service flow's traffic characteristics and scheduling requirements.

The DSA-REQ message MAY contain classifier parameters and payload header suppression parameters associated with the service flows specified in the message:

- **Classifier parameters**: Specification of the rules to be used to classify packets into a specific service flow.

- **Payload header suppression parameters**: Specification of the payload header suppression rules to be used with an associated classifier.

If privacy is enabled, the DSA-REQ message MUST contain:

- **Key sequence number**: The key sequence number of the auth key, which is used to calculate the HMAC-Digest (refer to clause C.1.4.3).
- **HMAC-Digest**: The HMAC-Digest attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest attribute MUST be the final attribute in the dynamic service message's attribute list (refer to clause C.1.4.1).

8.3.12.1 CM-initiated dynamic service addition

CM-initiated DSA-Requests MUST use the service flow reference to link classifiers to service flows. Values of the service flow reference are local to the DSA message; each service flow within the DSA-Request MUST be assigned a unique service flow reference. This value need not be unique with respect to the other service flows known by the sender.

CM-initiated DSA-Request MUST use the classifier reference and service flow reference to link payload header suppression parameters to classifiers and service flows. A DSA-Request MUST use the service flow reference to link classifier to service flow. Values of the classifier reference are local to the DSA message; each classifier within the DSA-Request MUST be assigned a unique classifier reference.

CM-initiated DSA-Requests MAY use the service class name (refer to clause C.2.2.3.4) in place of some, or all, of the QoS parameters.

8.3.12.2 CMTS-initiated dynamic service addition

CMTS-initiated DSA-Requests MUST use the service flow ID to link classifiers to service flows. Service flow identifiers are unique within the MAC domain. CMTS-initiated DSA-Requests for upstream service flows MUST also include a service ID.

CMTS-initiated DSA-Requests which include classifiers, MUST assign a unique classifier identifier on a per service flow basis.

CMTS-initiated DSA-Requests for named service classes MUST include the QoS parameter set associated with that service class.

8.3.13 Dynamic service addition response (DSA-RSP)

A dynamic service addition response MUST be generated in response to a received DSA-Request. The format of a DSA-RSP MUST be as shown in Figure 8-31.

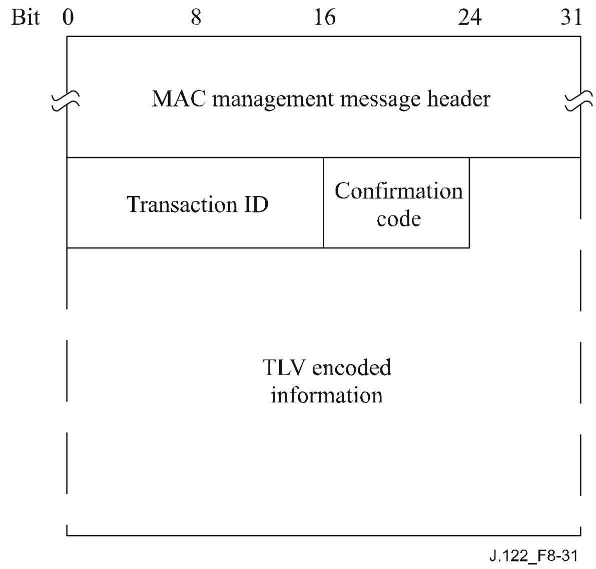

Figure 8-31 – Dynamic service addition response

Parameters MUST be as follows:

- **Transaction ID**: Transaction ID from corresponding DSA-REQ.
- **Confirmation code**: The appropriate confirmation code (refer to clause C.4) for the entire corresponding DSA-Request.

All other parameters are coded as TLV tuples as defined in Annex C.

If the transaction is successful, the DSA-RSP MAY contain one or more of the following:

- **Classifier parameters**: The CMTS MUST include the complete specification of the classifier in the DSA-RSP, including a newly assigned classifier identifier. The CM MUST NOT include the specification of the classifier in the DSA-RSP.
- **Service flow parameters**: The CMTS MUST include the complete specification of the service flow in the DSA-RSP, including a newly assigned service flow identifier and an expanded service class name if applicable. The CM MUST NOT include the specification of the service flow in the DSA-RSP.
- **Payload header suppression parameters**: The CMTS MUST include the complete specification of the PHS parameters in the DSA-RSP, including a newly assigned PHS index, a classifier identifier and a service flow identifier. The CM MUST NOT include the specification of the PHS parameters.

If the transaction is unsuccessful due to service flow parameters, classifier parameters or payload header suppression parameters, and the confirmation code is not one of the major error codes in clause C.4.2, the DSA-RSP MUST contain at least one of the following:

- **Service flow error set**: A service flow error set and identifying service flow reference/identifier MUST be included for at least one failed service flow in the corresponding DSA-REQ. Every service flow error set MUST include at least one specific failed QoS parameter of the corresponding service flow. This parameter MUST be omitted if the entire DSA-REQ is successful.
- **Classifier error set**: A classifier error set and identifying classifier reference/identifier and service flow reference/identifier pair MUST be included for at least one failed classifier in the corresponding DSA-REQ. Every classifier error set MUST include at least one specific failed classifier parameter of the corresponding classifier. This parameter MUST be omitted if the entire DSA-REQ is successful.

- **Payload header suppression error set**: A PHS error set and identifying classifier reference/identifier and service flow reference/identifier pair MUST be included for at least one failed PHS rule in the corresponding DSA-REQ. Every PHS error set MUST include at least one specific failed PHS parameter of the corresponding failed PHS rule. This parameter MUST be omitted if the entire DSA-REQ is successful.

If privacy is enabled, the DSA-RSP message MUST contain:

- **Key sequence number**: The key sequence number of the auth key, which is used to calculate the HMAC-Digest (refer to clause C.1.4.3).
- **HMAC-Digest**: The HMAC-Digest attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest attribute MUST be the final attribute in the dynamic service message's attribute list (refer to clause C.1.4.1).

8.3.13.1 CM-initiated dynamic service addition

The CMTS's DSA-Response for service flows that are successfully added MUST contain a service flow ID. The DSA-Response for successfully admitted or active upstream QoS parameter sets MUST also contain a service ID.

If the corresponding DSA-Request uses the service class name (refer to clause C.2.2.3.4) to request service addition, a DSA-Response MUST contain the QoS parameter set associated with the named service class. If the service class name is used in conjunction with other QoS parameters in the DSA-Request, the CMTS MUST accept or reject the DSA-Request using the explicit QoS parameters in the DSA-Request. If these service flow encodings conflict with the service class attributes, the CMTS MUST use the DSA-Request values as overrides for those of the service class.

If the transaction is successful, the CMTS MUST assign a classifier identifier to each requested classifier and a PHS index to each requested PHS rule. The CMTS MUST use the original classifier reference(s) and service flow reference(s) to link the successful parameters in the DSA-RSP.

If the transaction is unsuccessful, the CMTS MUST use the original classifier reference(s) and service flow reference(s) to identify the failed parameters in the DSA-RSP.

8.3.13.2 CMTS-initiated dynamic service addition

If the transaction is unsuccessful, the CM MUST use the classifier identifier(s) and service flow identifier(s) to identify the failed parameters in the DSA-RSP.

8.3.14 Dynamic service addition acknowledge (DSA-ACK)

A dynamic service addition acknowledge MUST be generated in response to a received DSA-RSP. The format of a DSA-ACK MUST be as shown in Figure 8-32.

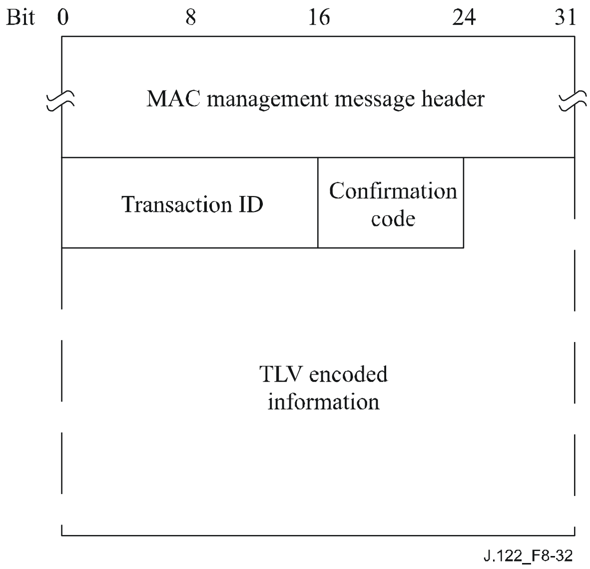

Figure 8-32 – Dynamic service addition acknowledge

Parameters MUST be as follows:

- **Transaction ID**: Transaction ID from corresponding DSA-Response.
- **Confirmation code**: The appropriate confirmation code (refer to clause C.4) for the entire corresponding DSA-Response[16].

All other parameters are coded TLV tuples.

- **Service flow error set**: The service flow error set of the DSA-ACK message encodes specifics of failed service flows in the DSA-RSP message. A service flow error set and identifying service flow reference/identifier MUST be included for at least one failed QoS parameter of at least one failed service flow in the corresponding DSA-REQ. This parameter MUST be omitted if the entire DSA-REQ is successful.

If privacy is enabled, the DSA-RSP message MUST contain:

- **Key sequence number**: The key sequence number of the auth key, which is used to calculate the HMAC-Digest (refer to clause C.1.4.3).
- **HMAC-Digest**: The HMAC-Digest attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest attribute MUST be the final attribute in the dynamic service message's attribute list (refer to clause C.1.4.1).

8.3.15 Dynamic service change request (DSC-REQ)

A dynamic service change request MAY be sent by a CM or CMTS to dynamically change the parameters of an existing service flow. DSCs changing classifiers MUST carry the entire classifier TLV set for that new classifier.

[16] The confirmation code is necessary particularly when a service class name (refer to clause 10.1.3) is used in the DSA-Request. In this case, the DSA-Response could contain service flow parameters that the CM is unable to support (either temporarily or as configured).

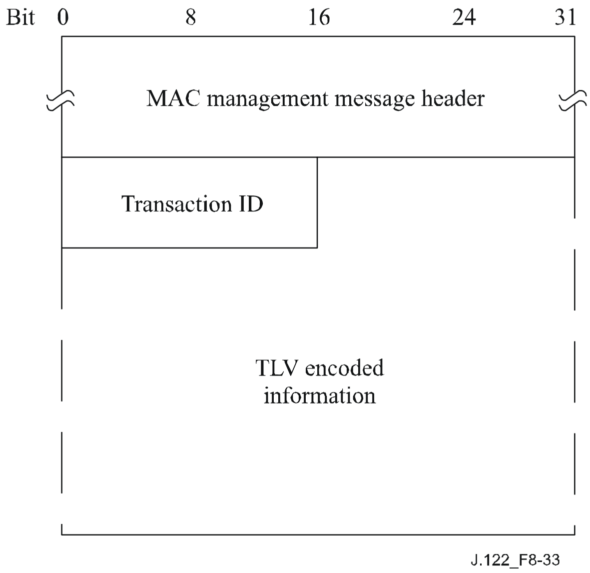

Figure 8-33 – Dynamic service change request

A CM or CMTS MUST generate DSC-REQ messages in the form shown in Figure 8-33 including the following parameters:

- **Transaction ID**: Unique identifier for this transaction assigned by the sender.

All other parameters are coded as TLV tuples as defined in Annex C.

A DSC-REQ message MUST NOT carry parameters for more than one service flow in each direction, i.e., a DSC-REQ message MUST contain parameters for either a single upstream service flow, or for a single downstream service flow, or for one upstream and one downstream service flow. A DSC-REQ MUST contain at least one of the following:

- **Classifier parameters**: Specification of the rules to be used to classify packets into a specific service flow – this includes the dynamic service change action TLV which indicates whether this classifier should be added, replaced or deleted from the service flow (refer to clause C.2.1.3.7). If included, the classifier parameters MUST contain the dynamic change action TLV, a classifier reference/identifier and a service flow identifier.
- **Service flow parameters**: Specification of the service flow's new traffic characteristics and scheduling requirements. The admitted and active quality of service parameter sets in this message replace the admitted and active quality of service parameter sets currently in use by the service flow. If the DSC message is successful and it contains service flow parameters, but does not contain replacement sets for both admitted and active quality of service parameter sets, the omitted set(s) MUST be set to null. If included, the service flow parameters MUST contain a service flow identifier.
- **Payload header suppression parameters**: Specification of the rules to be used for payload header suppression to suppress payload headers related to a specific classifier – this includes the dynamic service change action TLV which indicates whether this PHS rule should be added, set or deleted from the service flow or whether all the PHS rules for the service flow specified should be deleted (refer to clause C.2.2.8.5). If included, the PHS parameters MUST contain the dynamic service change action TLV, a classifier reference/identifier and a service flow identifier, unless the dynamic service change action is "delete all PHS rules". If the dynamic service change action is "delete all PHS rules", the PHS parameters MUST contain a service flow identifier along with the dynamic service change action, and no other PHS parameters need be present in this case. However, if other

PHS parameters are present, in particular payload header suppression index, they MUST be ignored by the receiver of the DSC-REQ message.

If privacy is enabled, a DSC-REQ MUST also contain:

- **Key sequence number**: The key sequence number of the auth key, which is used to calculate the HMAC-Digest (refer to clause C.1.4.3).
- **HMAC-Digest**: The HMAC-Digest attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest attribute MUST be the final attribute in the dynamic service message's attribute list (refer to clause C.1.4.1).

8.3.16 Dynamic service change response (DSC-RSP)

A dynamic service change response MUST be generated in response to a received DSC-REQ. The format of a DSC-RSP MUST be as shown in Figure 8-34.

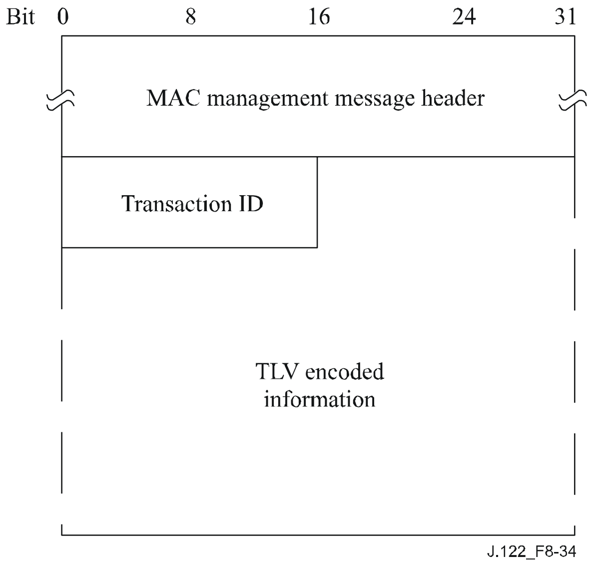

Figure 8-34 – Dynamic service change response

Parameters MUST be as follows:

- **Transaction ID**: Transaction ID from corresponding DSC-REQ.
- **Confirmation code**: The appropriate confirmation code (refer to clause C.4) for the corresponding DSC-Request.

All other parameters are coded as TLV tuples as defined in Annex C.

If the transaction is successful, the DSC-RSP MAY contain one or more of the following:

- **Classifier parameters**: The CMTS MUST include the complete specification of the classifier in the DSC-RSP, including a newly assigned classifier identifier for new classifiers. The CM MUST NOT include the specification of the classifier in the DSC-RSP.
- **Service flow parameters**: The CMTS MUST include the complete specification of the service flow in the DSC-RSP, including an expanded service class name if applicable. The CMTS MUST include a SID in the DSC-RSP if a service flow parameter set contained an upstream admitted QoS parameter set and this service flow does not have an associated SID. If a service flow parameter set contained a service class name and an admitted QoS parameter set, the CMTS MUST include the QoS parameter set corresponding to the named service class in the DSC-RSP. If specific QoS parameters were also included in the classed service flow request, the CMTS MUST include these QoS parameters in the DSC-RSP

instead of any QoS parameters of the same type of the named service class. The CM MUST NOT include the specification of the service flow in the DSC-RSP.

- **Payload header suppression parameters**: The CMTS MUST include the complete specification of the PHS parameters in the DSC-RSP, including a newly assigned PHS index for new PHS rules, a classifier identifier and a service flow identifier. The CM MUST NOT include the specification of the PHS parameters.

If the transaction is unsuccessful due to service flow parameters, classifier parameters or payload header suppression parameters, and the confirmation code is not one of the major error codes in clause C.4.2, the DSC-RSP MUST contain at least one of the following:

- **Classifier error set**: A classifier error set and identifying classifier reference/identifier and service flow reference/identifier pair MUST be included for at least one failed classifier in the corresponding DSC-REQ. Every classifier error set MUST include at least one specific failed classifier parameter of the corresponding classifier. This parameter MUST be omitted if the entire DSC-REQ is successful.
- **Service flow error set**: A service flow error set and identifying service flow ID MUST be included for at least one failed service flow in the corresponding DSC-REQ. Every service flow error set MUST include at least one specific failed QoS parameter of the corresponding service flow. This parameter MUST be omitted if the entire DSC-REQ is successful.
- **Payload header suppression error set**: A PHS error set and identifying service flow reference/identifier and classifier reference/identifier pair MUST be included for at least one failed PHS rule in the corresponding DSC-REQ, unless the dynamic service change action is "delete all PHS rules". If the dynamic service change action is "delete all PHS rules" the PHS error set(s) MUST include an identifying service flow ID. Every PHS error set MUST include at least one specific failed PHS parameter of the corresponding failed PHS rule. This parameter MUST be omitted if the entire DSC-REQ is successful.

Regardless of success or failure, if privacy is enabled for the CM, the DSC-RSP MUST contain:

- **Key sequence number**: The key sequence number of the auth key, which is used to calculate the HMAC-Digest (refer to clause C.1.4.3).
- **HMAC-Digest**: The HMAC-Digest attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest attribute MUST be the final attribute in the dynamic service message's attribute list (refer to clause C.1.4.1).

8.3.17 Dynamic service change acknowledge (DSC-ACK)

A dynamic service change acknowledge MUST be generated in response to a received DSC-RSP. The format of a DSC-ACK MUST be as shown in Figure 8-35.

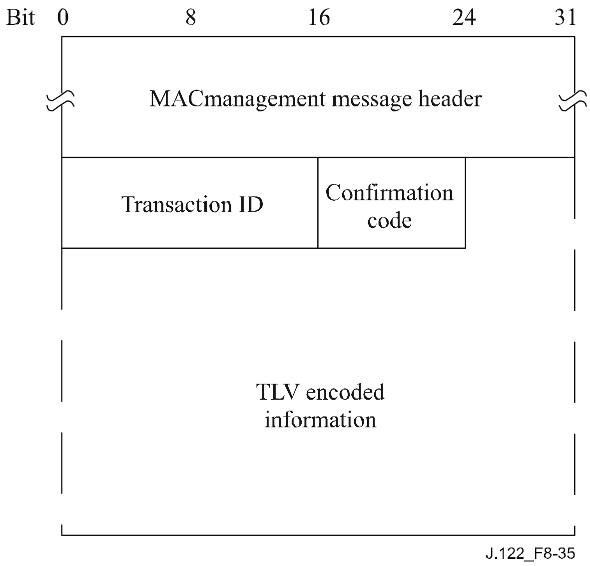

Figure 8-35 – Dynamic service change acknowledge

Parameters MUST be as follows:

- **Transaction ID**: Transaction ID from the corresponding DSC-REQ.
- **Confirmation code**: The appropriate confirmation code (refer to clause C.4) for the entire corresponding DSC-Response[17].

All other parameters are coded as TLV tuples.

- **Service flow error set**: The service flow error set of the DSC-ACK message encodes specifics of failed service flows in the DSC-RSP message. A service flow error set and identifying service flow identifier MUST be included for at least one failed QoS parameter of at least one failed service flow in the corresponding DSC-REQ. This parameter MUST be omitted if the entire DSC-REQ is successful.

If privacy is enabled, the DSC-ACK message MUST contain:

- **Key sequence number**: The key sequence number of the auth key, which is used to calculate the HMAC-Digest (refer to clause C.1.4.3).
- **HMAC-Digest**: The HMAC-Digest attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest attribute MUST be the final attribute in the dynamic service message's attribute list (refer to clause C.1.4.1).

8.3.18 Dynamic service deletion request (DSD-REQ)

A dynamic service deletion request MAY be sent by a CM or CMTS to delete a single existing upstream service flow and/or a single existing downstream service flow. The format of a DSD-Request MUST be as shown in Figure 8-36.

[17] The confirmation code and service flow error set are necessary particularly when a service class name (refer to clause 10.1.3) is used in the DSC-Request. In this case, the DSC-Response could contain service flow parameters that the CM is unable to support (either temporarily or as configured).

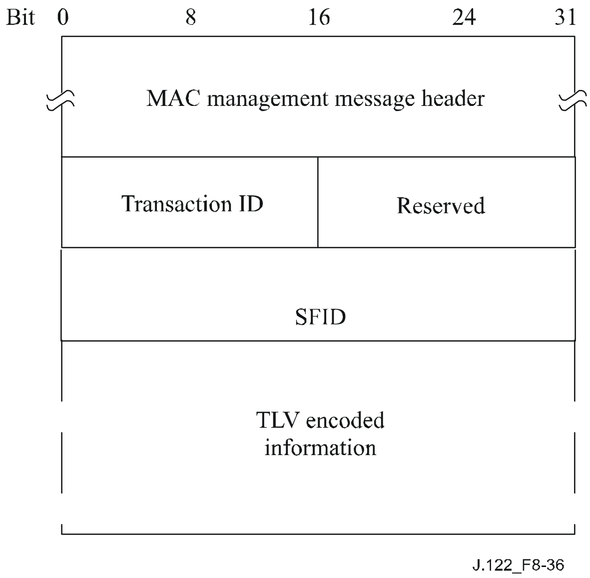

Figure 8-36 – Dynamic service deletion request

Parameters MUST be as follows:

- **Service flow identifier**: If this value is non-zero, it is the SFID of a single upstream or single downstream service flow to be deleted. If this value is zero, the service flow(s) to be deleted will be identified by SFID(s) in the TLVs. If this value is non-zero, any SFIDs included in the TLVs MUST be ignored.
- **Transaction ID**: Unique identifier for this transaction assigned by the sender.

All other parameters are coded as TLV tuples as defined in Annex C.

- **Service flow identifier**: The SFID(s) to be deleted, which MUST be encoded per clause C.2.2.3.2. The service flow identifier TLV MUST be the only service flow encoding sub-TLV used.

If privacy is enabled, the DSD-REQ MUST include:

- **Key sequence number**: The key sequence number of the auth key, which is used to calculate the HMAC-Digest (refer to clause C.1.4.3).
- **HMAC-Digest**: The HMAC-Digest attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest attribute MUST be the final attribute in the dynamic service message's attribute list (refer to clause C.1.4.1).

8.3.19 Dynamic service deletion response (DSD-RSP)

A dynamic service deletion response MUST be generated in response to a received DSD-REQ. The format of a DSD-RSP MUST be as shown in Figure 8-37.

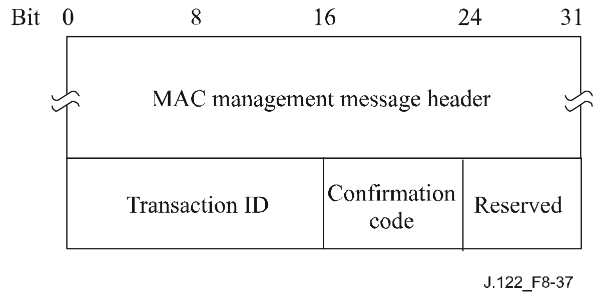

Figure 8-37 – Dynamic service deletion response

Parameters MUST be as follows:

- **Transaction ID**: Transaction ID from corresponding DSD-REQ.
- **Confirmation Code**: The appropriate confirmation code (refer to clause C.4) for the corresponding DSD-Request.

8.3.20 Dynamic channel change request (DCC-REQ)

A dynamic channel change request MAY be transmitted by a CMTS to cause a CM to change the upstream channel on which it is transmitting, the downstream channel on which it is receiving, or both.

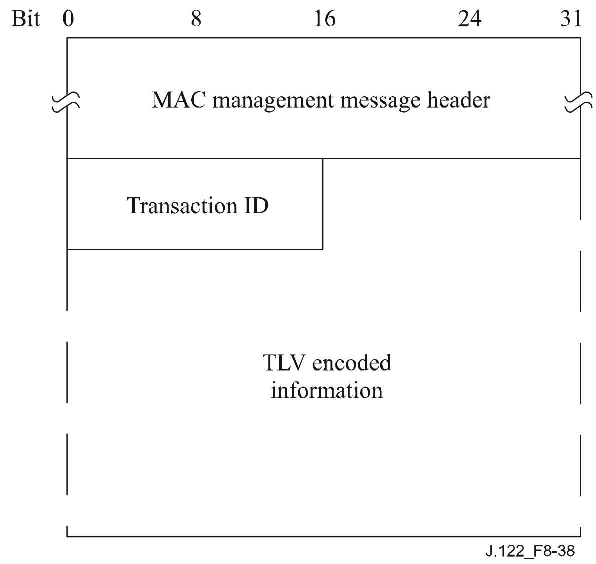

Figure 8-38 – Dynamic channel change request

A CMTS MUST generate DCC-REQ message in the form shown in Figure 8-38, including the following parameters:

- **Transaction ID**: A 16-bit unique identifier for this transaction assigned by the sender.

The following parameters are optional and are coded as TLV tuples:

- **Upstream channel ID**: The identifier of the upstream channel to which the CM is to switch for upstream transmissions.
- **Downstream parameters**: The frequency of the downstream channel to which the CM is to switch for downstream reception.

- **Initialization technique**: Directions for the type of initialization, if any, that the CM should perform once synchronized to the new channel(s).
- **UCD substitution**: Provides a copy of the UCD for the new channel. This TLV occurs as many times as necessary to contain one UCD.
- **SAID substitution**: A pair of security association identifiers (SAID) which contain the current SAID and the new SAID for the new channel. This TLV occurs once if the SAID requires substitution.
- **Service flow substitution**: A group of sub-TLVs which allows substitution in a service flow of the service flow identifier and service identifier. This TLV is repeated for every service flow which has parameters requiring substitution.

If privacy is enabled, a DCC-REQ MUST also contain:

- **Key sequence number**: The key sequence number of the auth key, which is used to calculate the HMAC-Digest (refer to clause C.1.4.3).
- **HMAC-Digest**: The HMAC-Digest attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest attribute MUST be the final attribute in the dynamic channel change message's attribute list (refer to clause C.1.4.1).

8.3.20.1 Encodings

The type values used MUST be those shown below. These are unique within the dynamic channel change request message, but not across the entire MAC message set.

8.3.20.1.1 Upstream channel ID

When present, this TLV specifies the new upstream channel ID that the CM MUST use when performing a dynamic channel change. It is an override for the current upstream channel ID. The CMTS SHOULD ensure that the upstream channel ID for the new channel is different from the upstream channel ID for the old channel. This TLV MUST be included if the upstream channel is changed, even if the UCD substitution TLV is included.

Type	Length	Value
1	1	0-255: Upstream channel ID

If this TLV is missing, the CM MUST NOT change its upstream channel ID. The CMTS MAY include this TLV. The CM MUST observe this TLV.

8.3.20.1.2 Downstream parameters

When present, this TLV specifies the operating parameters of the new downstream channel. The value field of this TLV contains a series of subtypes.

Type	Length	Value
2	n	

The CMTS MUST include this TLV when specifying a downstream channel change. If this TLV is missing, the CM MUST NOT change its downstream parameters.

8.3.20.1.2.1 Downstream frequency

This TLV specifies the new receive frequency that the CM MUST use when performing a dynamic channel change. It is an override for the current downstream channel frequency. This is the centre frequency of the downstream channel in Hz and is stored as a 32-bit binary number. The downstream frequency MUST be a multiple of 62 500 Hz.

Subtype	Length	Value
2.1	4	Rx frequency

The CMTS MUST include this sub-TLV. The CM MUST observe this sub-TLV.

8.3.20.1.2.2 Downstream modulation type

This TLV specifies the modulation type that is used on the new downstream channel.

Subtype	Length	Value
2.2	1	0: 64QAM 1: 256QAM 2-255: Reserved

The CMTS SHOULD include this sub-TLV. The CM SHOULD observe this sub-TLV.

8.3.20.1.2.3 Downstream symbol rate

This TLV specifies the symbol rate that is used on the new downstream channel.

Subtype	Length	Value
2.3	1	0: 5.056941 Msymb/s 1: 5.360537 Msymb/s 2: 6.952 Msymb/s 3-255: Reserved

The CMTS SHOULD include this sub-TLV. The CM SHOULD observe this sub-TLV.

8.3.20.1.2.4 Downstream interleaver depth

This TLV specifies the parameters "I" and "J" of the downstream interleaver.

Subtype	Length	Value
2.4	2	I: 0-255 J: 0-255

The CMTS SHOULD include this sub-TLV. The CM SHOULD observe this sub-TLV.

8.3.20.1.2.5 Downstream channel identifier

This TLV specifies the 8-bit downstream channel identifier of the new downstream channel.

Subtype	Length	Value
2.5	1	0-255: Downstream channel ID

The CMTS SHOULD include this sub-TLV. The CM SHOULD observe this sub-TLV.

8.3.20.1.2.6 SYNC Substitution

When present, this TLV allows the CMTS to inform the CM to wait or not wait for a SYNC message before proceeding. The CMTS MUST have synchronized timestamps between the old and new channel(s) if it instructs the CM not to wait for a SYNC message before transmitting on the new channel. Synchronized timestamps implies that the timestamps are derived from the same clock and contain the same value.

Subtype	Length	Value
2.6	1	0: Acquire SYNC message on the new downstream channel before proceeding 1: Proceed without first obtaining the SYNC message 2-255: Reserved

If this TLV is absent, the CM MUST wait for a SYNC message on the new channel before proceeding. If the CM has to wait for a new SYNC message when changing channels, then operation may be suspended for a time up to the "SYNC interval" (see Annex B) or longer, if the SYNC message is lost or is not synchronized with the old channel(s).

The CM MUST observe this TLV.

8.3.20.1.3 Initialization technique

When present, this TLV allows the CMTS to direct the CM as to what level of re-initialization, if any, it MUST perform before it can commence communications on the new channel(s). The CMTS can make this decision based upon its knowledge of the differences between the old and new MAC domains and the PHY characteristics of their upstream and downstream channels.

Typically, if the move is between upstream and/or downstream channels within the same MAC domain, then the connection profile values may be left intact. If the move is between different MAC domains, then a complete initialization may be performed.

If a complete re-initialization is not required, some re-ranging may still be required. For example, areas of upstream spectrum are often configured in groups. A DCC-REQ to an adjacent upstream channel within a group may not warrant re-ranging. Alternatively, a DCC-REQ to a non-adjacent upstream channel might require unicast initial ranging whereas a DCC-REQ from one upstream channel group to another might require broadcast initial ranging. Re-ranging may also be required if there is any difference in the PHY parameters between the old and new channels.

Subtype	Length	Value
3	1	0: Re-initialize the MAC 1: Perform broadcast initial ranging on new channel before normal operation 2: Perform unicast initial ranging on new channel before normal operation 3: Perform either broadcast initial ranging or unicast initial ranging on new channel before normal operation 4: Use the new channel(s) directly without re-initializing or ranging 5-255: Reserved

The CM MUST first select the new upstream and downstream channels based upon the upstream channel ID TLV (refer to clause 8.3.20.1.1) and the downstream frequency TLV (refer to clause 8.3.20.1.2.1). The CM MUST follow the directives of these TLVs. For Option 0, the CM MUST begin with the initialization SID. For Options 1 to 4 the CM MUST continue to use the primary SID for ranging. A SID substitution TLV may specify a new primary SID for use on the new channel (refer to clause 8.3.20.1.6.2).

- **Option 0**: This option directs the CM to perform all the operations associated with initializing the CM (refer to clause 11.2). This includes all the events after acquiring downstream QAM, FEC and MPEG lock and before standard operation (refer to clause 11.3), including obtaining a UCD, ranging, establishing IP connectivity, establishing time of day, transfer of operational parameters, registration and baseline privacy initialization. When this option is used, the only other TLVs in DCC-REQ that are relevant are the upstream channel ID TLV and the downstream parameters TLV. All other DCC-REQ TLVs are irrelevant.
- **Option 1**: If broadcast initial ranging is specified, operation on the new channel could be delayed by several ranging intervals (see Annex B).
- **Option 2**: If unicast initial ranging is specified, operation on the new channel could be delayed by the value of T4 (see Annex B).
- **Option 3**: This value authorizes a CM to use an initial maintenance or station maintenance region, whichever the CM selects. This value might be used when there is uncertainty when the CM may execute the DCC command and thus a chance that it might miss station maintenance slots.
- **Option 4**: This option provides for the least interruption of service, and the CM may continue its normal operation as soon as it has achieved synchronization on the new channel.

If this TLV is absent, the CM MUST re-initialize the MAC. The CMTS MAY include this TLV. The CM MUST observe this TLV.

8.3.20.1.4 UCD substitution

When present, this TLV allows the CMTS to send an upstream channel descriptor message to the CM. This UCD message is intended to be associated with the new upstream and/or downstream channel(s). The CM stores this UCD message in its cache, and uses it after synchronizing to the new channel(s).

Type	Length	Value
4	n	UCD for the new upstream channel

This TLV includes all parameters for the UCD message as described in clause 8.3.3 except for the MAC management message header. The CMTS MUST ensure that the change count in the UCD matches the change count in the UCDs of the new channel(s). The CMTS SHOULD ensure that the upstream channel ID for the new channel is different from the upstream channel ID for the old channel. If the upstream channel IDs for the old and new channels are identical, the CMTS MUST include this TLV. The ranging required parameter in the new UCD does not apply in this context, since the functionality is covered by the initialization technique TLV.

If the length of the UCD exceeds 254 bytes, the UCD MUST be fragmented into two or more successive type 4 elements. Each fragment, except the last, MUST be 254 bytes in length. The CM reconstructs the UCD substitution by concatenating the contents (value of the TLV) of successive type 4 elements in the order in which they appear in the DCC-REQ message. For example, the first byte following the length field of the second type 4 element is treated as if it immediately follows the last byte of the first type 4 element.

If the CM has to wait for a new UCD message when changing channels, then operation may be suspended for a time up to the "UCD interval" (see Annex B) or longer, if the UCD message is lost.

The CM MUST observe this TLV, even if the upstream channel ID and the UCD change count match the old channel.

8.3.20.1.5 Security association identifier (SAID) substitution

When present, this TLV allows the CMTS to replace the security association identifier (SAID) in the current service flow with a new security association identifier. The baseline privacy keys associated with the SAID MUST remain the same. The CM does not have to simultaneously respond to the old and new SAID.

Subtype	Length	Value
6	4	Current SAID (lower-order 14 bits of a 16-bit field), new SAID (lower-order 14 bits of a 16-bit field)

If this TLV is absent, the current security association identifier assignment is retained. The CMTS MAY include this TLV. The CM MUST observe this TLV.

8.3.20.1.6 Service flow substitutions

When present, this TLV allows the CMTS to replace specific parameters within the current service flows on the current channel assignment with new parameters for the new channel assignment. One TLV is used for each service flow that requires changes in parameters. The CMTS may choose to do this to help facilitate setting up new QoS reservations on the new channel before deleting QoS reservations on the old channel. The CM does not have to simultaneously respond to the old and new service flows.

This TLV allows resource assignments and services to be moved between two independent ID value spaces and scheduling entities by changing the associated IDs and indices. ID value spaces that may differ between the two channels include the service flow identifier and the service ID. This TLV does not allow changes to service flow QoS parameters.

The service class names used within the service flow ID should remain identical between the old and new channels.

Type	Length	Value
7	n	List of subtypes

If this TLV is absent for a particular service flow, then the current service flow and its attributes are retained. The CMTS MAY include this TLV. The CM MUST observe this TLV.

8.3.20.1.6.1 Service flow identifier substitution

This TLV allows the CMTS to replace the current service flow identifier (SFID) with a new service flow identifier. Refer to clause C.2.2.3.2 for usage details.

This TLV MUST be present if any other service flow subtype substitutions are made. If this TLV is included and the service flow ID is not changing, then the current and new service flow ID will be set to the same value.

Subtype	Length	Value
7.1	8	Current service flow ID, new service flow ID

The CMTS MUST include this Sub-TLV. The CM MUST observe this Sub-TLV.

8.3.20.1.6.2 Service identifier substitution

When present, this TLV allows the CMTS to replace the service identifier (SID) in the current upstream service flow with a new service identifier. Refer to clause C.2.2.3.3 for usage details.

Subtype	Length	Value
7.2	4	Current SID (lower-order 14 bits of a 16-bit field), new SID (lower-order 14 bits of a 16-bit field)

If this TLV is absent, the current service identifier assignments are retained. The CMTS MAY include this TLV. The CM MUST observe this TLV.

8.3.20.1.6.3 Unsolicited grant time reference substitution

When present, this TLV allows the CMTS to replace the current unsolicited grant time reference with a new unsolicited grant time reference. Refer to clause C.2.2.6.11 for usage details.

This TLV is useful if the old and new upstream use different time bases for their timestamps. This TLV is also useful if the unsolicited grant transmission window is moved to a different point in time. Changing this value may cause operation to temporarily exceed the jitter window specified by clause C.2.2.6.8.

Subtype	Length	Value
7.5	4	New reference

If this TLV is absent, the current unsolicited grant time reference is retained. The CMTS MAY include this TLV. The CM MUST observe this TLV.

8.3.20.1.7 CMTS MAC address

When present, this TLV allows the current CMTS to send the MAC address of the destination CMTS corresponding to the target downstream frequency.

Type	Length	Value
8	6	MAC address of destination CMTS

The CMTS MUST include this TLV if the CM is changing downstream channels and UCD substitution is specified or if the CM is changing downstream channels and using initialization technique 4. The CM SHOULD observe this TLV.

8.3.21 Dynamic channel change response (DCC-RSP)

A dynamic channel change response MUST be transmitted by a CM in response to a received dynamic channel change request message to indicate that it has received and is complying with the DCC-REQ. The format of a DCC-RSP message MUST be as shown in Figure 8-39.

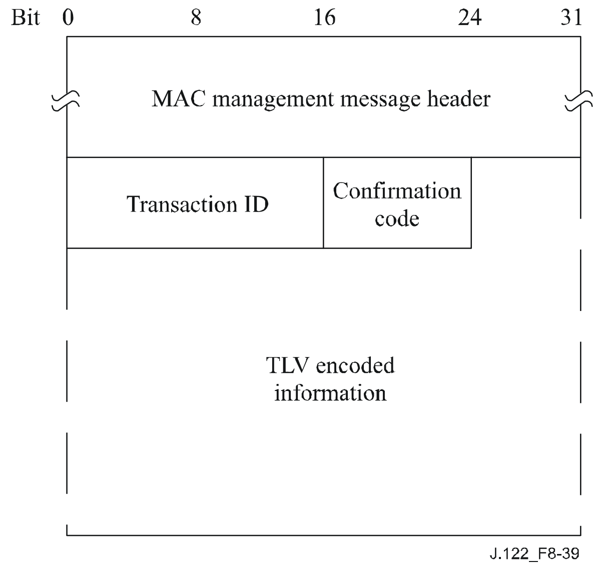

Figure 8-39 – Dynamic channel change response

Before it begins to switch to a new upstream or downstream channel, a CM MUST transmit a DCC-RSP on its existing upstream channel. When a CM receives a DCC-REQ message requesting that it switch to an upstream and downstream channel that it is already using or requesting that it switch to only an upstream or downstream channel that it is already using, the CM MUST respond with a DCC-RSP message on that channel indicating that it is already using the correct channel.

A CM MAY ignore a DCC-REQ message while it is in the process of performing a channel change.

After switching to a new channel, if the MAC was not re-initialized per DCC-REQ initialization TLV, option 0, the CM MUST send a DCC-RSP message to the CMTS. A DCC-RSP MUST NOT be sent if the CM re-initializes its MAC.

The full procedure for changing channels is described in clause 11.4.5.

Parameters MUST be as follows:

- **Transaction ID**: A 16-bit transaction ID from corresponding DCC-REQ.
- **Confirmation code**: An 8-bit confirmation code as described in clause C.4.1.

The following parameters are optional and are coded as TLV tuples:

- **CM jump time**: Timing parameters describing when the CM will make the jump.

Regardless of success or failure, if privacy is enabled for the CM, the DCC-RSP MUST contain:

- **Key sequence number**: The key sequence number of the auth key, which is used to calculate the HMAC-Digest (refer to clause C.1.4.3).
- **HMAC-Digest**: The HMAC-Digest attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest attribute MUST be the final attribute in the dynamic channel change message's attribute list (refer to clause C.1.4.1).

8.3.21.1 Encodings

The type values used MUST be those shown below. These are unique within the dynamic channel change response message, but not across the entire MAC message set.

8.3.21.1.1 CM jump time

When present, this TLV allows the CM to indicate to the CMTS when the CM plans to perform its jump and be disconnected from the network. With this information, the CMTS MAY take preventative measures to minimize or to eliminate packet drops in the downstream due to the channel change.

Type	Length	Value
1	n	

The time reference and units of time for these sub-TLVs is based upon the same 32-bit time base used in the SYNC message on the current downstream channel. This timestamp is incremented by a 10.24-MHz clock.

The CM SHOULD include this TLV. The CMTS SHOULD observe this TLV.

8.3.21.1.1.1 Length of jump

This TLV indicates to the CMTS the length of the jump from the previous channel to the new channel. Specifically, it represents the length of time that the CM will not be able to receive data in the downstream.

Subtype	Length	Value
1	4	Length (based upon timestamp)

The CM MUST include this sub-TLV.

8.3.21.1.1.2 Start time of jump

When present, this TLV indicates to the CMTS the time in the future that the CM is planning on making the jump.

Subtype	Length	Value
2	8	Start time (based upon timestamp), accuracy of start time (based upon timestamp)

The 32-bit, 10.24 MHz time base rolls over approximately every 7 minutes. If the value of the start time is less than the current timestamp, the CMTS will assume one roll-over of the timestamp counter has elapsed. The accuracy of the start time is an absolute amount of time before and after the start time.

The potential jump window is from (start time – accuracy) to (start time + accuracy + length).

The CM SHOULD include this TLV.

8.3.22 Dynamic channel change acknowledge (DCC-ACK)

A dynamic channel change acknowledge MUST be transmitted by a CMTS in response to a received dynamic channel change response message on the new channel with its confirmation code set to arrive(181). The format of a DCC-ACK message MUST be as shown in Figure 8-40.

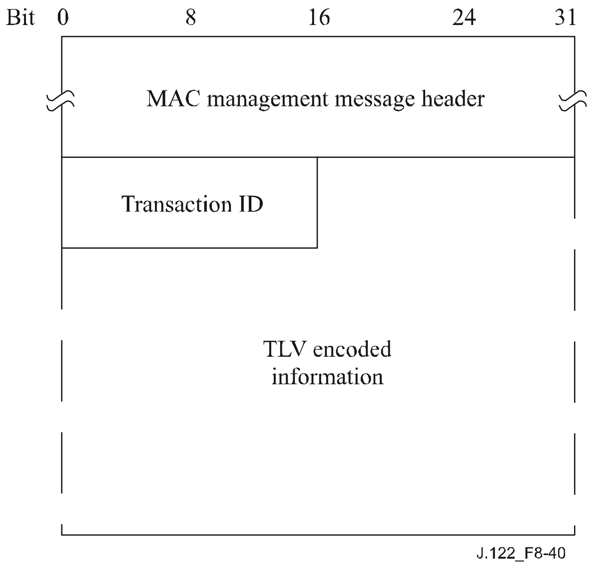

Figure 8-40 – Dynamic channel change acknowledge

Parameters MUST be as follows:

- **Transaction ID**: A 16-bit transaction ID from corresponding DCC-RSP.

If privacy is enabled, the DCC-ACK message MUST contain:

- **Key sequence number**: The key sequence number of the auth key, which is used to calculate the HMAC-Digest (refer to clause C.1.4.3).
- **HMAC-Digest**: The HMAC-Digest attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest attribute MUST be the final attribute in the dynamic channel change message's attribute list (refer to clause C.1.4.1).

8.3.23 Device class identification request (DCI-REQ)

A CM MAY implement the DCI-REQ message.

When implemented, a CM MUST transmit a DCI-REQ immediately following receipt of a ranging complete indication from the CMTS. A CM MUST NOT continue with initialization until a DCI-RSP message is received from the CMTS. Timeout and retry information is provided in Annex B, "parameters and constants".

The DCI-REQ MUST be formatted as shown in Figure 8-41.

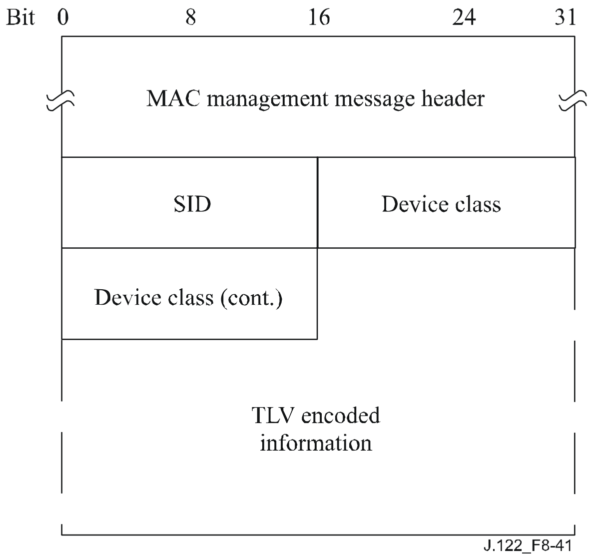

Figure 8-41 – Device class identification request

Parameters MUST be as follows:

- **SID**: The temporary SID assigned during ranging.
- **Device class**: This is a 32-bit field where individual bits represent individual attributes of the CM. Bit #0 is the LSB of the field. Bits are set to 1 to select the attributes defined below:
  - bit #0: CPE controlled cable modem (CCCM).
  - bits #1-31: Reserved and must be set to zero.

8.3.24 Device class identification response (DCI-RSP)

A DCI-RSP MUST be transmitted by a CMTS in response to a received DCI-REQ.

The DCI-RSP MUST be formatted as shown in Figure 8-42.

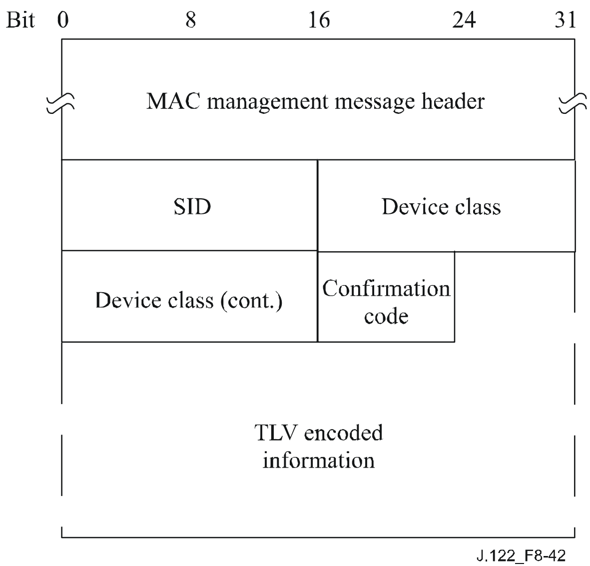

Figure 8-42 – Device class identification response

Parameters MUST be as follows:

- **SID**: The SID received in the associated DCI-REQ.
- **Device class**: The device class field as received in the associated DCI-REQ.
- **Confirmation code**: Refer to clause C.4.

The CMTS MUST use only one of three confirmation codes in the DCI-RSP.

– If the response is reject-temporary(3), the CM MUST reset its DCI-REQ retry counter to zero and MUST resend the DCI-REQ and wait for the DCI-RSP before proceeding.

– If the response is reject-permanent(4), the CM MUST abort this registration attempt and MUST begin re-scanning for a different downstream channel. The CM MUST NOT retry this channel until it has tried all other DOCS downstream channels on the network.

– If the response is success(0), the CM MUST continue with registration.

The CMTS MUST retain the device class information for use in the DHCP process. The CMTS MUST create a DHCP agent option 82 tuple with the device class information and MUST insert this tuple in the DHCPDISCOVER from the corresponding CM before forwarding that DHCPDISCOVER to the DHCP server.

8.3.25 Upstream transmitter disable (UP-DIS) MAC management message

The UP-DIS message provides additional functionality to permanently or temporarily disable the modem, as well as to disable the modem for a specified period of time. It is used to control the admission of certain modem types and groups to the network as early as immediately before registration. It can also be used for network troubleshooting, disabling modems that violate network policies, or for avoiding request floods in a large network when the CMTS goes on-line.

This message is stateless and can be issued by the CMTS at any time. The UP-DIS message is sent from a CMTS to a CM; there is no response from the CM transmitted back to the CMTS. UP-DIS messages may be unicast, in which case the destination address in the MAC header is the address of the selected CM; or multicast, in which case the destination address is a well-known MAC multicast address (see Annex A for details on well-known addresses).

The CMTS MUST be capable of transmitting the UP-DIS message. The CMTS can transmit UP-DIS messages either as a result of a triggering event detected by the CMTS internally, or in response to a remote management command. Mechanisms for setting up, detecting and reporting situations where the transmission of an UP-DIS message might be appropriate are implementation-dependent. Similarly, signalling, which remotely instructs the CMTS to transmit a UP-DIS message, is outside the scope of this Recommendation. One of the possible implementations may be SNMP commands sent to the CMTS over the network.

CMs SHOULD support the UP-DIS message in order to ease network management.

Since the UP-DIS mechanism at the CM is stateless and the CMs do not retain disabled status after a power cycle, the CMTS MAY incorporate mechanisms to track disabled CMs by their MAC addresses. The CMTS would resend an UP-DIS message as appropriate to the modems that were permanently disabled by the network operator, and then power-cycled by the user in an attempt to re-register. However, the same function may also be implemented by the provisioning infrastructure on modem registration; therefore, if the CMTS is unable to track disabled modems autonomously, it SHOULD be able to send a UP-DIS in response to an external command.

The UP-DIS message MUST be formatted as shown in Figure 8-43.

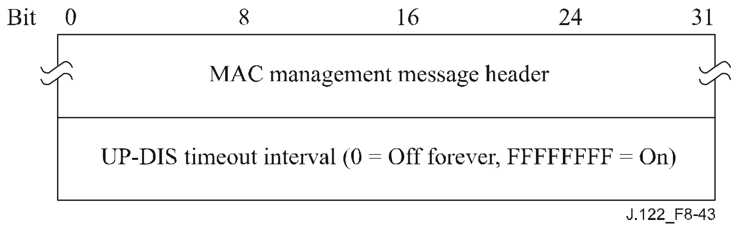

Figure 8-43 – UP-DIS message format

The only parameter is UP-DIS timeout interval, which MUST be encoded as follows.

UP-DIS timeout interval is a 32-bit, unsigned integer representing the disable timeout interval in milliseconds. There are two special values defined:

– 00000000 permanently disables the upstream of the modem, as described below.

– FFFFFFFF remotely re-initializes the MAC, which resumes the normal operation of the modem.

The CM MUST autonomously disable its upstream transmitter immediately upon receipt of an UP-DIS message with UP-DIS timeout interval = 0, regardless of any other transaction state (refer to clause 11), or the state of its control programme. The modem stops all transmissions, but continues to listen to the MAC messages sent in the downstream. Once disabled, the CM upstream transmitter MUST only be re-enabled by power cycling the CM, or by an UP-DIS message with UP-DIS timeout interval = FFFFFFFF. All other UP-DIS messages MUST be ignored when the upstream is disabled.

If supported, the CM MUST autonomously reset its upstream transmitter upon receipt of an UP-DIS message with UP-DIS time-out interval = FFFFFFFF, regardless of any other transaction state (refer to clause 11), or the state of its control program. Resetting allows the modem to resume transmissions.

Additional non-zero timeout values in the UP-DIS message SHOULD be supported. If supported, the CM MUST autonomously disable its upstream transmitter immediately upon receipt of an UP-DIS message with UP-DIS timeout interval T > 0 for a period of T ms, regardless of any other transaction state (refer to clause 11), or the state of its control program. Although the timeout T is specified in ms, the CM MAY extend the specified timeout by up to 100 ms. When timeout expires, the CM SHOULD re-initialize the MAC as appropriate, starting with the initial ranging process and registration, because there is no guarantee that the CMTS has not de-registered it. In the disabled state, all other UP-DIS messages MUST be ignored, except for an UP-DIS message with UP-DIS timeout interval = FFFFFFFF or 00000000.

8.3.26 Initial ranging request (INIT-RNG-REQ)

A ranging request MUST be transmitted by a CM at initialization to determine network delay and request power adjustment. This message MUST use an FC_TYPE = MAC-specific header and FC_PARM = timing MAC header. This MUST be followed by a packet PDU in the format shown in Figure 8-44.

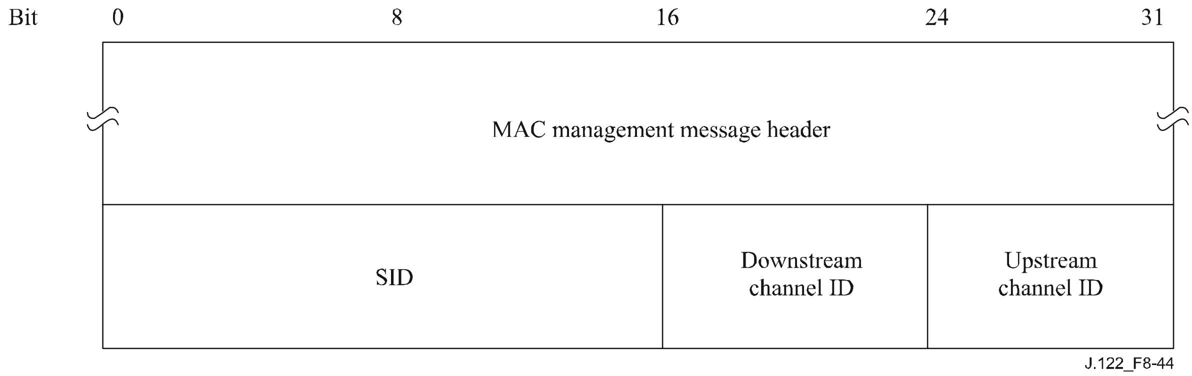

Figure 8-44 – Packet PDU following the timing header

The INIT-RNG-REQ differs from the RNG-REQ in that it is only transmitted in broadcast initial ranging opportunities and MUST NOT be transmitted on a logical upstream that is not a DOCS 2.0-only upstream. It also has an upstream channel ID in place of the pending till complete field in a RNG-REQ.

Parameters MUST be as follows:

- **SID**:
  - Initialization SID if modem is attempting to join the network.
  - Initialization SID if modem has not yet registered and is changing downstream, or both downstream and upstream channels, as directed by a downloaded parameter file.
  - Primary SID (previously assigned in REG-RSP) if modem is registered and is changing upstream channels; if the CM is redoing initial ranging as a result of a DCC, UCC, or UCD change (see clauses 8.3.3 and 11.3.2); or if the CM is redoing initial ranging as a result of a temporary loss of downstream signal, while in S-CDMA mode (see clause 11.2.1).

This is a 16-bit field of which the lower 14 bits define the SID with bits 14 and 15 defined to be 0.

- **Downstream channel ID**: The identifier of the downstream channel on which the CM received the UCD which described this upstream. This is an 8-bit field.
- **Upstream channel ID**: The upstream channel ID from the UCD the CM is using to transmit this INIT-RNG-REQ. In the case where multiple logical upstreams are sharing the same spectrum, and the broadcast initial ranging opportunities of some of these logical channels are aligned, this allows the CMTS to know which logical channel the CM is using.

A CM MUST set the RSVD field of the MAC management message header to report support of the S-CDMA maximum scheduled codes if, and only if, the CMTS indicated that it supports the S-CDMA maximum scheduled codes from TLV-17 of the UCD for the upstream channel on which the CM is ranging. In this case, the CM MUST report the maximum ratio of number of active codes to maximum scheduled codes that the CM can support; this ratio is indicated by a bit mask in the reserved field as shown below. For example, if the number of active codes on the channel is 128 and the CM supports a minimum of 64 scheduled codes (the minimum number of allowed active codes), the CM would report a ratio of 2. The CMTS will use this value in calculating an appropriate value for maximum scheduled codes to assign to the CM. The CM SHOULD support a maximum ratio of 32.

When the CM reports support for the S-CDMA maximum scheduled codes, the CM MUST also report its current transmit power shortfall (in dB). The CM power shortfall is the difference between the current target transmit power of the ranging request and the maximum S-CDMA spreader-on transmit power of 53 dBmV. The CM MUST report a power shortfall of 0 if the current target transmit power of the ranging request is less than or equal to 53 dBmV. This value will be used by the CMTS for calculating appropriate values for S-CDMA maximum scheduled codes and S-CDMA power headroom for the CM.

The format of the RSVD field is:

Bit 7: 1= S-CDMA maximum scheduled codes supported	Bits 6 to 5: CM maximum ratio of: <Anchor6> 00 = 2 01 = 8 10 = 16 11 = 32	Bits 4 to 0: CM power shortfall (1/4 dB)

8.3.27 Test request (TST-REQ)

The test request is used to force a CM to enter or leave one of two test modes. The TST-REQ message with Mode!= 0 MUST NOT be sent by the CMTS except in response to an explicit command from the operator.

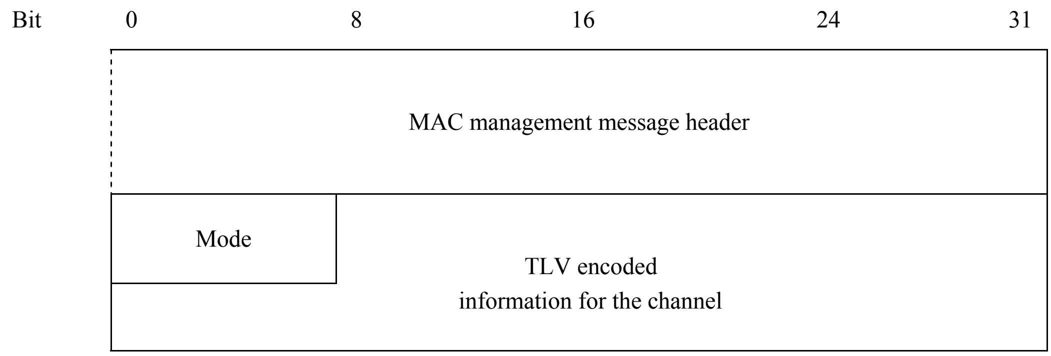

Figure 8-45 – Test request

Parameters MUST be as follows:

- **Mode**:
  - 0 = Disable all test modes and reboot.
  - 1 = Transmit a continuous (non-bursted) upstream signal at the commanded modulation rate, carrier frequency and power level. The chip sequence at the spreader output is replaced with an alternating binary sequence (1, –1, 1, –1, 1, –1, ...) at nominal amplitude, equal on I and Q. The CM tracks the downstream symbol clock and uses it to generate the upstream symbol clock as in normal synchronous operation.
  - 2 = Transmit a continuous (non-bursted), unmodulated (CW) upstream signal at the commanded carrier frequency, modulation rate and power 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.

In normal operation, the CM MUST ignore any TST-REQ message it receives subsequent to receiving the first SYNC message. Note that this makes it less convenient to use this test mode with a CMTS, since the CMTS may send a SYNC message before the CM sees a TST-REQ message.

After acquiring a downstream signal, and prior to receiving a SYNC message, if the CM receives a TST-REQ message (either unicast to the CM itself, or broadcast) with Mode != 0, the CM MUST begin the test mode indicated in the mode parameter, using the channel parameters included in the TST-REQ message.

In test mode, if the CM receives a TST-REQ message (either unicast to the CM itself, or broadcast) with Mode = 0, the CM MUST reboot. The CM MUST reboot after the expiration of the T16 timer, a 30 minute test mode timer.

The TST-REQ message MUST be generated in the format shown in Figure 8-45, including all of the parameters coded as TLV-tuples defined in Table 8-22.

Table 8-22 – Channel TLV parameters

Name	Type (1 byte)	Length (1 byte)	Value (variable length)
Modulation rate	1	1	Multiples of base rate of 160 kHz (value is 1, 2, 4, 8, 16 or 32). This TLV MUST be present if the test mode is 1.
Frequency	2	4	Upstream carrier frequency (Hz). This TLV MUST be present if the test mode is 1 or 2.
Power	3	1	This TLV specifies the power (unsigned 8-bit, dBmV units) at which the CM MUST transmit the TST-REQ message. This TLV MUST be present if the test mode is 1 or 2.
S-CDMA US ratio numerator 'M'	4	2	The numerator (M) of the M/N ratio relating the downstream symbol clock to the upstream modulation clock. This TLV MUST be present if the test mode is 1. This TLV MUST be present if the test mode is 2 and the operation is synchronous.
S-CDMA US ratio denominator 'N'	5	2	The denominator (N) of the M/N ratio relating the downstream symbol clock to the upstream modulation clock. This TLV MUST be present if the test mode is 1. This TLV MUST be present if the test mode is 2 and the operation is synchronous.

9 Media access control protocol operation

9.1 Upstream bandwidth allocation

The upstream channel is modelled as a stream of mini-slots. The CMTS MUST generate the time reference for identifying these slots. It MUST also control access to these slots by the cable modems. For example, it MAY grant some number of contiguous slots to a CM for it to transmit a data PDU. The CM MUST time its transmission so that the CMTS receives it in the time reference specified. This clause describes the elements of protocol used in requesting, granting and using upstream bandwidth. The basic mechanism for assigning bandwidth management is the allocation MAP. Please refer to Figure 9-1.

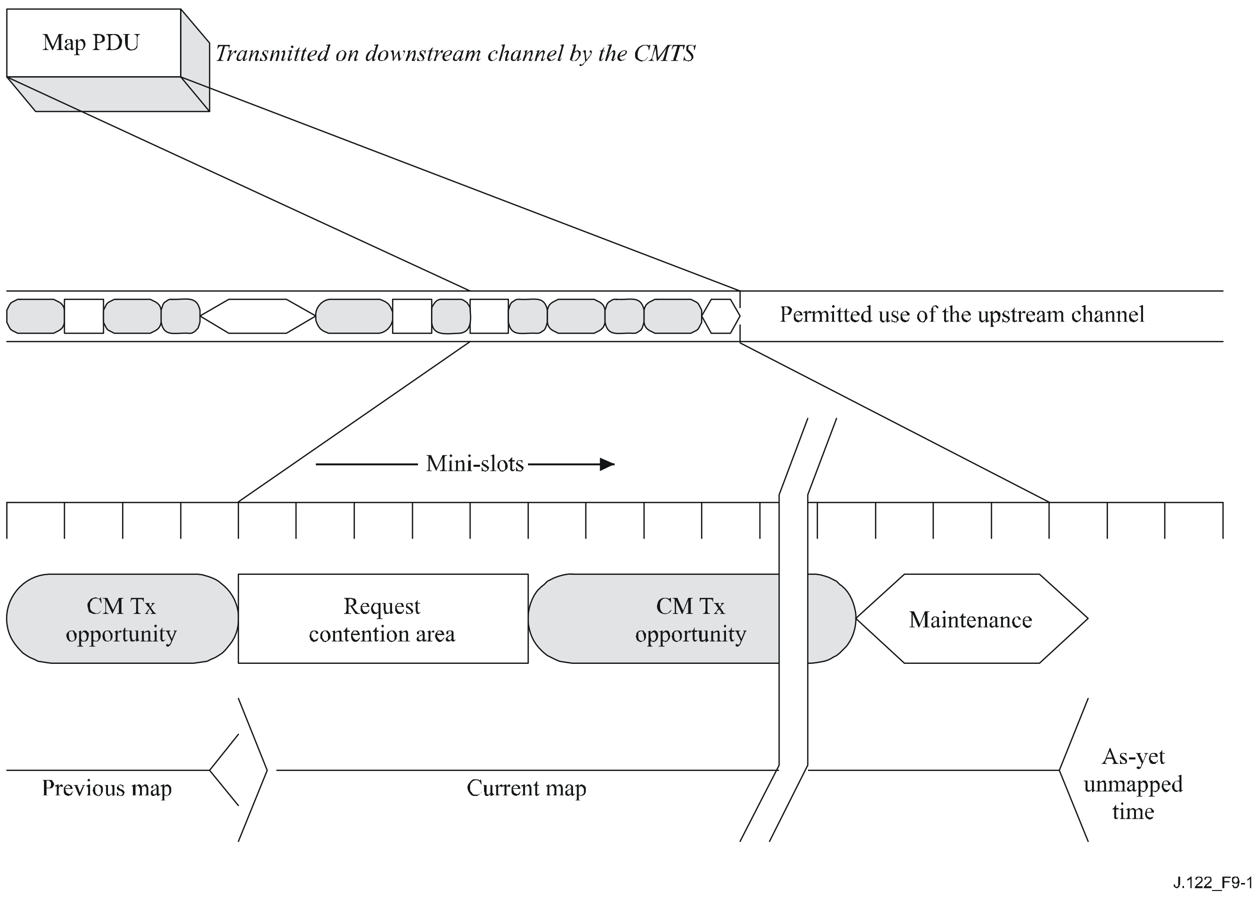

Figure 9-1 – Allocation MAP

The allocation MAP is a MAC management message transmitted by the CMTS on the downstream channel which describes, for some interval, the uses to which the upstream mini-slots MUST be put. A given MAP MAY describe some slots as grants for particular stations to transmit data in, other slots as available for contention transmission, and other slots as an opportunity for new stations to join the link.

Many different scheduling algorithms MAY be implemented in the CMTS by different vendors; this Recommendation does not mandate a particular algorithm. Instead, it describes the protocol elements by which bandwidth is requested and granted.

The bandwidth allocation includes the following basic elements:

- Each CM has one or more short (14-bit) service identifiers (SIDs) as well as a 48-bit address.
- Upstream bandwidth is divided into a stream of mini-slots. Each mini-slot is numbered relative to a master reference maintained by the CMTS. The master reference is distributed to the CMs by means of SYNC and UCD packets (see clause 6.2.11.2).
- CMs may issue requests to the CMTS for upstream bandwidth.

The CMTS MUST transmit allocation MAP PDUs on the downstream channel defining the allowed usage of each mini-slot. The MAP is described below.

9.1.1 The allocation MAP MAC management message

The allocation MAP is a varying-length MAC management message that is transmitted by the CMTS to define transmission opportunities on the upstream channel. It includes a fixed-length header followed by a variable number of information elements (IEs) in the format shown in clause 8.3.4. Each information element defines the allowed usage for a range of mini-slots.

NOTE – For TDMA channels, it should be understood by both CM and CMTS that the lower (26-M) bits of alloc start and ack times MUST be used as the effective MAP start and ack times, where M is defined in clause 8.3.3. The relationship between alloc start/ack time counters and the timestamp counter is further described in clause 9.3.4. For DOCS 2.0 S-CDMA channels, the alloc start/ack time counters are defined in mini-slots that are related to the timestamp counter, frame counter and S-CDMA timestamp snapshot as described in clause 6.2.11.2.

9.1.2 Information elements

Each IE consists of a 14-bit service ID, a 4-bit type code, and a 14-bit starting offset as defined in clause 8.3.4. Since all stations MUST scan all IEs, it is critical that IEs be short and relatively fixed format. IEs within the MAP are strictly ordered by starting offset. For most purposes, the duration described by the IE is inferred by the difference between the IE's starting offset and that of the following IE. For this reason, a null IE MUST terminate the list. Refer to Table 8-20.

Five types of service IDs are defined:

1) 0x3FFF – Broadcast; intended for all stations.
2) 0x3E00-0x3FFE – Multicast; purpose is defined administratively. Refer to Annex A.
3) 0x2000-0x3DFF – Expanded unicast; intended for a particular CM or a particular service within that CM, when supported by both the CM and CMTS.
4) 0x0001-0x1FFF – Unicast; intended for a particular CM or a particular service within that CM.
5) 0x0000 – Null address; addressed to no station.

A CM MUST support the expanded unicast SID space. A CMTS MAY support the expanded unicast SID space.

Unicast SIDs (including expanded unicast SIDs) MUST be unique on a given logical upstream. The CMTS MAY support unicast SIDs assignments which overlap across multiple logical upstreams within a single MAC-sublayer domain.

All of the information elements defined below MUST be supported by conformant CMs. Conformant CMTSs MAY use any of these information elements when creating bandwidth allocation MAPs.

9.1.2.1 The request IE

The request IE provides an upstream interval in which requests MAY be made for bandwidth for upstream data transmission. The character of this IE changes depending on the class-of-service ID. If broadcast, this is an invitation for CMs to contend for requests. Clause 9.4 describes which contention transmit opportunity may be used. If unicast, this is an invitation for a particular CM to request bandwidth. Unicasts MAY be used as part of a quality of service scheduling scheme (refer to clause 10.2). Packets transmitted in this interval MUST use the request MAC frame format (refer to clause 8.2.5.3).

A small number of priority request SIDs are defined in clause A.2.3. These allow contention for request IEs to be limited to service flows of a given traffic priority (refer to clause C.2.2.5.1).

9.1.2.2 The request/data IE

The request/data IE provides an upstream interval in which requests for bandwidth or short data packets MAY be transmitted. This IE is distinguished from the request IE in that:

- It provides a means by which allocation algorithms MAY provide for "immediate" data contention under light loads, and a means by which this opportunity can be withdrawn as network loading increases.

- Multicast service IDs MUST be used to specify maximum data length, as well as allowed random starting points within the interval. For example, a particular multicast ID may specify a maximum of 64-byte data packets, with transmit opportunities every fourth slot.

A small number of well-known multicast Service IDs are defined in Annex A. Others are available for vendor-specific algorithms.

Since data packets transmitted within this interval may collide, the CMTS MUST acknowledge any that are successfully received. The data packet MUST indicate in the MAC header that a data acknowledgment is desired (see Table 8-13).

9.1.2.3 The initial maintenance IE

The initial maintenance IE, when used with the broadcast SID, provides an interval in which new stations may join the network. A long interval, equivalent to the maximum round-trip propagation delay plus the transmission time of a ranging request (RNG-REQ) message (see clause 9.3.3), MUST be provided to allow new stations to perform initial ranging. Packets transmitted in this interval MUST use the RNG-REQ or the INIT-RNG-REQ MAC management message format (refer to clauses 8.3.5 and 8.3.26).

On DOCS 2.0-only upstream channels, the initial maintenance IE MAY be used with a unicast SID. This is done to provide unicast initial maintenance opportunities in place of station maintenance opportunities at the discretion of the CMTS. This may be useful if the first unicast ranging opportunity on an S-CDMA channel needs to have spreader-off just like initial maintenance, but it is not desirable to impose the overhead of having the spreader-off on routine station maintenance. Unicast initial maintenance opportunities only need to be large enough to allow transmission of the ranging request. The CMTS MUST NOT provide unicast initial maintenance opportunities on any logical upstream which is not a DOCS 2.0-only upstream.

9.1.2.4 The station maintenance IE

The station maintenance IE provides an interval in which stations are expected to perform some aspect of routine network maintenance, such as ranging or power adjustment. The CMTS MAY request that a particular CM perform some task related to network maintenance, such as periodic transmit power adjustment. In this case, the station maintenance IE is unicast to provide upstream bandwidth in which to perform this task. Packets transmitted in this interval MUST use the RNG-REQ MAC management message format (see clause 8.3.5).

9.1.2.5 Short and long data grant IEs

The short and long data grant IEs provide an opportunity for a CM to transmit one or more upstream PDUs. These IEs are issued either in response to a request from a station, or because of an administrative policy providing some amount of bandwidth to a particular station (see class-of-service discussion below). These IEs MAY also be used with an inferred length of zero mini-slots (a zero length grant), to indicate that a request has been received and is pending (a data grant pending).

Short data grants are used with intervals less than or equal to the maximum burst size for this usage specified in the upstream channel descriptor. If short data burst profiles are defined in the UCD, then all long data grants MUST be for a larger number of mini-slots than the maximum for short data. The distinction between long and short data grants may be exploited in physical-layer forward-error-correction coding; otherwise, it is not meaningful to the bandwidth allocation process.

If this IE is a data grant pending (a zero length grant), it MUST follow the null IE. This allows cable modems to process all actual allocations first, before scanning the MAP for data grants pending and data acknowledgments.

9.1.2.6 Data acknowledge IE

The data acknowledge IE acknowledges that a data PDU was received. The CM MUST have requested this acknowledgment within the data PDU (normally this would be done for PDUs transmitted within a contention interval in order to detect collisions).

This IE MUST follow the null IE. This allows cable modems to process all actual interval allocations first, before scanning the MAP for data grants pending and data acknowledgments.

9.1.2.7 Expansion IE

The expansion IE provides for extensibility, if more than 16 code points or 32 bits are needed for future IEs.

9.1.2.8 Null IE

A null IE terminates all actual allocations in the IE list. It is used to infer a length for the last interval. All data acknowledge IEs and all data grant pending IEs (data grants with an inferred length of 0) must follow the null IE.

9.1.2.9 Advanced PHY short and long data grant IEs

These IEs are the advanced PHY channel equivalent of the short and long data grant IEs in clause 9.1.2.5. In addition, these IEs allow DOCS 2.0 modems operating in DOCS 2.0 TDMA mode to share the same upstream channel with DOCS 1.x modems. Modems registered in DOCS 1.x mode MUST NOT use these intervals.

For upstream channels supporting a mixture of DOCS 1.x and DOCS 2.0 TDMA CMs, the CMTS MUST use the SID in the request and the operational state of the CM to distinguish between requests for IUC 5 and 6 data grants and requests for IUC 9 and 10 data grants (refer to clause 11.2.9 for additional information). Once this distinction has been made, the CMTS then uses the request size to distinguish between a long grant and a short grant.

Once a CMTS has received a REG_ACK from a 2.0 CM on a type 2 channel, the CMTS MUST NOT send data grants using IUCs 5 or 6 if either IUC 9 or 10 is defined for that upstream channel. This restriction allows the CM to support only 7 burst profiles simultaneously.

9.1.2.10 Advanced PHY unsolicited grant IE

This IE can be used by the CMTS to make unsolicited grants of bandwidth to DOCS 2.0 CMs. If a significant portion of the traffic for an upstream is going to consist of unsolicited grants of a particular size, this IE provides a way for the CMTS to provide a set of physical layer parameters (such as code word length and FEC length) well tailored to that traffic, without compromising the general usefulness of the advanced phy short or advanced phy long data grant IEs. It is never used by the CM to calculate the size of a bandwidth request. The CMTS MUST NOT use it to make grants to DOCS 1.x CMs.

9.1.3 Requests

Requests refer to the mechanism that CMs use to indicate to the CMTS that it needs upstream bandwidth allocation. A Request MAY come as a stand-alone request frame transmission (refer to clause 8.2.5.3) or it MAY come as a piggyback request in the EHDR of another frame transmission (refer to clause 8.2.6).

The request frame MAY be transmitted during any of the following intervals:

- request IE;
- request/data IE;
- short data grant IE;
- long data grant IE;

- adv PHY short data grant IE;
- adv PHY long data grant IE;
- adv PHY unsolicited grant IE.

A piggyback request MAY be contained in the following extended headers:

- request EH element;
- upstream privacy EH element;
- upstream privacy EH element with fragmentation.

The request MUST include:

- the service ID making the request;
- the number of mini-slots requested.

The CM MUST request the number of mini-slots needed to transmit an entire frame, or a fragment containing the entire remaining portion of a frame that a previous grant has caused to be fragmented. The frame may be a single MAC frame, or a MAC frame that has been formed by the concatenation of multiple MAC frames (see clause 8.2.5.5). The request from the CM MUST be large enough to accommodate the entire necessary physical layer overhead (see clause 6.2) for transmitting the MAC frame or fragment. The CM MUST NOT make a request that would violate the limits on data grant sizes in the UCD message (see clause 8.3.3) or any limits established by QoS parameters associated with the service flow.

The CM MUST NOT request more mini-slots than are necessary to transmit the MAC frame. This means that if the CM is using short and long data IUCs to transmit data and the frame can fit into a short data grant, the CM MUST use the short data grant IUC attributes to calculate the amount of bandwidth to request and MUST make a request less than or equal to the short data maximum burst size. If the CM is using advanced PHY short and long data IUCs to transmit data and the frame can fit into an advanced PHY short data grant, the CM MUST use the advanced PHY short data grant IUC attributes to calculate the amount of bandwidth to request and MUST make a request less than or equal to the advanced PHY short data maximum burst size.

The CM MUST have only one request outstanding at a time per service ID. If the CMTS does not immediately respond with a data grant, the CM is able to unambiguously determine that its request is still pending because the CMTS MUST continue to issue a data grant pending in every MAP that has an ack time indicating the request has already been processed until the request is granted or discarded.

In MAPs, the CMTS MUST NOT make a data grant greater than 255 mini-slots to any assigned service ID. This puts an upper bound on the grant size the CM has to support.

9.1.4 Information element feature usage summary

Table 9-1 summarizes what types of frames the CM can transmit using each of the MAP IE types that represent transmit opportunities. A "MUST" entry in the table means that, if appropriate, a compliant CM implementation has to be able to transmit that type of frame in that type of opportunity. A "MAY" entry means that compliant CM implementation does not have to be able to transmit that type of frame in that type of opportunity but that it is legal for it to do so, if appropriate. A "MUST NOT" entry means that a compliant CM will never transmit that type of frame in that type of opportunity.

Table 9-1 – IE feature compatibility summary

Information element	Transmit request frame	Transmit concatenated MAC frame	Transmit fragmented MAC frame	Transmit RNG-REQ	Transmit any other MAC frame
Request IE	MUST	MUST NOT	MUST NOT	MUST NOT	MUST NOT
Request/data IE	MUST	MAY	MUST NOT	MUST NOT	MAY
Initial maintenance IE	MUST NOT	MUST NOT	MUST NOT	MUST	MUST NOT
Station maintenance IE	MUST NOT	MUST NOT	MUST NOT	MUST	MUST NOT
Short data grant IE	MAY	MUST	MUST	MUST NOT	MUST
Long data grant IE	MAY	MUST	MUST	MUST NOT	MUST
Adv PHY short data grant IE	MAY	MUST	MUST	MUST NOT	MUST
Adv PHY long data grant IE	MAY	MUST	MUST	MUST NOT	MUST
Adv PHY unsolicited grant IE	MAY	MUST	MUST	MUST NOT	MUST

9.1.5 MAP transmission and timing

The allocation MAP MUST be transmitted in time to propagate across the physical cable and be received and handled by the receiving CMs. As such, it MAY be transmitted considerably earlier than its effective time. The components of the delay are:

- worst-case round-trip propagation delay – may be network-specific, but on the order of hundreds of microseconds;
- queuing delays within the CMTS – implementation-specific;
- processing delays within the CMs – MUST allow a minimum processing time by each CM as specified in Annex B (CM MAP processing time), which has to include any upstream delays caused by the DOCS 2.0 TDMA;
- downstream delays caused by the PMD-layer framer and the FEC interleaver.

Within these constraints, vendors may wish to minimize this delay so as to minimize latency of access to the upstream channel.

The number of mini-slots described MAY vary from MAP to MAP. At minimum, a MAP MAY describe a single mini-slot. This would be wasteful in both downstream bandwidth and in processing time within the CMs. At maximum, a MAP MAY stretch to tens of milliseconds. Such a MAP would provide poor upstream latency. Allocation algorithms MAY vary the size of the maps over time to provide a balance of network utilization and latency under varying traffic loads.

At minimum, a MAP MUST contain two information elements: one to describe an interval, and a null IE to terminate the list. At a maximum, a MAP MUST be bounded by a limit of 240 information elements. Maps are also bounded in that they MUST NOT describe more than 4096 mini-slots into the future. The latter limit is intended to bound the number of future mini-slots that each CM is required to track. A CM MUST be able to support multiple outstanding MAPs. Even though multiple MAPs may be outstanding, the sum of the number of mini-slots they describe MUST NOT exceed 4096.

The set of all MAPs, taken together, MUST describe every mini-slot in the upstream channel. If a CM fails to receive a MAP describing a particular interval, it MUST NOT transmit during that interval.

9.1.6 Protocol example

This clause illustrates the interchange between the CM and the CMTS when the CM has data to transmit (Figure 9-2). Suppose a given CM has a data PDU available for transmission.

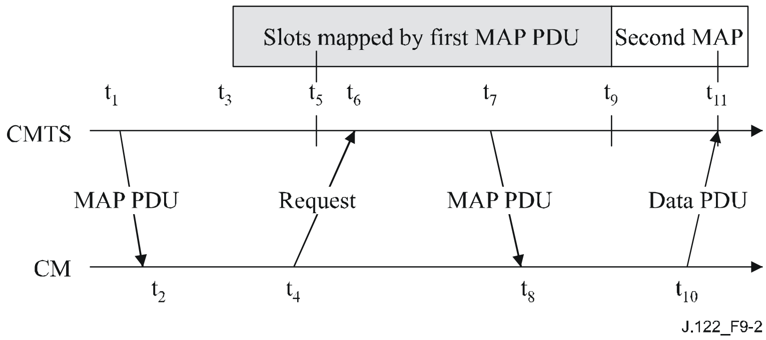

Figure 9-2 – Protocol example

Description:

1) At time t_1, the CMTS transmits a MAP whose effective starting time is t_3. Within this MAP is a request IE which will start at t_5. The difference between t_1 and t_3 is needed to allow for all the delays discussed in clause 9.1.5.
2) At t_2, the CM receives this MAP and scans it for request opportunities. In order to minimize request collisions, it calculates t_6 as a random offset based on the data backoff start value in the most recent MAP (see clause 9.4; also, the multicast SID definitions in clause A.2).
3) At t_4, the CM transmits a request for as many mini-slots as needed to accommodate the PDU. Time t_4 is chosen based on the ranging offset (see clause 9.3.3) so that the request will arrive at the CMTS at t_6.
4) At t_6, the CMTS receives the request and schedules it for service in the next MAP (the choice of which requests to grant will vary with the class of service requested, any competing requests, and the algorithm used by the CMTS).
5) At t_7, the CMTS transmits a MAP whose effective starting time is t_9. Within this MAP, a data grant for the CM will start at t_{11}.
6) At t_8, the CM receives the MAP and scans for its data grant.
7) At t_{10}, the CM transmits its data PDU so that it will arrive at the CMTS at t_{11}. Time t_{10} is calculated from the ranging offset as in step 3.

Steps 1 and 2 need not contribute to access latency if CMs routinely maintain a list of request opportunities.

At Step 3, the request may collide with requests from other CMs and be lost. The CMTS does not directly detect the collision. The CM determines that a collision (or other reception failure) occurred when the next MAP with an ack time indicating that the request would have been received and processed fails to include an acknowledgment of the request. The CM MUST then perform a back-off algorithm and retry (refer to clause 9.4.1).

At Step 4, the CMTS scheduler MAY fail to accommodate the request within the next MAP. If so, it MUST reply with a zero-length grant in that MAP or discard the request by giving no grant at all. It MUST continue to report this zero-length grant in all succeeding maps until the request can be granted or is discarded. This MUST signal to the CM that the request is still pending. So long as the CM is receiving a zero-length grant, it MUST NOT issue new requests for that service queue.

9.1.7 MAP generation example – Two logical upstreams

This clause illustrates the timing requirements for scheduling an S-CDMA and a TDMA logical channel on the same physical channel (Figure 9-3).

For simplicity, it is assumed that:

- the duration of the S-CDMA frames is an integral multiple of the duration of the TDMA mini-slots;
- both TDMA and S-CDMA maps begin and end on frame boundaries;
- for the duration of the example, there are no S-CDMA bursts with the spreader-off, and there are no broadcast initial ranging regions where both channels are active.

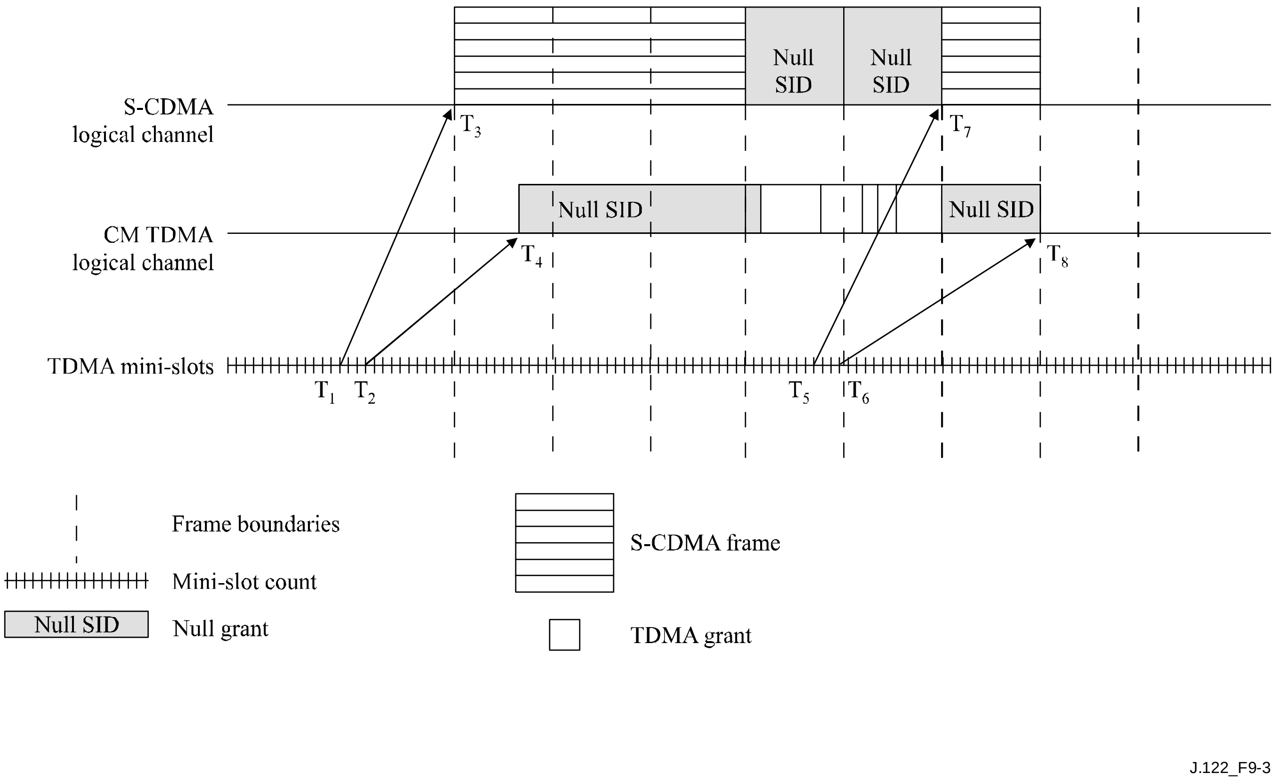

Figure 9-3 – Logical S-CDMA TDMA channels

Description:

1) The example begins at T_1 and the first MAP discussed takes effect at T_3.
2) At time T_1, the CMTS transmits a S-CDMA map whose effective starting time is T_3 and end time is T_7.
3) At time T_2, the CMTS transmits a TDMA map whose effective starting time is T_4 and end time is T_8.
4) At time T_3 the S-CDMA map has three frames of S-CDMA grants. The CMTS upstream scheduler must not allow TDMA transmissions to occur at the same time. To prevent the two channels from interfering with each other, the scheduler will mute the TDMA upstream (by granting mini-slots to the null SID for the TDMA channel) during the time S-CMDA is active.
5) At time T_4, on a frame boundary, the TDMA channel becomes active. In this example, it has one empty mini-slot (null SID) to guarantee sufficient guard time for the following

TDMA burst. Then it proceeds with usable TDMA grants. At the same time, the S-CDMA upstream is muted by granting mini-slots to the null SID in every frame.

6) At T_5 and T_6, the TDMA logical channel and S-CDMA logical channel transmit the next map for the upstream. Note that Figure 9-3 does not continue to detail the complete maps beginning at T_7 and T_8.

7) At time T_7, the S-CDMA map sends a group of S-CDMA grants in a frame. Note that when switching from TDMA to S-CDMA, there is no requirement for additional guard time.

9.2 Support for multiple channels

Vendors may choose to offer various combinations of upstream and downstream channels within one MAC service access point. The upstream bandwidth allocation protocol allows for multiple upstream channels to be managed via one or many downstream channels. Some or all of these multiple upstream channels may even co-exist on the same upstream transmit centre frequency.

If multiple upstream channels are associated with a single downstream channel, then the CMTS MUST send one allocation MAP per upstream channel. The MAP's channel identifier, taken with the upstream channel descriptor message (see clause 8.3.3), MUST specify to which channel each MAP applies. There is no requirement that the maps be synchronized across channels. Appendix III provides an example.

If multiple logical upstream channels are associated with the same upstream centre frequency on the same cable segment, then the CMTS MUST ensure that the MAP allocations to each logical upstream channel sharing the same spectrum do not coincide with the potential transmit opportunities of the other logical upstream channels with the possible exception of broadcast initial ranging opportunities. When sharing the upstream between S-CDMA and TDMA channels, the CMTS MUST take into account the lack of guard time on the synchronous physical layer upstreams. Annex G provides more information on the co-existence of DOCS 1.x and DOCS 2.0 channels.

If multiple downstream channels are associated with a single upstream channel, the CMTS MUST ensure that the allocation MAP reaches all CMs. That is, if some CMs are attached to a particular downstream channel, then the MAP MUST be transmitted on that channel. This may necessitate that multiple copies of the same MAP be transmitted. The alloc start time in the MAP header MUST always relate to the timebase reference on the downstream channel on which it is transmitted.

If multiple downstream channels are associated with multiple upstream channels, the CMTS may need to transmit multiple copies of multiple maps to ensure both that all upstream channels are mapped and that all CMs have received their needed maps.

9.3 Timing and synchronization

One of the major challenges in designing a MAC protocol for a cable network is compensating for the large delays involved. These delays are an order of magnitude larger than the transmission burst time in the upstream. To compensate for these delays, the cable modem MUST be able to time its transmissions precisely to arrive at the CMTS at the start of the assigned mini-slot.

To accomplish this, two pieces of information are needed by each cable modem:

- a global timing reference sent downstream from the CMTS to all cable modems;
- a timing offset, calculated during a ranging process, for each cable modem.

9.3.1 Global timing reference

For TDMA channels, the CMTS MUST create a global timing reference by transmitting the time synchronization (SYNC) MAC management message downstream at a nominal frequency. The message contains a timestamp that exactly identifies when the CMTS transmitted the message. Cable modems MUST then compare the actual time the message was received with the timestamp and adjust their local clock references accordingly.

For S-CDMA channels, the CMTS also creates a global timing reference by transmitting the time synchronization (SYNC) and upstream channel descriptor (UCD) MAC messages downstream at a nominal frequency (see clause 6.2.11.2).

The transmission convergence sublayer must operate closely with the MAC sublayer to provide an accurate timestamp for the SYNC message. As mentioned in the ranging clause below (see 9.3.3), the model assumes that the timing delays through the remainder of the PHY layer MUST be relatively constant with the exception of the timing offsets specified in clause 6.2.19.4 related to modulation rate changes to accommodate a legacy DOCSIS upstream receiver implementation. For TDMA, any variation in the PHY delays MUST be accounted for in the guard time of the PHY overhead.

It is intended that the nominal interval between SYNC messages be tens of milliseconds and the nominal interval between UCD messages be no more than 2 seconds. This imposes very little downstream overhead while letting cable modems acquire their global timing synchronization quickly.

9.3.2 CM channel acquisition

Any cable modem MUST NOT use the upstream channel until it has successfully synchronized to the downstream.

First, the cable modem MUST establish PMD sublayer synchronization. This implies that it has locked onto the correct frequency, equalized the downstream channel, recovered any PMD sublayer framing and the FEC is operational (refer to clause 11.2.2). At this point, a valid bit stream is being sent to the transmission convergence sublayer. The transmission convergence sublayer performs its own synchronization (see clause 7). On detecting the well-known DOCS PID, along with a payload unit start indicator per [ITU-T H.222.0], it delivers the MAC frame to the MAC sublayer.

The MAC sublayer MUST now search for the timing synchronization (SYNC) MAC management messages. For TDMA channels, the cable modem achieves MAC synchronization once it has received at least two SYNC messages and has verified that its clock tolerances are within specified limits. For S-CDMA channels, the cable modem achieves MAC synchronization once it has received at least two SYNC messages, received one UCD message, and has locked to the downstream symbol clock and verified that its clock tolerances are within specified limits.

A cable modem remains in "SYNC" as long as it continues to successfully receive the SYNC messages. If the lost SYNC interval (refer to Annex B) has elapsed without a valid SYNC message, a cable modem MUST NOT use the upstream and MUST try to re-establish synchronization again.

9.3.3 Ranging

Ranging is the process of acquiring the correct timing offset such that the cable modem's transmissions are aligned to the correct mini-slot boundary. The timing delays through the PHY layer MUST be relatively constant with the exception of the timing offsets specified in clause 6.2.19.4 related to modulation rate changes to accommodate a legacy DOCSIS upstream receiver implementation. For TDMA, any variation in the PHY delays MUST be accounted for in the guard time of the upstream PMD overhead.

9.3.3.1 Broadcast initial ranging

First, a cable modem MUST synchronize to the downstream and learn the upstream channel characteristics through the upstream channel descriptor MAC management message. At this point, the cable modem MUST scan the bandwidth allocation MAP message to find a broadcast initial maintenance region. Refer to clause 9.1.2.4. The CMTS MUST make a broadcast initial maintenance region large enough to account for the variation in delays between any two CMs. On S-CDMA channels, the CMTS MUST schedule broadcast initial maintenance transmit opportunities such that they align with S-CDMA frames and span an integral number of S-CDMA frames. Refer to clause 6.2.11.5.

The cable modem MUST put together either an initial ranging request message or a ranging request message to be sent in a broadcast initial maintenance region. An INIT-RNG-REQ MUST be transmitted if the upstream is a DOCS 2.0-only upstream, which can be determined from the UCD. Otherwise a RNG-REQ MUST be transmitted. The SID field MUST be set to the non-initialized CM value (zero), unless this initial ranging is a result of a UCD ranging required TLV or a DCC or UCC request in which the CM has been instructed to retain its existing SIDs.

The CM MUST set its initial timing offset to the amount of internal fixed delay equivalent to putting this CM next to the CMTS. This amount includes delays introduced through a particular implementation, and MUST include the downstream PHY interleaving latency. When the broadcast initial maintenance transmit opportunity occurs, the cable modem MUST send the INIT-RNG-REQ or RNG-REQ message. Thus, the cable modem sends the message as if it were physically right at the CMTS.

Once the CMTS has successfully received the RNG-REQ or INIT-RNG-REQ message, it MUST return a ranging response message addressed to the individual cable modem. Within the ranging response message MUST be a temporary SID assigned to this cable modem (unless the CM has retained a previous primary SID during a UCC, DCC or UCD change) until it has completed the registration process. The message MUST also contain information on RF power level adjustment and offset frequency adjustment as well as any timing offset corrections. Ranging adjusts each CM's timing offset such that it appears to be located right next to the CMTS.

9.3.3.2 Unicast initial ranging

The cable modem MUST now wait for an individual station maintenance or unicast initial maintenance region assigned to its temporary SID (or previous primary SID if ranging as a result of a UCC, DCC or UCD change). It MUST now transmit a ranging request message at this time using the temporary SID (or previous primary SID) along with any power level and timing offset corrections.

The CMTS MUST return another ranging response message to the cable modem with any additional fine tuning required. The ranging request/response steps MUST be repeated until the response contains a ranging successful notification or the CMTS aborts ranging. Once successfully ranged, the cable modem MUST join normal data traffic in the upstream. See clause 11 for complete details on the entire initialization sequence. In particular, state machines and the applicability of retry counts and timer values for the ranging process are defined in clause 11.2.4.

NOTE – The burst type to use for any transmission is defined by the interval usage code (IUC). Each IUC is mapped to a burst type in the UCD message.

9.3.4 Timing units and relationships

The SYNC message conveys a time reference with a resolution of 6.25/64 microseconds (10.24 MHz) to allow the CM to track the CMTS clock with a small phase offset. Since this timing reference is decoupled from particular upstream channel characteristics, a single SYNC time reference may be used for all upstream channels associated with the downstream channel.

The bandwidth allocation MAP uses time units of "mini-slots." A mini-slot represents the time needed for transmission of a fixed number of symbols. For some modulations (e.g., QPSK), an integer number of bytes can be transmitted in a mini-slot. For these channels, the mini-slot is expected to represent 16 byte-times, although other values could be chosen.

A "mini-slot" is the unit of granularity for upstream transmission opportunities; there is no implication that any PDU can actually be transmitted in a single mini-slot.

9.3.4.1 TDMA timing units and relationships

9.3.4.1.1 Mini-slot capacity

On TDMA channels, the size of the mini-slot, expressed as a multiple of the SYNC time reference, is carried in the upstream channel descriptor. The example in Table 9-2 relates mini-slots to the SYNC time ticks (assuming QPSK modulation):

Table 9-2 – Example relating mini-slots to time ticks

Parameter	Example value
Time tick	6.25 microseconds
Bytes per mini-slot	16 (nominal, when using QPSK modulation)
Symbols/byte	4 (assuming QPSK)
Symbols/second	2 560 000
Mini-slots/second	40 000
Microseconds/mini-slot	25
Ticks/mini-slot	4

Note that the symbols/byte is a characteristic of an individual burst transmission, not of the channel. A mini-slot in this instance could represent a minimum of 16 or a maximum of 48 bytes, depending on the modulation choice.

In a channel allocated exclusively to DOCS 2.0 TDMA modems, the mini-slot size field of the UCD MAY take on the value 0, in which case the mini-slot size is 1 timebase tick. If a channel is to be accessible to both DOCS 1.x and 2.0 TDMA cable modems, the UCD MUST follow the DOCS 1.x requirements for timing units and relationships.

9.3.4.1.2 Mini-slot numbering

The MAP counts mini-slots in a 32-bit counter that normally counts to ($2^{32} - 1$) and then wraps back to zero. The least-significant bits (i.e., bit 0 to bit 25-M) of the mini-slot counter MUST match the most-significant bits (i.e., bit 6 + M to bit 31) of the SYNC timestamp counter. That is, mini-slot N begins at timestamp reference ($N \times T \times 64$), where $T = 2^M$ is the UCD multiplier that defines the mini-slot (i.e., the number of time ticks per mini-slot).

NOTE 1 – The unused upper bits of the 32-bit mini-slot counter (i.e., bit 26-M to bit 31) are not needed by the CM and MAY be ignored.

NOTE 2 – The constraint that the UCD multiplier be a power of two has the consequence that the number of bytes per mini-slot must also be a power of two.

9.3.4.2 S-CDMA timing units and relationships

9.3.4.2.1 Mini-slot capacity

On S-CDMA channels, the size of the mini-slot is dependent on the modulation rate, the codes per mini-slot, and the spreading intervals per frame, which are all carried in the upstream channel descriptor. The timing units and relationships for S-CDMA are covered in detail in clause 6.2.11, "S-CDMA framer and interleaver". An example of the timing relationships (assuming 64QAM modulation) is shown in Table 9-3:

Table 9-3 – Example of mini-slot capacity in S-CDMA mode

Parameter	Example value
Spreading intervals per frame	10
Active code length	128
Codes per mini-slot	4
Mini-slots per frame	32
Symbols per mini-slot	40
Bytes per mini-slot	30 (nominal, when using 64QAM modulation)
Bits/symbol	6 (assuming 64QAM)
Symbols/second	5 120 000
Mini-slots/second	128 000
Microseconds/mini-slot	7.8125

9.3.4.2.2 Mini-slot numbering

Mini-slot numbering in S-CDMA mode is described in detail in clause 6.2.11.2.

9.4 Upstream transmission and contention resolution

The CMTS controls assignments on the upstream channel through the MAP and determines which mini-slots are subject to collisions. The CMTS MAY allow collisions on either requests or data PDUs.

This clause provides an overview of upstream transmission and contention resolution. For simplicity, it refers to the decisions a CM makes; however, this is just a pedagogical tool. Since a CM can have multiple upstream service flows (each with its own SID) it makes these decisions on a per service queue or per SID basis. Refer to Appendix IV for a state transition diagram and more detail.

9.4.1 Contention resolution overview

The mandatory method of contention resolution which MUST be supported is based on a truncated binary exponential back-off, with the initial back-off window and the maximum back-off window controlled by the CMTS. The values are specified as part of the bandwidth allocation map (MAP) MAC message and represent a power-of-two value. For example, a value of 4 indicates a window between 0 and 15; a value of 10 indicates a window between 0 and 1023.

Every time a CM wants to transmit in a contention region, it MUST enter the contention resolution process by setting its internal back-off window equal to the data backoff start defined in the MAP currently in effect[18].

The CM MUST randomly select a number within its back-off window. This random value indicates the number of contention transmit opportunities that the CM MUST defer before transmitting. A CM MUST only consider contention transmit opportunities for which this transmission would have been eligible. These are defined by either request IEs or request/data IEs in the MAP.

NOTE 1 – Each IE can represent multiple transmission opportunities.

As an example, consider a CM whose initial back-off window is 0 to 15 and it randomly selects the number 11. The CM must defer a total of 11 contention transmission opportunities. If the first available request IE is for 6 requests, the CM does not use this and has 5 more opportunities to

[18] The MAP currently in effect is the MAP whose allocation start time has occurred but which includes IEs that have not occurred.

defer. If the next request IE is for 2 requests, the CM has 3 more to defer. If the third request IE is for 8 requests, the CM transmits on the fourth request, after deferring for 3 more opportunities.

After a contention transmission, the CM waits for a data grant (data grant pending) or data acknowledge in a subsequent MAP. Once either is received, the contention resolution is complete. The CM determines that the contention transmission was lost when it finds a MAP without a data grant (data grant pending) or data acknowledge for it and with an ack time more recent than the time of transmission[19]. The CM MUST now increase its back-off window by a factor of two, as long as it is less than the maximum back-off window. The CM MUST randomly select a number within its new back-off window and repeat the deferring process described above.

This re-try process continues until the maximum number of retries (16) has been reached, at which time the PDU MUST be discarded.

NOTE 2 – The maximum number of retries is independent of the initial and maximum back-off windows that are defined by the CMTS.

If the CM receives a unicast request or data grant at any time while deferring for this SID, it MUST stop the contention resolution process and use the explicit transmit opportunity.

The CMTS has much flexibility in controlling the contention resolution. At one extreme, the CMTS may choose to set up the data backoff start and end to emulate an Ethernet-style back-off with its associated simplicity and distributed nature, but also its fairness and efficiency issues. This would be done by setting data backoff start = 0 and end = 10 in the MAP. At the other end, the CMTS may make the data backoff start and end identical and frequently update these values in the MAP so all cable modems are using the same, and hopefully optimal, back-off window.

A CM transmitting a RNG-REQ in the initial maintenance IE MUST perform truncated binary exponential backoff using the ranging backoff start and ranging backoff end to control the backoff window. The algorithm works similarly to data transmissions, except on the calculation of transmit opportunities, which is described in the next clause.

9.4.2 Transmit opportunities

A transmit opportunity is defined as any mini-slot in which a CM may be allowed to start a transmission. Transmit opportunities typically apply to contention opportunities and are used to calculate the proper amount to defer in the contention resolution process.

The number of transmit opportunities associated with a particular IE in a MAP is dependent on the total size of the region as well as the allowable size of an individual transmission. As an example, assume a contention REQ IE defines a region of 12 mini-slots. If the UCD defines a REQ burst size that fits into a single mini-slot then there are 12 transmit opportunities associated with this REQ IE, i.e., one for each mini-slot. If the UCD defines a REQ that fits in two mini-slots, then there are six transmit opportunities and a REQ can start on every other mini-slot.

As another example, assume a REQ/data IE that defines a 24 mini-slot region. If it is sent with an SID of 0x3FF4 (refer to Annex A), then a CM can potentially start a transmit on every fourth mini-slot; so this IE contains a total of six transmit opportunities (TX OPs). Similarly, a SID of 0x3FF6 implies four TX OPs; 0x3FF8 implies three TX OPs; and 0x3FFC implies two TX OPs.

For a broadcast initial maintenance IE, a CM MUST start its transmission in the first mini-slot of the region; therefore it has a single transmit opportunity. The remainder of the region is used to compensate for the round-trip delays since the CM has not yet been ranged.

[19] Data acknowledge IEs are intended for collision detection only and are not designed for providing reliable transport (that is the responsibility of higher layers). If a MAP is lost or damaged, a CM waiting for a data acknowledge MUST assume that its contention data transmission was successful and MUST NOT retransmit the data packet. This prevents the CM from sending duplicate packets unnecessarily.

Station maintenance IEs, short/long data grant IEs, adv PHY short/long data grant IEs, adv PHY unsolicited grant IEs, unicast initial maintenance, and unicast request IEs are unicast and thus are not typically associated with contention transmit opportunities. They represent a single dedicated, or reservation-based, transmit opportunity.

In summary:

Table 9-4 – Transmit opportunity

Interval	SID type	Transmit opportunity
Request	Broadcast	# mini-slots required for a request
Request	Multicast	# mini-slots required for a request
Request/data	Broadcast	Not allowed
Request/data	Well-known multicast	As defined by SID in Annex A
Request/data	Multicast	Vendor-specific algorithms
Initial maintenance	Broadcast	Entire interval is a single TX OP
NOTE – Transmit opportunity should not be confused with burst size. Burst size requirements are specified in Table 6-1.		

9.4.3 CM bandwidth utilization

The following rules govern the response a CM makes when processing maps.

NOTE – These standard behaviours can be overridden by the CM's request/transmission policy (refer to clause C.2.2.6.3):

1) A CM MUST first use any grants assigned to it. Next, the CM MUST use any unicast REQ for it. Finally, the CM MUST use the next available broadcast/multicast REQ or REQ/data IEs for which it is eligible.
2) A CM MUST NOT have more than one request outstanding at a time for a particular service ID.
3) If a CM has a request pending, it MUST NOT use intervening contention intervals for that service ID.

9.5 Data link encryption support

The procedures to support data link encryption are defined in [DOCS8]. The interaction between the MAC layer and the security system is limited to the items defined below.

9.5.1 MAC messages

MAC management messages (see clause 8.3) MUST NOT be encrypted, except for certain cases where such a frame is included in a fragmented concatenated burst on the upstream (refer to clause 8.2.7.1)

9.5.2 Framing

The following rules MUST be followed when encryption is applied to a data PDU:

- Privacy EH element of [DOCS8] MUST be in the extended header and MUST be the first EH element of the extended header field (EHDR).
- Encrypted data are carried as data PDUs to the cable MAC transparently.

10 Quality of service and fragmentation

This Recommendation introduces several new quality of service (QoS) related concepts not present in the original March 1998 version of J.112 Annex B. These include:

- packet classification and flow identification;
- service flow QoS scheduling;
- dynamic service establishment;
- fragmentation;
- two-phase activation model.

10.1 Theory of operation

The various DOCS protocol mechanisms described in this Recommendation can be used to support quality of service (QoS) for both upstream and downstream traffic through the CM and the CMTS. This clause provides an overview of the QoS protocol mechanisms and their part in providing end-to-end QoS.

The requirements for quality of service include:

- a configuration and registration function for pre-configuring CM-based QoS **service flows** and traffic parameters;
- a signalling function for dynamically establishing QoS-enabled service flows and traffic parameters;
- a traffic-shaping and traffic-policing function for service flow-based traffic management, performed on traffic arriving from the upper layer service interface and outbound to the RF;
- utilization of MAC scheduling and traffic parameters for upstream service flows;
- utilization of QoS traffic parameters for downstream service flows;
- classification of packets arriving from the upper layer service interface to a specific active service flow;
- grouping of service flow properties into named **service classes**, so upper layer entities and external applications (at both the CM and CMTS) can request service flows with desired QoS parameters in a globally consistent way.

The principal mechanism for providing enhanced QoS is to classify packets traversing the RF MAC interface into a **service flow**. A service flow is a unidirectional flow of packets that is provided a particular quality of service. The CM and CMTS provide this QoS by shaping, policing and prioritizing traffic according to the **QoS parameter set** defined for the service flow.

The primary purpose of the quality of service features defined here is to define transmission ordering and scheduling on the radio frequency interface. However, these features often need to work in conjunction with mechanisms beyond the RF interface in order to provide end-to-end QoS or to police the behaviour of cable modems. For example, the following behaviours are permitted:

- Policies may be defined by CM MIBs which overwrite the ToS byte. Such policies are outside the scope of the RFI Recommendation. In the upstream direction, the CMTS polices the ToS byte setting regardless of how the ToS byte is derived or by whom it is written (originator or CM policy).
- The queueing of service flow packets at the CMTS in the downstream direction may be based on the ToS byte.
- Downstream service flows can be reclassified by the CM to provide enhanced service onto the subscriber-side network.

Service flows exist in both the upstream and downstream direction, and may exist without actually being activated to carry traffic. Service flows have a 32-bit **service flow identifier** (SFID) assigned by the CMTS. All service flows have an SFID; active and admitted upstream service flows also have a 14-bit service identifier (SID).

At least two service flows must be defined in each configuration file: one for upstream and one for downstream service. The first upstream service flow describes the **primary upstream service flow,** and is the default service flow used for otherwise unclassified traffic, including both MAC management messages and data PDUs. The first downstream service flow describes service to the primary downstream service flow. Additional service flows defined in the configuration file create service flows that are provided QoS services.

Conceptually, incoming packets are matched to a **classifier** that determines to which QoS service flow the packet is forwarded. The classifier can examine the LLC header of the packet, the IP/TCP/UDP header of the packet or some combination of the two. If the packet matches one of the classifiers, it is forwarded to the service flow indicated by the SFID attribute of the classifier. If the packet is not matched to a classifier, it is forwarded on the primary service flow.

10.1.1 Concepts

10.1.1.1 Service flows

A service flow is a MAC-layer transport service that provides unidirectional transport of packets either to upstream packets transmitted by the CM or to downstream packets transmitted by the CMTS[20]. A service flow is characterized by a set of **QoS parameters** such as latency, jitter and throughput assurances. In order to standardize operation between the CM and CMTS, these attributes include details of how the CM requests upstream mini-slots and the expected behaviour of the CMTS upstream scheduler.

A service flow is partially characterized by the following attributes[21]:

- **ServiceFlowID**: Exists for all service flows.
- **ServiceID**: Only exists for admitted or active upstream service flows.
- **ProvisionedQosParamSet**: Defines a set of QoS parameters that appears in the configuration file and is presented during registration. This MAY define the initial limit for authorizations allowed by the authorization module. The ProvisionedQosParamSet is defined once when the service flow is created via registration[22].
- **AdmittedQosParamSet**: Defines a set of QoS parameters for which the CMTS (and possibly the CM) are reserving resources. The principal resource to be reserved is bandwidth, but this also includes any other memory or time-based resource required to subsequently activate the flow.
- **ActiveQosParamSet**: Defines a set of QoS parameters defining the service actually being provided to the service flow. Only an active service flow may forward packets.

---

[20] A service flow, as defined here, has no direct relationship to the concept of a "flow" as defined by the IETF's integrated services (intserv) working group [b-RFC 2212]. An intserv flow is a collection of packets sharing transport-layer endpoints. Multiple intserv flows can be served by a single service flow. However, the classifiers for a service flow may be based on 802.1P/Q criteria, and so may not involve intserv flows at all.

[21] Some attributes are derived from the above attribute list. The service class name is an attribute of the ProvisionedQoSParamSet. The activation state of the service flow is determined by the ActiveQoSParamSet. If the ActiveQoSParamSet is null then the service flow is inactive.

[22] The ProvisionedQoSParamSet is null when a flow is created dynamically.

A service flow exists when the CMTS assigns a service flow ID (SFID) to it. The SFID serves as the principal identifier in the CM and CMTS for the service flow. A service flow which exists has at least an SFID, and an associated direction.

The **authorization module** is a logical function within the CMTS that approves or denies every change to QoS parameters and classifiers associated with a service flow. As such, it defines an "envelope" that limits the possible values of the AdmittedQosParamSet and ActiveQosParamSet.

The relationship between the QoS parameter sets is as shown in Figures 10-1 and 10-2. The ActiveQosParamSet is always a subset[23] of the AdmittedQosParamSet which is always a subset of the authorized "envelope." In the dynamic authorization model, this envelope is determined by the authorization module (labelled as the AuthQosParamSet). In the provisioned authorization model, this envelope is determined by the ProvisionedQosParamSet (refer to clause 10.1.4 for further information on the authorization models).

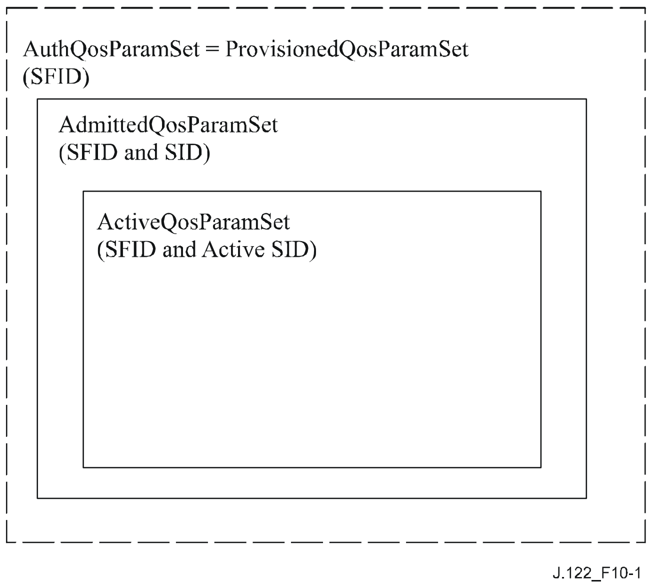

Figure 10-1 – Provisioned authorization model "envelopes"

[23] To say that QoS parameter set A is a subset of QoS parameter set B, the following MUST be true for all QoS parameters in A and B:

– If a smaller QoS parameter value indicates less resources (e.g., maximum traffic rate), A is a subset of B if the parameter in A less than or equal to the same parameter in B.

– If a larger QoS parameter value indicates less resources (e.g., tolerated grant jitter), A is a subset of B if the parameter in A is greater than or equal to the same parameter in B.

– If the QoS parameter specifies a periodic interval (e.g., nominal grant interval), A is a subset of B if the parameter in A is an integer multiple of the same parameter in B.

– If the QoS parameter is not quantitative (e.g., service flow scheduling type), A is a subset of B if the parameter in A is equal to the same parameter in B.

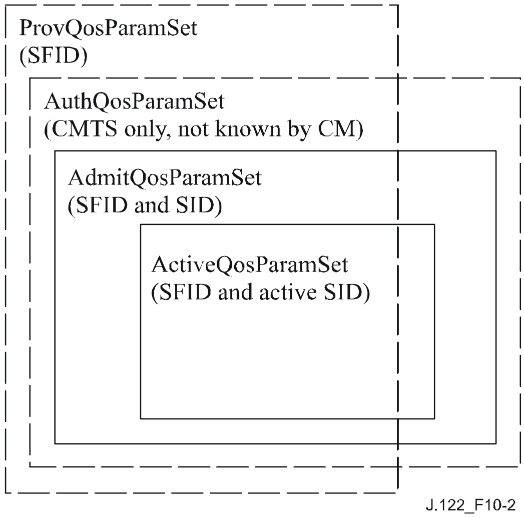

Figure 10-2 – Dynamic authorization model "envelopes"

It is useful to think of three types of service flows:

- **Provisioned**: this type of service flow is known via provisioning through the configuration file, its AdmittedQosParamSet and ActiveQosParamSet are both null. A **provisioned service flow** may or may not have associated classifiers. If a provisioned service flow has associated classifiers, the classifiers MUST NOT be used to classify packets onto the flow, regardless of the classifier's activation state.
- **Admitted**: this type of service flow has resources reserved by the CMTS for its AdmittedQosParamSet, but these parameters are not active (its ActiveQosParamSet is null). **Admitted service flows** may have been provisioned or may have been signalled by some other mechanism. Generally, admitted service flows have associated classifiers, however, it is possible for admitted service flows to use policy-based classification. If admitted service flows have associated classifiers, the classifiers MUST NOT be used to classify packets onto the flow, regardless of the classifier's activation state.
- **Active**: this type of service flow has resources committed by the CMTS for its QoS parameter set, (e.g., is actively sending MAPs containing unsolicited grants for a UGS-based service flow). Its ActiveQosParamSet is non-null. Generally, active service flows have associated classifiers; however, it is possible for active service flows to use policy-based classification. primary service flows may have associated classifiers(s), but in addition to any packets matching such classifiers, all packets that fail to match any classifier will be sent on the primary service flow for that direction.

10.1.1.2 Classifiers

A **classifier** is a set of matching criteria applied to each packet entering the cable network. It consists of some packet-matching criteria (destination IP address, for example), a **classifier priority,** and a reference to a service flow. If a packet matches the specified packet matching criteria, it is then delivered on the referenced service flow.

Several classifiers may all refer to the same service flow. The classifier priority is used for ordering the application of classifiers to packets. Explicit ordering is necessary because the patterns used by classifiers may overlap. The priority need not be unique, but care must be taken within a classifier priority to prevent ambiguity in classification (refer to clause 10.1.6.1). **Downstream classifiers** are applied by the CMTS to packets it is transmitting, and **upstream classifiers** are applied at the CM and may be applied at the CMTS to police the classification of upstream packets. Figure 10-3 illustrates the mappings discussed above.

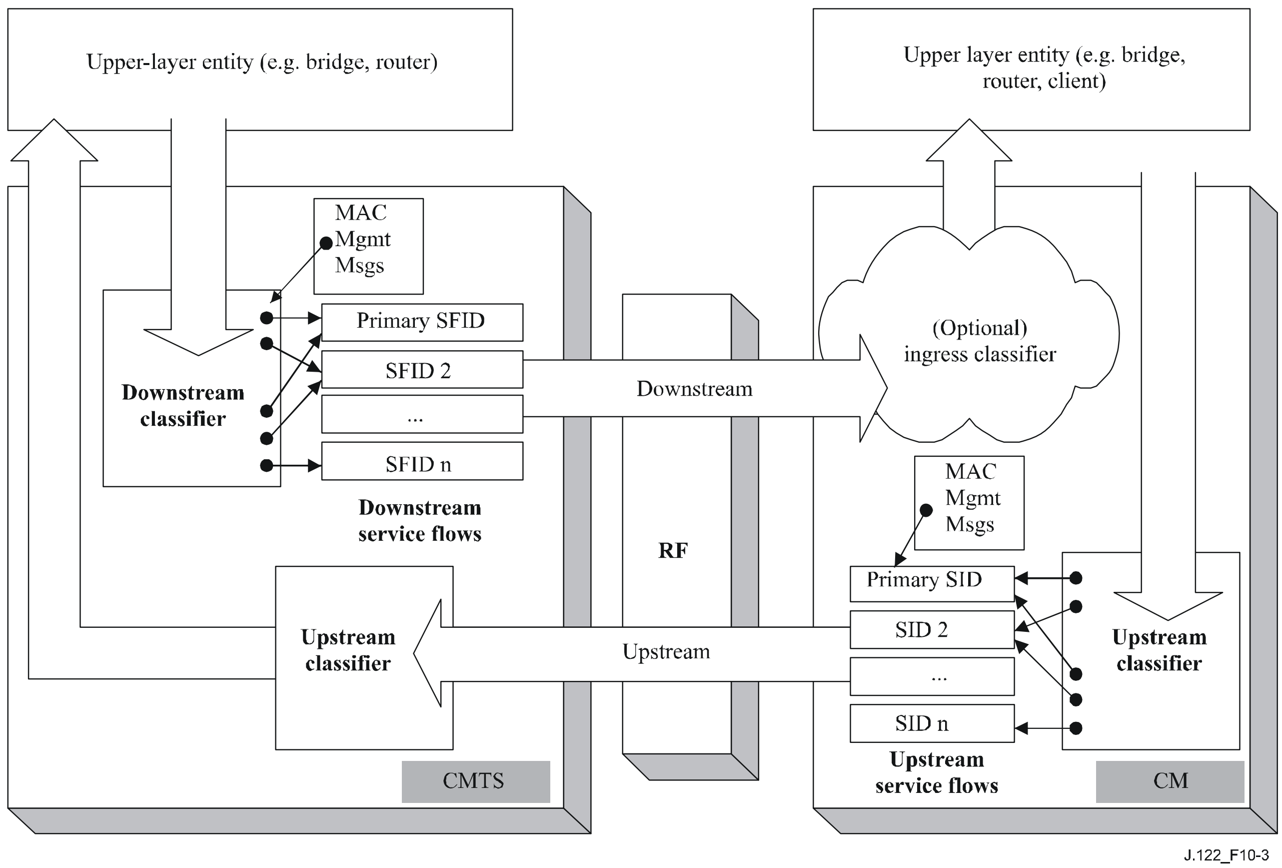

Figure 10-3 – Classification within the MAC layer

CM and CMTS packet classification consists of multiple classifiers. Each classifier contains a priority field which determines the search order for the classifier. The highest priority classifier MUST be applied first. If a classifier is found in which all parameters match the packet, the classifier MUST forward the packet to the corresponding service flow (irrelevant parameters – as defined in clause C.2.1 – have no impact on packet classification decisions). If a classifier contains no relevant parameters for a given packet (i.e., all parameters are irrelevant), then that packet cannot match the classifier, and the classifier MUST NOT forward the packet to the corresponding service flow. If a packet does not match any classifier and, as a result, has not been classified to any other flow, then it MUST be classified to the primary service flow.

The packet classification table contains the following fields:

- Priority – Determines the search order for the table. Higher priority classifiers are searched before lower priority classifiers.
- IP classification parameters – Zero or more of the IP classification parameters (IP ToS range/mask, IP protocol, IP source address/mask, IP destination address/mask, TCP/UDP source port start, TCP/UDP source port end, TCP/UDP destination port start, TCP/UCP destination port end).
- LLC classification parameters – Zero or more of the LLC classification parameters (destination MAC address, source MAC address, ethertype/SAP).
- IEEE 802.1P/Q parameters – Zero or more of the IEEE classification parameters (802.1P priority range, 802.1Q VLAN ID).
- Service flow identifier – Identifier of a specific flow to which this packet is to be directed.

Classifiers can be added to the table either via management operations (configuration file, registration) or via dynamic operations (dynamic signalling, DOCS MAC sublayer service interface). SNMP-based operations can view classifiers that are added via dynamic operations, but cannot modify or delete classifiers that are created by dynamic operations. The format for classification table parameters defined in the configuration file, registration message or dynamic signalling message is contained in Annex C.

Classifier attributes include an activation state (see clause C.2.1.3.6). The 'inactive' setting may be used to reserve resources for a classifier which is to be activated later. The actual activation of the classifier depends both on this attribute and on the state of its service flow. If the service flow is not active then the classifier is not used, regardless of the setting of this attribute.

10.1.2 Object model

The major objects of the architecture are represented by named rectangles in Figure 10-4. Each object has a number of attributes; the attribute names that uniquely identify the object are underlined. Optional attributes are denoted with brackets. The relationship between the number of objects is marked at each end of the association line between the objects. For example, a service flow may be associated with from 0 to 65535 classifiers, but a classifier is associated with exactly one service flow.

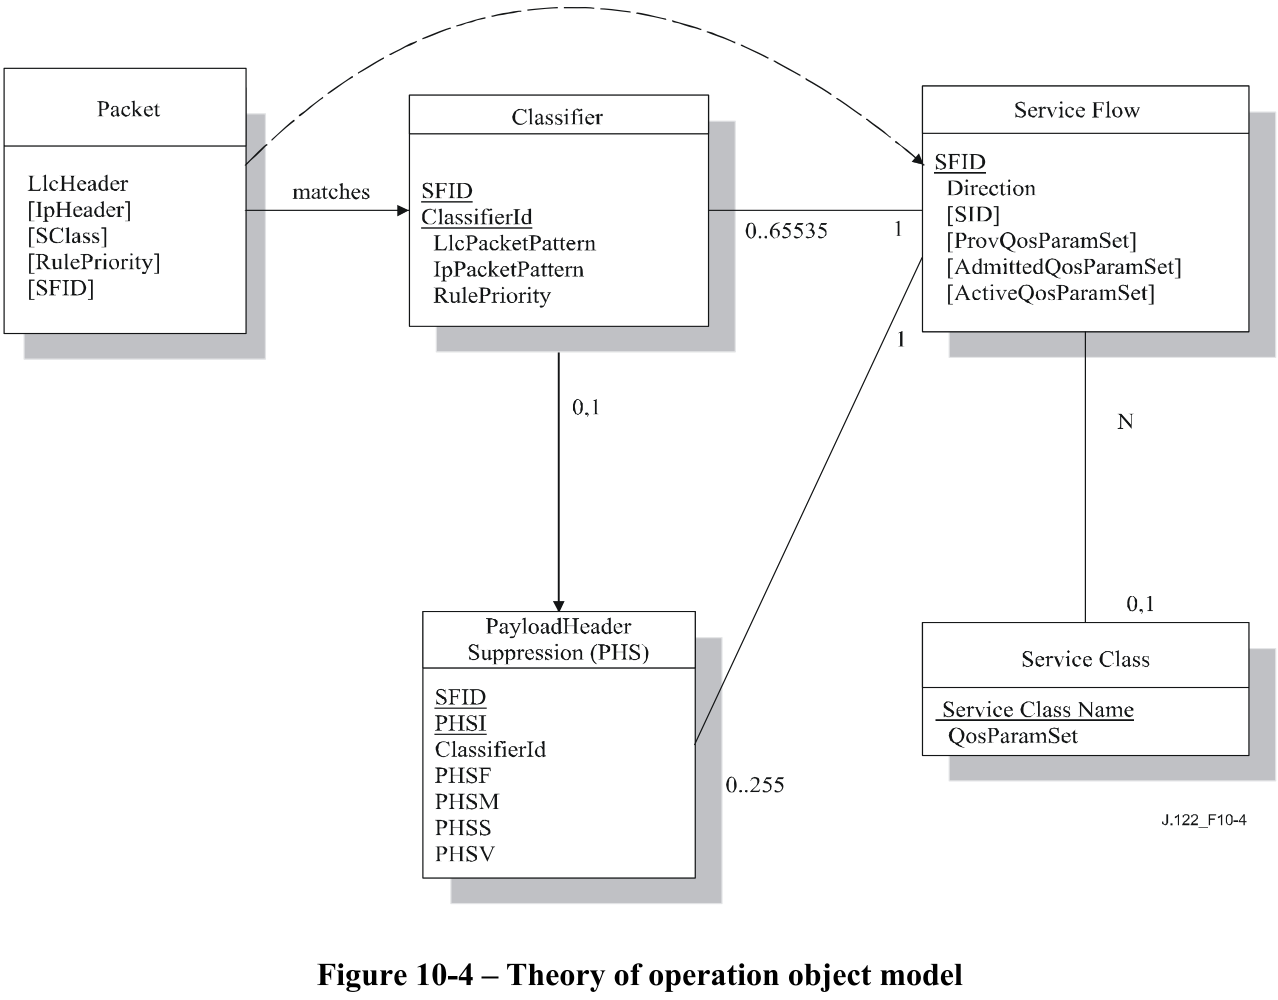

Figure 10-4 – Theory of operation object model

The service flow is the central concept of the MAC protocol. It is uniquely identified by a 32-bit service flow ID (SFID) assigned by the CMTS. Service flows may be in either the upstream or downstream direction. A unicast service identifier (SID) is a 14-bit index, assigned by the CMTS, which is associated with one, and only one, admitted upstream service flow.

Typically, an outgoing user data packet is submitted by an upper layer protocol (such as the forwarding bridge of a CM) for transmission on the cable MAC interface. The packet is compared against a set of classifiers. The matching classifier for the packet identifies the corresponding service flow via the service flow ID (SFID). In the case where more than one classifier matches the packet, the highest priority classifier is chosen.

The classifier matching a packet may be associated with a payload header suppression rule. A PHS rule provides details on how header bytes of a packet PDU can be omitted, replaced with a payload header suppression index for transmission, and subsequently regenerated at the receiving end. PHS rules are indexed by the combination of {SFID, PHSI} (refer to clause 10.4). When a service flow is deleted, all classifiers and any associated PHS rules referencing it MUST also be deleted.

The service class is an optional object that MUST be implemented at the CMTS. It is referenced by an ASCII name which is intended for provisioning purposes. A service class is defined in the CMTS to have a particular QoS parameter Set. A service flow may contain a reference to the service class name that selects all of the QoS parameters of the service class. The service flow QoS parameter sets may augment and even override the QoS parameter settings of the service class, subject to authorization by the CMTS (refer to clause C.2.2.5).

If a packet has already been determined by upper layer policy mechanisms to be associated with a particular service class name/priority combination, that combination associates the packet with a particular service flow directly (refer to clause 10.1.6.1). The upper layer may also be aware of the particular service flows in the MAC sublayer, and may have assigned the packet directly to a service flow. In these cases, a user data packet is considered to be directly associated with a service flow as selected by the upper layer. This is depicted with the dashed arrow in Figure 10-4 (refer to Appendix I).

10.1.3 Service classes

The QoS attributes of a service flow may be specified in two ways: either by explicitly defining all attributes, or implicitly by specifying a **service class name.** A service class name is a string which the CMTS associates with a QoS parameter set. It is described further below.

The service class serves the following purposes:

1) It allows operators, who so wish, to move the burden of configuring service flows from the provisioning server to the CMTS. Operators provision the modems with the service class name; the implementation of the name is configured at the CMTS. This allows operators to modify the implementation of a given service to local circumstances without changing modem provisioning. For example, some scheduling parameters may need to be tweaked differently for two different CMTSs to provide the same service. As another example, service profiles could be changed by time of day.
2) It allows CMTS vendors to provide class-based-queuing if they choose, where service flows compete within their class and classes compete with each other for bandwidth.
3) It allows higher-layer protocols to create a service flow by its service class name. For example, telephony signalling may direct the CM to instantiate any available provisioned service flow of class "G711".
4) It allows packet classification policies to be defined which refer to a desired service class, without having to refer to a particular service flow instance of that class.

NOTE – The service class is optional: the flow scheduling specification may always be provided in full; a service flow may belong to no service class whatsoever. CMTS implementations MAY treat such "unclassed" flows differently from "classed" flows with equivalent parameters.

Any service flow MAY have each of its QoS parameter sets specified in any of three ways:

- by explicitly including all traffic parameters;

- by indirectly referring to a set of traffic parameters by specifying a service class name;
- by specifying a service class name along with modifying parameters.

The service class name is "expanded" to its defined set of parameters at the time the CMTS successfully admits the service flow. The service class expansion can be contained in the following CMTS-originated messages: Registration response, DSA-REQ, DSC-REQ, DSA-RSP and DSC-RSP. In all of these cases, the CMTS MUST include a service flow encoding that includes the service class name and the QoS parameter set of the service class. If a CM-initiated request contained any supplemental or overriding service flow parameters, a successful response MUST also include these parameters.

When a service class name is given in an admission or activation request, the returned QoS parameter set may change from activation to activation. This can happen because of administrative changes to the service class' QoS parameter set at the CMTS. If the definition of a service class name is changed at the CMTS (e.g., its associated QoS parameter set is modified), it has no effect on the QoS parameters of existing service flows associated with that service class. A CMTS MAY initiate DSC transactions to existing service flows that reference the service class name to affect the changed service class definition.

When a CM uses the service class name to specify the admitted QoS parameter set, the expanded set of TLV encodings of the service flow will be returned to the CM in the response message (REG-RSP, DSA-RSP or DSC-RSP). Use of the service class name later in the activation request may fail if the definition of the service class name has changed and the new required resources are not available. Thus, the CM SHOULD explicitly request the expanded set of TLVs from the response message in its later activation request.

10.1.4 Authorization

Every change to the service flow QoS parameters MUST be approved by an authorization module. This includes every REG-REQ or DSA-REQ message to create a new service flow, and every DSC-REQ message to change a QoS parameter set of an existing service flow. Such changes include requesting an admission control decision (e.g., setting the AdmittedQosParamSet) and requesting activation of a service flow (e.g., setting the ActiveQosParamSet). Reduction requests regarding the resources to be admitted or activated are also checked by the authorization module, as are requests to add or change the classifiers.

In the static authorization model, the authorization module receives all registration messages and stores the provisioned status of all "deferred" service flows. Admission and activation requests for these provisioned service flows will be permitted as long as the admitted QoS parameter set is a subset of the provisioned QoS parameter set, and the active QoS parameter set is a subset of the admitted QoS parameter set. Requests to change the provisioned QoS parameter set will be refused, as will requests to create new dynamic service flows. This defines a static system where all possible services are defined in the initial configuration of each CM.

In the dynamic authorization model, the authorization module not only receives all registration messages, but also communicates through a separate interface to an independent policy server. This policy server may provide to the authorization module advance notice of upcoming admission and activation requests, and specifies the proper authorization action to be taken on those requests. Admission and activation requests from a CM are then checked by the authorization module to ensure that the ActiveQosParamSet being requested is a subset of the set provided by the policy server. Admission and activation requests from a CM that are signalled in advance by the external policy server are permitted. Admission and activation requests from a CM that are not pre-signalled by the external policy server may result in a real-time query to the policy server, or may be refused.

During registration, the CM MUST send to the CMTS the authenticated set of TLVs derived from its configuration file which defines the provisioned QoS parameter set. Upon receipt and verification at the CMTS, these are handed to the authorization module within the CMTS. The CMTS MUST be capable of caching the provisioned QoS parameter set, and MUST be able to use this information to authorize dynamic flows which are a subset of the provisioned QoS parameter set. The CMTS SHOULD implement mechanisms for overriding this automated approval process (such as described in the dynamic authorization model). For example:

- Deny all requests whether or not they have been pre-provisioned.
- Define an internal table with a richer policy mechanism but seeded by the configuration file information.
- Refer all requests to an external policy server.

10.1.5 Types of service flows

It is useful to think about three basic types of service flows. This clause describes these three types of service flows in more detail. However, it is important to note that there are more than just these three basic types (refer to clause C.2.2.3.5).

10.1.5.1 Provisioned service flows

A service flow may be provisioned but not immediately activated (sometimes called "deferred"). That is, the description of any such service flow in the TFTP configuration file contains an attribute which provisions but defers activation and admission (refer to clause C.2.2.3.5). During registration, the CMTS assigns a service flow ID for such a service flow but does not reserve resources. The CMTS MAY also require an exchange with a policy module prior to admission.

As a result of external action beyond the scope of this Recommendation (e.g., [ITU-T J.162]), the CM MAY choose to activate a provisioned service flow by passing the service flow ID and the associated QoS parameter sets. The CM MUST also provide any applicable classifiers. If authorized and resources are available, the CMTS MUST respond by assigning a unique unicast SID for the upstream service flow. The CMTS MAY deactivate the service flow, but SHOULD NOT delete the service flow during the CM registration epoch.

As a result of external action beyond the scope of this Recommendation (e.g., [ITU-T J.162]), the CMTS MAY choose to activate a service flow by passing the service flow ID as well as the SID and the associated QoS parameter sets. The CMTS MUST also provide any applicable classifiers. The CMTS MAY deactivate the service flow, but SHOULD NOT delete the service flow during the CM registration epoch. Such a provisioned service flow MAY be activated and deactivated many times (through DSC exchanges). In all cases, the original service flow ID MUST be used when reactivating the service flow.

10.1.5.2 Admitted service flows

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

a) conserving network resources until a complete end-to-end connection has been established;

b) performing policy checks and admission control on resources as quickly as possible, and, in particular, before informing the far end of a connection request; and

c) preventing several potential theft-of-service scenarios.

For example, if an upper-layer service were using unsolicited grant service, and the addition of upper-layer flows could be adequately provided by increasing the grants per interval QoS parameter, then the following might be used. When the first upper-layer flow is pending, the CM issues a DSA-Request with the admit grants per interval parameter equal to one, and the activate grants per interval parameter equal to zero. Later, when the upper-layer flow becomes active, it issues a DSC-Request with the instance of the activate grants per interval parameter equal to one. Admission control was performed at the time of the reservation, so the later DSC-Request, having the activate parameters within the range of the previous reservation, is guaranteed to succeed. Subsequent upper-layer flows would be handled in the same way. If there were three upper-layer flows establishing connections, with one flow already active, the service flow would have admit(ted) grants per interval equal to four, and active grants per interval equal to one.

An activation request of a service flow where the new ActiveQosParamSet is a subset of the AdmittedQosParamSet and no new classifiers are being added MUST be allowed (except in the case of catastrophic failure). An admission request where the AdmittedQosParamSet is a subset of the previous AdmittedQosParamSet, so long as the ActiveQosParamSet remains a subset of the AdmittedQosParamSet, MUST succeed.

A service flow that has resources assigned to its AdmittedQosParamSet, but whose resources are not yet completely activated, is in a transient state. A timeout value MUST be enforced by the CMTS that requires service flow activation within this period (refer to clause C.2.2.5.7). If service flow activation is not completed within this interval, the assigned resources in excess of the active QoS parameters MUST be released by the CMTS.

It is possible in some applications that a long-term reservation of resources is necessary or desirable. For example, placing a telephone call on hold should allow any resources in use for the call to be temporarily allocated to other purposes, but these resources must be available for resumption of the call later. The AdmittedQosParamSet is maintained as "soft state" in the CMTS; this state must be refreshed periodically for it to be maintained without the above timeout releasing the non-activated resources. This refresh MAY be signalled with a periodic DSC-REQ message with identical QoS parameter sets, or MAY be signalled by some internal mechanism within the CMTS outside of the scope of this Recommendation (e.g., by the CMTS monitoring RSVP refresh messages). Every time a refresh is signalled to the CMTS, the CMTS MUST refresh the "soft state."

10.1.5.3 Active service flows

A service flow that has a non-NULL set of ActiveQosParameters is said to be an active service flow. It is requesting[24] and being granted bandwidth for transport of data packets. An admitted service flow may be made active by providing an ActiveQosParamSet, signalling the resources actually desired at the current time. This completes the second stage of the two-phase activation model (refer to clause 10.1.5.2).

A service flow may be provisioned and immediately activated. This is the case for the primary service flows. It is also typical of service flows for monthly subscription services, etc. These service flows are established at registration time and MUST be authorized by the CMTS based on the CMTS MIC. These service flows MAY also be authorized by the CMTS authorization module.

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

10.1.6 Service flows and classifiers

The basic model is that the classifiers associate packets into exactly one service flow. The service flow encodings provide the QoS parameters for treatment of those packets on the RF interface. These encodings are described in clause C.2.

In the upstream direction, the CM MUST classify upstream packets to active service flows. The CMTS MUST classify downstream traffic to active downstream service flows. There MUST be a default downstream service flow for otherwise unclassified broadcast and multicast traffic.

---

[24] According to its request/transmission policy (refer to clause C.2.2.6.3).

The CMTS polices packets in upstream service flows to ensure the integrity of the QoS parameters and the packet's ToS value. When the rate at which packets are sent is greater than the policed rate at the CMTS, then these packets MAY be dropped by the CMTS (refer to clause C.2.2.5.2). When the value of the ToS byte is incorrect, the CMTS (based on policy) MUST police the stream by overwriting the ToS byte (refer to clause C.2.2.6.10).

It may not be possible for the CM to forward certain upstream packets on certain service flows. In particular, a service flow using unsolicited grant service with fragmentation disabled cannot be used to forward packets larger than the grant size. If a packet is classified to a service flow on which it cannot be transmitted, the CM MUST either transmit the packet on the primary service flow or discard the packet depending on the request/transmission policy of the service flow to which the packet was classified.

MAC management messages may only be matched by a classifier that contains a clause C.2.1.6.3 "ethertype/DSAP/MacType" parameter encoding and when the "type" field of the MAC management message header (see clause 8.3.1) matches that parameter. One exception is that the primary SID MUST be used for periodic ranging, as specified in clause 8.1.2.3, even if a classifier matches the upstream RNG-REQ message of periodic ranging. In the absence of any classifier matching a MAC management message, it SHOULD be transmitted on the primary service flow. Other than those MAC message types precluded from classification in clause C.2.1.6.3, a CM or CMTS MAY forward an otherwise unclassified MAC message on any service flow in an implementation-specific manner.

Although MAC management messages are subject to classification, they are not considered part of any service flow. Transmission of MAC management messages MUST NOT influence any QoS calculations of the service flow to which they are classified. Delivery of MAC management messages is implicitly influenced by the attributes of the associated service flow.

10.1.6.1 Policy-based classification and service classes

As noted in Appendix I, there are a variety of ways in which packets may be enqueued for transmission at the MAC service interface. At one extreme are embedded applications that are tightly bound to a particular payload header suppression rule (refer to clause 10.4) and which forego more general classification by the MAC. At the other extreme are general transit packets of which nothing is known until they are parsed by the MAC classification rules. Another useful category is traffic to which policies are applied by a higher-layer entity and then passed to the MAC for further classification to a particular service flow.

Policy-based classification is, in general, beyond the scope of this Recommendation. One example might be the docsDevFilterIpPolicyTable defined in the Cable Device MIB [b-RFC 2669]. Such policies may tend to be longer-lived than individual service flows and MAC classifiers and so it is appropriate to layer the two mechanisms, with a well-defined interface between policies and MAC service flow classification.

The interface between the two layers is the addition of two parameters at the MAC transmission request interface. The two parameters are a service class name and a rule priority that is applied to matching the service class name. The policy priority is from the same number space as the packet classifier priority of the packet-matching rules used by MAC classifiers. The MAC classification algorithm is now:

```
MAC_DATA.request( PDU, ServiceClassName, RulePriority )
TxServiceFlowID = FIND_FIRST_SERVICE_FLOW_ID (ServiceClassName)
SearchID = SEARCH_CLASSIFIER_TABLE (All Priority Levels)
IF (SearchID not NULL and Classifier.RulePriority >= MAC_DATA.RulePriority)
            TxServiceFlowID = SearchID
IF (TxServiceFlowID = NULL)
            TRANSMIT_PDU (PrimaryServiceFlowID)
```

```
ELSE
            TRANSMIT_PDU (TxServiceFlowID)
```

While policy priority competes with packet classifier priority and its choice might in theory be problematic, it is anticipated that well-known ranges of priorities will be chosen to avoid ambiguity. In particular, dynamically-added classifiers MUST use the priority range 64-191. Classifiers created as part of registration, as well as policy-based classifiers, may use zero through 255, but SHOULD avoid the dynamic range.

NOTE – Classification within the MAC sublayer is intended to simply associate a packet with a service flow. If a packet is intended to be dropped it MUST be dropped by the higher-layer entity and not delivered to the MAC sublayer.

10.1.7 General operation

10.1.7.1 Static operation

Static configuration of classifiers and service flows uses the registration process. A provisioning server provides the CM with configuration information. The CM passes this information to the CMTS in a registration request. The CMTS adds information and replies with a registration response. The CM sends a registration acknowledge to complete registration.

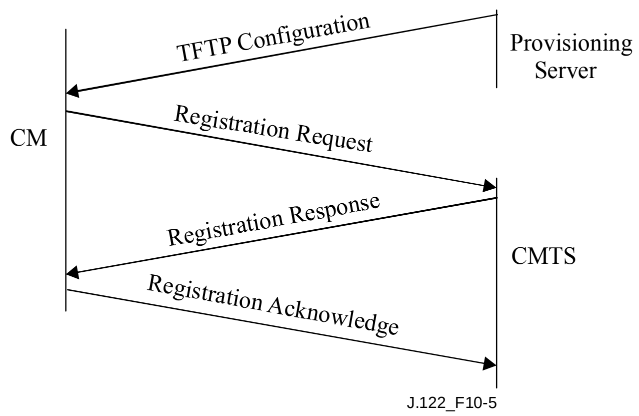

Figure 10-5 – Registration message flow

A TFTP configuration file consists of one or more instances of classifiers and service flow encodings. Classifiers are loosely ordered by 'priority'. Each classifier refers to a service flow via a 'service flow reference'. Several classifiers may refer to the same service flow. Additionally, more than one classifier may have the same priority, and in this case, the particular classifier used is not defined.

Table 10-1 – TFTP file contents

Items	Point to service flow reference	Service flow reference	Service flow ID
Upstream classifiers Each containing a service flow reference (pointer)	1..n		
Downstream classifiers Each containing a service flow reference (pointer)	(n + 1)..q		
Service flow encodings Immediate activation requested, upstream		1..m	None yet
Service flow encodings Provisioned for later activation requested, upstream		(m + 1)..n	None yet
Service flow encodings Immediate activation requested, downstream		(n + 1)..p	None yet
Service flow encodings Provisioned for later activation requested, downstream		(p + 1)..q	None yet

Service flow encodings contain either a full definition of service attributes (omitting defaultable items if desired) or a service class name. A service class name is an ASCII string which is known at the CMTS and which indirectly specifies a set of QoS parameters (refer to clauses 10.1.3 and C.2.2.3.4)

NOTE – At the time of the TFTP configuration file, service flow references exist as defined by the provisioning server. Service flow identifiers do not yet exist because the CMTS is unaware of these service flow definitions.

The registration request packet contains downstream classifiers (if to be immediately activated) and all inactive service flows. The configuration file, and thus the registration request, generally does not contain a downstream classifier if the corresponding service flow is requested with deferred activation. This allows for late binding of the classifier when the flow is activated.

Table 10-2 – Registration request contents

Items	Point to service flow reference	Service flow reference	Service flow ID
Upstream classifiers Each containing a service flow reference (pointer)	1..n		
Downstream classifiers Each containing a service flow reference (pointer)	(n + 1)..p		
Service flow encodings Immediate activation requested, upstream. May specify explicit attributes or service class name.		1..m	None yet
Service flow encodings Provisioned for later activation requested, upstream. Explicit attributes or service class name.		(m + 1)..n	None yet
Service flow encodings Immediate activation requested, downstream. Explicit attributes or service name.		(n + 1)..p	None yet
Service flow encodings Provisioned for later activation requested, downstream. Explicit attributes or service name.		(p + 1)..q	None yet

The registration response sets the QoS parameter sets according to the quality of service parameter set type in the registration request.

The registration response preserves the service flow reference attribute, so that the service flow reference can be associated with SFID and/or SID.

Table 10-3 – Registration response contents

Items	Service flow reference	Service flow identifier	Service identifier
Active upstream service flows explicit attributes	1..m	SFID	SID
Provisioned upstream service flows explicit attributes	(m + 1)..n	SFID	Not yet
Active downstream service flows explicit attributes	(n + 1)..p	SFID	N/A
Provisioned downstream service flows explicit attributes	(p + 1)..q	SFID	N/A

The SFID is chosen by the CMTS to identify a downstream or upstream service flow that has been authorized but not activated. A DSC-Request from a modem to admit or activate a provisioned service flow contains its SFID. If it is a downstream flow then the downstream classifier is also included.

10.1.7.2 Dynamic service flow Creation – CM-initiated

Service flows may be created by the dynamic service addition process, as well as through the registration process outlined above. The dynamic service addition may be initiated by either the CM or the CMTS, and may create one upstream and/or one downstream dynamic service flow(s). A three-way handshake is used to create service flows. The CM-initiated protocol is illustrated in Figure 10-6 and described in detail in clause 11.4.2.1.

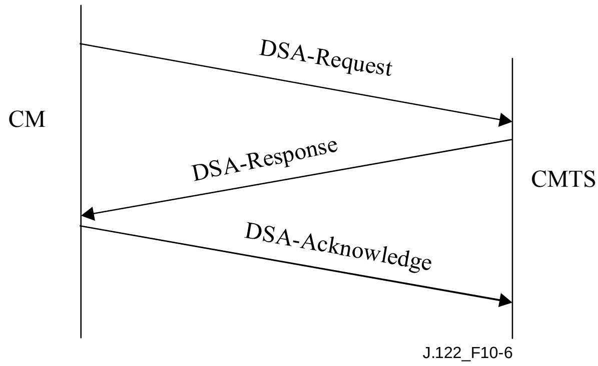

Figure 10-6 – Dynamic service addition message flow – CM-initiated

A DSA-Request from a CM contains service flow reference(s), QoS parameter set(s) (marked either for admission-only or for admission and activation) and any required classifiers.

10.1.7.3 Dynamic service flow creation – CMTS-initiated

A DSA-Request from a CMTS contains service flow identifier(s) for one upstream and/or one downstream service flow, possibly a SID, set(s) of active or admitted QoS parameters, and any required classifier(s). The protocol is as illustrated in Figure 10-7 and is described in detail in clause 11.4.2.2.

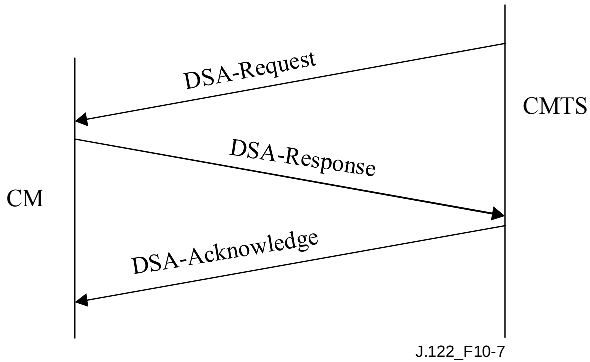

Figure 10-7 – Dynamic service addition message flow – CMTS-initiated

10.1.7.4 Dynamic service flow modification and deletion

In addition to the methods presented above for creating service flows, protocols are defined for modifying and deleting service flows. Refer to clauses 11.4.3 and 11.4.4.

Both provisioned and dynamically created service flows are modified with the DSC message, which can change the admitted and active QoS parameter sets of the flow. The DSC can also add, replace, or delete classifiers, and add, add parameters to, or delete PHS Rules.

A successful DSC transaction changes a service flow's QoS parameters by replacing both the admitted and active QoS parameter sets. If the message contains only the admitted set, the active set is set to null and the flow is deactivated. If the message contains neither set ('000' value used for quality of service parameter set type, see clause C.2.2.3.5) then both sets are set to null and the flow is de-admitted. When the message contains both QoS parameter sets, the admitted set is checked first and, if admission control succeeds, the active set in the message is checked against the admitted set in the message to ensure that it is a subset (see clause 10.1.1.1). If all checks are successful, the QoS parameter sets in the message become the new admitted and active QoS parameter sets for the service flow. If either of the checks fails, the DSC transaction fails and the service flow QoS parameter sets are unchanged.

10.2 Upstream service flow scheduling services

The following clauses define the basic upstream service flow scheduling services and list the QoS parameters associated with each service. A detailed description of each QoS parameter is provided in Annex C. Annex C also discusses how these basic services and QoS parameters can be combined to form new services, such as committed information rate (CIR) service.

Scheduling services are designed to improve the efficiency of the poll/grant process. By specifying a scheduling service and its associated QoS parameters, the CMTS can anticipate the throughput and latency needs of the upstream traffic and provide polls and/or grants at the appropriate times.

Each service is tailored to a specific type of data flow as described below. The basic services comprise: unsolicited grant service (UGS), real-time polling service (rtPS), unsolicited grant service with activity detection (UGS-AD), non-real-time polling service (nrtPS) and best effort (BE) service. Table 10-4 shows the relationship between the scheduling services and the related QoS parameters.

10.2.1 Unsolicited grant service

The unsolicited grant service (UGS) is designed to support real-time service flows that generate fixed size data packets on a periodic basis, such as voice over IP. The service offers fixed-size grants on a real-time periodic basis, which eliminate the overhead and latency of CM requests and assure that grants will be available to meet the flow's real-time needs. The CMTS MUST provide fixed-size data grants at periodic intervals to the service flow. In order for this service to work correctly, the request/transmission policy (refer to clause C.2.2.6.3) setting MUST be such that the CM is prohibited from using any contention request or request/data opportunities and the CMTS SHOULD NOT provide any unicast request opportunities. The request/transmission policy MUST also prohibit piggyback requests. This will result in the CM only using unsolicited data grants for upstream transmission. All other bits of the request/transmission policy are not relevant to the fundamental operation of this scheduling service and should be set according to network policy. The key service parameters are the unsolicited grant size, the nominal grant interval, the tolerated grant jitter and the request/transmission policy (refer to Appendix VI).

The unsolicited grant synchronization header (UGSH) in the service flow EH element (refer to clause 8.2.6.3.2) is used to pass status information from the CM to the CMTS regarding the state of the UGS service flow. The most significant bit of the UGSH is the queue indicator (QI) flag. When the QI flag is set, it indicates a rate overrun condition for the service flow. When the QI flag is clear, it indicates a rate non-overrun condition for the service flow. The QI flag allows the CMTS to provide for a dynamic rate compensation function by issuing additional grants.

The CM MUST set the QI flag when it detects that the packet reception rate is greater than the upstream transmission rate. The CM MUST clear the QI flag when it detects that the packet reception rate is equal to or less than the upstream transmission rate and the queued packet backlog is cleared.

The number of packets already queued for upstream transmission is a measure of the rate differential between received and transmitted packets. The CM SHOULD set the QI flag when the number of packets queued is greater than the number of grants per interval parameter of the active QoS set. The CM SHOULD clear the QI flag when the number of packets queued is less than or equal to the number of grants per interval parameter of the active QoS set. The QI flag of each packet MAY be set either at the time the packet is received and queued or at the time the packet is dequeued and transmitted.

The CM MAY set/clear the QI flag using a threshold of two times the number of grants per interval parameter of the active QoS set. Alternatively, the CM MAY provide hysteresis by setting the QI flag using a threshold of two times the number of grants per interval, then clearing it using a threshold of one times the number of grants per interval.

The CMTS MUST NOT allocate more grants per nominal grant interval than the grants per interval parameter of the active QoS parameter set, excluding the case when the QI bit of the UGSH is set. In this case, the CMTS SHOULD grant up to 1% additional bandwidth for clock rate mismatch compensation. If the CMTS grants additional bandwidth, it MUST limit the total number of bytes forwarded on the flow during any time interval to Max(T), as described in the expression:

$$Max(T) = T \times (R \times 1.01) + 3B$$

where:

Max(T) = the maximum number of bytes transmitted on the flow over a time T (in units of seconds)

R = (grant_size × grants_per_interval)/nominal_grant_interval

B = grant_size × grants_per_interval

The active grants field of the UGSH is ignored with UGS service. The CMTS policing of the service flow remains unchanged.

10.2.2 Real-time polling service

The real-time polling service (rtPS) is designed to support real-time service flows that generate variable-size data packets on a periodic basis, such as MPEG video. The service offers real-time, periodic, unicast request opportunities, which meet the flow's real-time needs and allow the CM to specify the size of the desired grant. This service requires more request overhead than UGS, but supports variable grant sizes for optimum data transport efficiency.

The CMTS MUST provide periodic unicast request opportunities. In order for this service to work correctly, the request/transmission policy setting (refer to clause C.2.2.6.3) SHOULD be such that the CM is prohibited from using any contention request or request/data opportunities. The request/transmission policy SHOULD also prohibit piggyback requests. The CMTS MAY issue unicast request opportunities as prescribed by this service even if a grant is pending. This will result in the CM using only unicast request opportunities in order to obtain upstream transmission opportunities (the CM could still use unsolicited data grants for upstream transmission as well). All other bits of the request/transmission policy are not relevant to the fundamental operation of this scheduling service and should be set according to network policy. The key service parameters are the nominal polling interval, the tolerated poll jitter and the request/transmission policy.

10.2.3 Unsolicited grant service with activity detection

The unsolicited grant service with activity detection (UGS-AD) is designed to support UGS flows that may become inactive for substantial portions of time (i.e., tens of milliseconds or more), such as voice over IP with silence suppression. The service provides unsolicited grants when the flow is active and unicast polls when the flow is inactive. This combines the low overhead and low latency of UGS with the efficiency of rtPS. Though USG/AD combines UGS and rtPS, only one scheduling service is active at a time.

The CMTS MUST provide periodic unicast grants when the flow is active but MUST revert to providing periodic unicast request opportunities when the flow is inactive. (The CMTS can detect flow inactivity by detecting unused grants. However, the algorithm for detecting a flow changing from an active to an inactive state is dependent on the CMTS implementation.) In order for this service to work correctly, the request/transmission policy setting (refer to clause C.2.2.6.3) MUST be such that the CM is prohibited from using any contention request or request/data opportunities. The request/transmission policy MUST also prohibit piggyback requests. This results in the CM using only unicast request opportunities in order to obtain upstream transmission opportunities. However, the CM will use unsolicited data grants for upstream transmission as well. All other bits of the request/transmission policy are not relevant to the fundamental operation of this scheduling service and should be set according to network policy. The key service parameters are the nominal polling interval, the tolerated poll jitter, the nominal grant interval, the tolerated grant jitter, the unsolicited grant size and the request/transmission policy.

In UGS-AD service, when restarting UGS after an interval of rtPS, the CMTS SHOULD provide additional grants in the first (and/or second) grant interval such that the CM receives a total of one grant for each grant interval from the time the CM requested restart of UGS, plus one additional grant (refer to Appendix VI). Because the service flow is provisioned as a UGS flow with a specific grant interval and grant size, when restarting UGS, the CM MUST NOT request a different sized grant than the already provisioned UGS flow. As with any service flow, changes can only be requested with a DSC command. If the restarted activity requires more than one grant per interval, the CM MUST indicate this in the active grants field of the UGSH beginning with the first packet sent.

The service flow extended header element allows for the CM to dynamically state how many grants per interval are required to support the number of flows with activity present. In UGS-AD, the CM MAY use the queue indicator bit in the UGSH. The remaining seven bits of the UGSH define the active grants field. This field defines the number of grants within a nominal grant interval that this service flow currently requires. When using UGS-AD, the CM MUST indicate the number of requested grants per nominal grant interval in this field. The active grants field of the UGSH is ignored with UGS without activity detection. This field allows the CM to signal to the CMTS to dynamically adjust the number of grants per interval that this UGS service flow is actually using. The CM MUST NOT request more than the number of grants per interval in the ActiveQosParamSet.

If the CMTS allocates additional bandwidth in response to the QI bit, it MUST use the same rate-limiting formula as UGS, but the formula only applies to steady state periods where the CMTS has adjusted the grants_per_interval to match the active_grants requested by the CM.

When the CM is receiving unsolicited grants and it detects no activity on the service flow, it MAY send one packet with the active grants field set to zero grants and then cease transmission. Because this packet may not be received by the CMTS, when the service flow goes from inactive to active the CM MUST be able to restart transmission with either polled requests or unsolicited grants.

10.2.4 Non-real-time polling service

The non-real-time polling service (nrtPS) is designed to support non real-time service flows that require variable size data grants on a regular basis, such as high bandwidth FTP. The service offers unicast polls on a regular basis, which assures that the flow receives request opportunities even during network congestion. The CMTS typically polls nrtPS SIDs on an (periodic or non-periodic) interval on the order of one second or less.

The CMTS MUST provide timely unicast request opportunities. In order for this service to work correctly, the request/transmission policy setting (refer to clause C.2.2.6.3) SHOULD be such that the CM is allowed to use contention request opportunities. This will result in the CM using contention request opportunities as well as unicast request opportunities and unsolicited data grants. All other bits of the request/transmission policy are not relevant to the fundamental operation of this scheduling service and should be set according to network policy. The key service parameters are nominal polling interval, minimum reserved traffic rate, maximum sustained traffic rate, request/transmission policy and traffic priority.

10.2.5 Best effort service

The intent of the best effort (BE) service is to provide efficient service to best effort traffic. In order for this service to work correctly, the request/transmission policy setting SHOULD be such that the CM is allowed to use contention request opportunities. This will result in the CM using contention request opportunities as well as unicast request opportunities and unsolicited data grants. All other bits of the request/transmission policy are not relevant to the fundamental operation of this scheduling service and should be set according to network policy. The key service parameters are the minimum reserved traffic rate, the maximum sustained traffic rate, and the traffic priority.

10.2.6 Other services

10.2.6.1 Committed information rate (CIR)

A committed information rate (CIR) service can be defined a number of different ways. For example, it could be configured by using a best effort service with a minimum reserved traffic rate or a nrtPS with a minimum reserved traffic rate.

10.2.7 Parameter applicability for upstream service scheduling

Table 10-4 summarizes the relationship between the scheduling services and key QoS parameters. A detailed description of each QoS parameter is provided in Annex C.

Table 10-4 – Parameter applicability for upstream service scheduling

Service flow parameter	Best effort	Non-real-time polling	Real-time polling	Unsolicited grant	Unsolicited grant with activity detection
Miscellaneous					
• Traffic priority	Optional, default = 0	Optional default = 0	N/A[a)]	N/A	N/A
• Max. concatenated burst	Optional	Optional	Optional	N/A	N/A
• Upstream scheduling service type	Optional, default = 2	Mandatory	Mandatory	Mandatory	Mandatory
• Request/transmission policy	Optional, default = 0	Mandatory	Mandatory	Mandatory	Mandatory
Maximum rate					
• Max. sustained traffic rate	Optional, default = 0	Optional default = 0	Optional default = 0	N/A	N/A
• Max. traffic burst	Optional, default = 1522	Optional default = 1522	Optional default = 1522	N/A	N/A
Minimum rate					
• Min. reserved traffic rate	Optional, default = 0	Optional default = 0	Optional default = 0	N/A	N/A
• Assumed minimum packet size	Optional[c)]	Optional[c)]	Optional[c)]	Optional[c)]	Optional[c)]
Grants					
• Unsolicited grant size	N/A	N/A	N/A	Mandatory	Mandatory
• Grants per interval	N/A	N/A	N/A	Mandatory	Mandatory
• Nominal grant interval	N/A	N/A	N/A	Mandatory	Mandatory

Table 10-4 – Parameter applicability for upstream service scheduling

Service flow parameter	Best effort	Non-real-time polling	Real-time polling	Unsolicited grant	Unsolicited grant with activity detection
• Tolerated grant jitter	N/A	N/A	N/A	Mandatory	Mandatory
Polls					
• Nominal polling interval	N/A	Optional[c)]	Mandatory	N/A	Optional[b)]
• Tolerated poll jitter	N/A	N/A	Optional[c)]	N/A	Optional[c)]

a) N/A means not applicable to this service flow scheduling type. If included in a request for a service flow of this service flow scheduling type, this request MUST be denied.

b) Default is same as nominal grant interval.

c) Default is CMTS-specific.

10.2.8 CM transmit behaviour

In order for these services to function correctly, all that is required of the CM in regards to its transmit behaviour for a service flow is for it to follow the rules specified in clause 9.4.3 and the request/transmission policy specified for the service flow.

10.3 Fragmentation

Fragmentation is an upstream CM "modem capability". The CMTS MUST enable or disable this capability on a per-modem basis with a TLV in the registration response. The per-modem basis provides compatibility with DOCS 1.0 CMs. Once fragmentation is enabled for a DOCS 1.1 modem, fragmentation is enabled on a per-service flow basis via the request/transmission policy configuration settings. When enabled for a service flow, fragmentation is initiated by the CMTS when it grants bandwidth to a particular CM with a grant size that is smaller than the corresponding bandwidth request from the CM. This is known as a **partial grant.**

10.3.1 CM fragmentation support

Fragmentation is essentially encapsulation of a portion of a MAC frame within a fixed size fragmentation header and a fragment CRC. Concatenated PDUs, as well as single PDUs, are encapsulated in the same manner. Baseline privacy, if enabled, is performed on each fragment as opposed to the complete original MAC frame.

The CM MUST perform fragmentation according to the flow diagram in Figure 10-8. The phrase "untransmitted portion of packet" in the flow diagram refers to the entire MAC frame when fragmentation has not been initiated, and to the remaining untransmitted portion of the original MAC frame when fragmentation has been initiated.

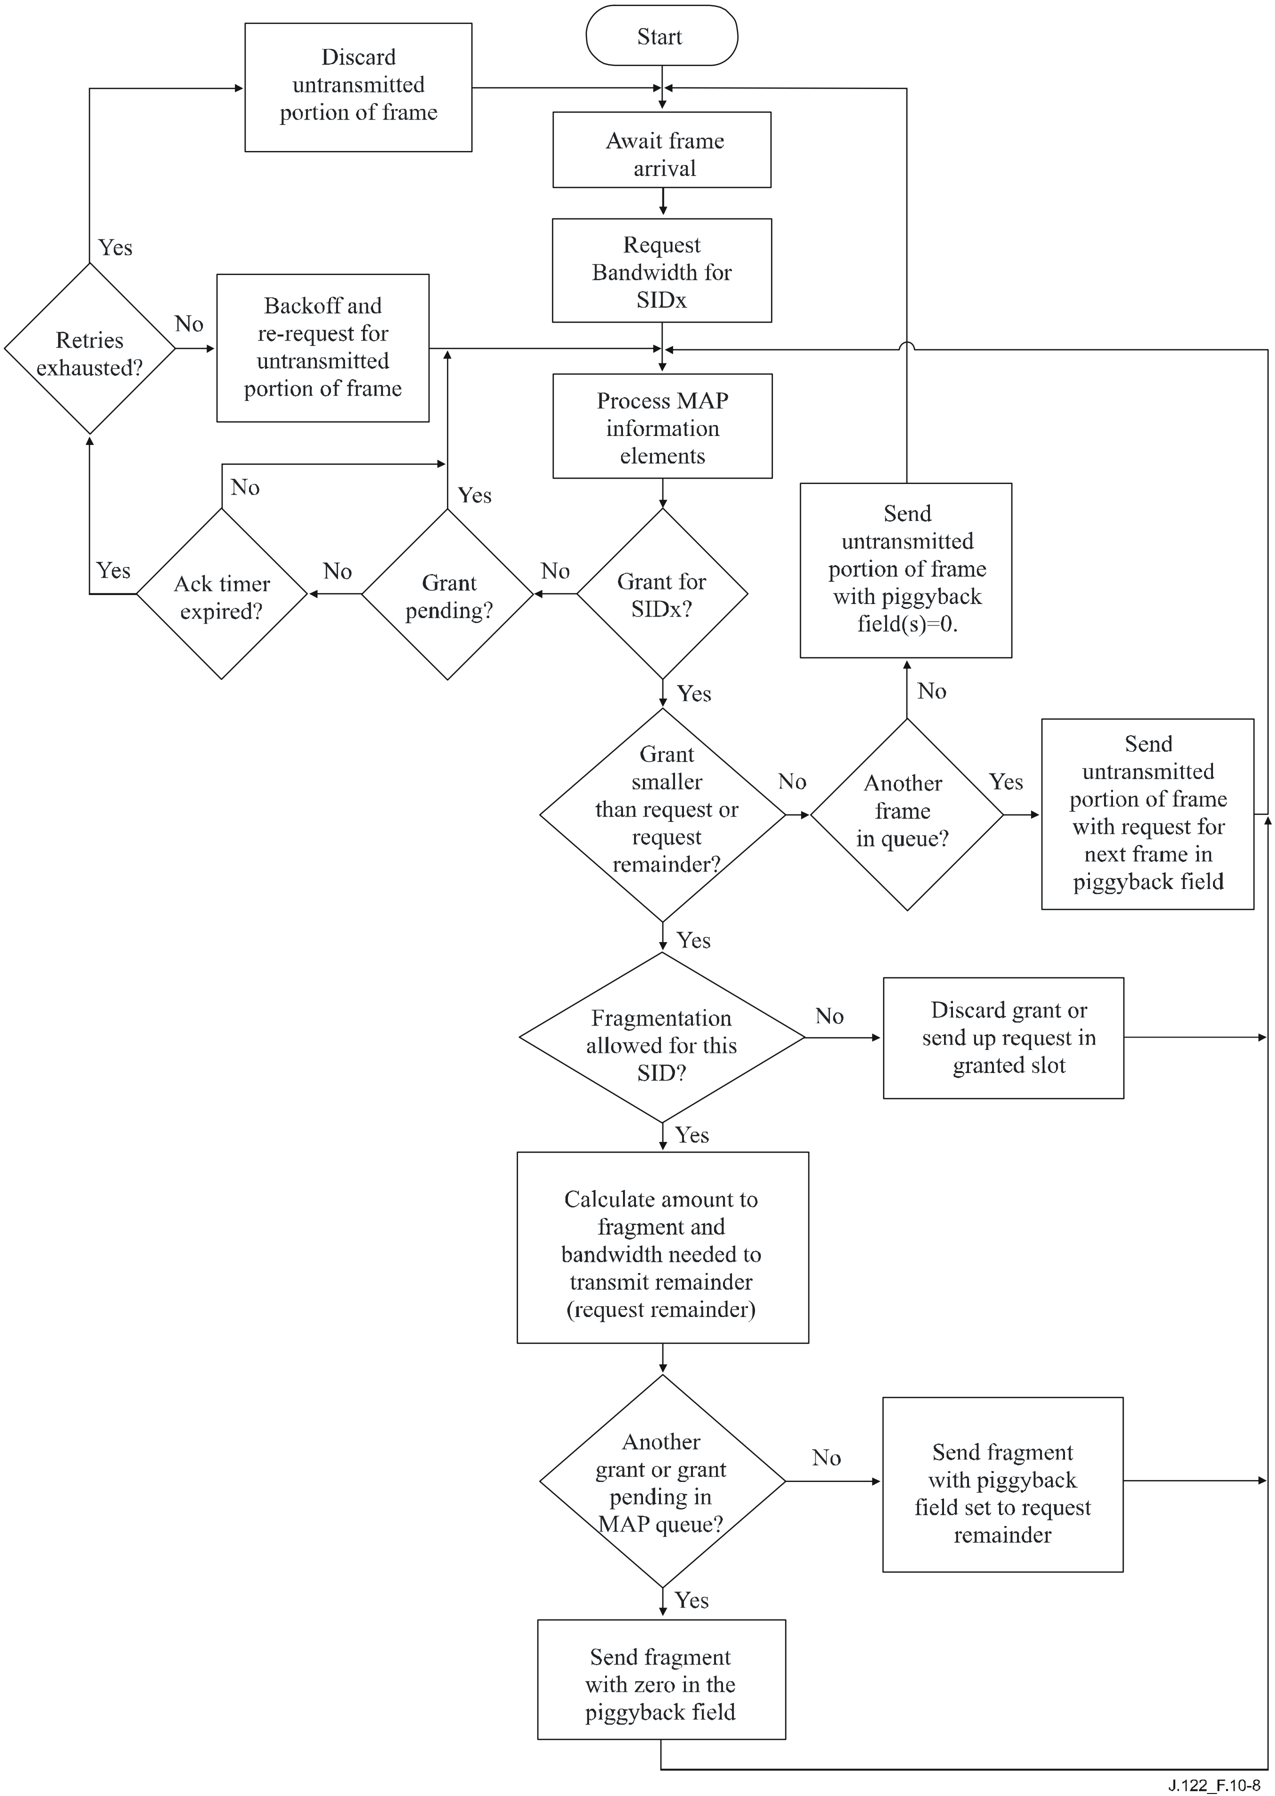

Figure 10-8 – CM fragmentation flowchart

10.3.1.1 Fragmentation rules

1) Any time fragmentation is enabled and the grant size is smaller than the request, the CM MUST fill the partial grant it receives with the maximum amount of data (fragment payload) possible accounting for fragmentation overhead and physical layer overhead.
2) The CM MUST send up a piggyback request any time there is no later grant or grant pending for that SID in MAPs that have been received at the CM.
3) If the CM is fragmenting a frame[25], any piggyback request for the next fragment MUST be made in the BPI EHDR portion of the fragment header. Any piggyback request for a subsequent frame SHOULD be made in the BPI EHDR portion of the last fragment, but MAY be made in one of the extended headers inside the original frame. However, the same request MUST NOT be made in more than one place. Because the CMTS could ignore a request inside the original frame, making the request in the original frame may cause a loss of the request.
4) In calculating bandwidth requests for the remainder of the frame (concatenated frame, if concatenated) that has been fragmented, the CM MUST request enough bandwidth to transmit the entire remainder of the frame plus the 16-byte fragment overhead and all associated physical layer overhead.
5) If the CM does not receive a grant or grant pending within the ack time of sending a request, the CM MUST backoff and re-request for the untransmitted portion of the frame until the bandwidth is granted or the CM exceeds its retry threshold.
6) If the CM exceeds its retry threshold while requesting bandwidth, the CM discards whatever portion of the frame was not previously transmitted.
7) The CM MUST set the F bit and clear the L bit in the first fragment of a frame.
8) The CM MUST clear the F and L bits in the fragment header for any fragments that occur between the first and last fragments of a frame.
9) The CM MUST set the L bit and clear the F bit in the last fragment of a frame.
10) The CM MUST increment the fragment sequence number sequentially for each fragment of a frame transmitted.
11) If a frame is to be encrypted and the frame is fragmented, the frame is encrypted only at the fragment layer with encryption beginning immediately after the fragment header HCS and continuing through the fragment CRC.
12) Frames sent in immediate data (request/data) regions MUST NOT be fragmented.

10.3.2 CMTS fragmentation support

At the CMTS, the fragment is processed similarly to an ordinary packet with the exception that the baseline privacy encryption starts just after the fragmentation header as opposed to being offset by 12 bytes.

The CMTS has two modes it can use to perform fragmentation. The multiple grant mode assumes that the CMTS retains the state of the fragmentation. This mode allows the CMTS to have multiple partial grants outstanding for any given SID. The piggyback mode assumes the CMTS does NOT retain any fragmentation state. Only one partial grant is outstanding, so that the CM inserts the remaining amount into the piggyback field of the fragment header. The type of mode being used is determined by the CMTS. In all cases, the CM operates with a consistent set of rules.

[25] 'Frame' always refers to either frames with a single packet PDU or concatenated frames.

10.3.2.1 Multiple grant mode

A CMTS MUST support multiple grant mode or piggyback mode for performing fragmentation. A CMTS MAY support both fragmentation modes.

Multiple grant mode allows the CMTS to break a request up into two or more grants in a single or over successive maps and it calculates the additional overhead required in the remaining partial grants to satisfy the request. In multiple grant mode, if the CMTS cannot grant the remainder in the current MAP, it MUST send a grant pending (zero length grant) in the current MAP and all subsequent MAPs to the CM until it can grant additional bandwidth. If there is no grant or grant pending in subsequent MAPs, the CM MUST re-request for the remainder. This re-request mechanism is the same as that used when a normal REQ does not receive a grant or grant pending within the ack time.

If a CM receives a grant pending IE along with a fragment grant, it MUST NOT piggyback a request in the extended header of the fragment transmitted in that grant.

In the case where the CM misses a grant and re-requests the remaining bandwidth, the CMTS MUST recover without dropping the frame.

Due to the imprecision of the mini-slot to byte conversion process, the CMTS may not be able to calculate exactly the number of extra mini-slots needed to allow for fragmentation overhead. Also because it is possible for a CM to have missed a map with a partial grant, and thus to be requesting to send an unsent fragment rather than a new PDU, the CMTS cannot be certain whether the CM has already accounted for fragmentation overhead in a request. Therefore, the CMTS MUST make sure that any fragment payload remainder is at least one mini-slot greater than the number of mini-slots needed to contain the overhead for a fragment (16 bytes) plus the physical layer overhead necessary to transmit a minimum-sized fragment. Failure to do this may cause the CMTS to issue a grant that is not needed as the CM has completed transmission of the fragment payload remainder using the previous partial grant. This may cause the CM to get out of sync with the CMTS by inadvertently starting a new fragmentation. Also the CMTS needs to deal with the fact that with certain sets of physical layer parameters, the CM may request one more mini-slot than the maximum size of a short data grant, but not actually need that many mini-slots. This happens in the case where the CM needs to push the request size beyond the short data grant limit. The CMTS needs a policy to ensure that fragmenting such requests in multiple grant mode does not lead to unneeded fragmentary grants.

10.3.2.2 Piggyback mode

If the CMTS does not put another partial grant or a grant pending in the MAP in which it initiates fragmentation on a SID, the CM MUST automatically piggyback for the remainder. The CM calculates how much of a frame can be sent in the granted bandwidth and forms a fragment to send it. The CM utilizes the piggyback field in the fragment extended header to request the bandwidth necessary to transfer the remainder of the frame. Since the CMTS did not indicate a multiple grant in the first fragment MAP, the CM MUST keep track of the remainder to send. The request length, including physical-layer and fragmentation overhead, for the remainder of the original frame is inserted into the piggyback request byte in the fragmentation header.

If the fragment HCS is correct, the piggybacked request, if present, is passed on to the bandwidth allocation process while the fragment itself is enqueued for reassembly. Once the complete MAC frame is reassembled, and it has been determined that the HCS is correct, the CMTS processes the frame as though it had been received unfragmented except that the CMTS MUST ignore the decryption-related portion of any privacy EHDRs. However, the bandwidth requests in privacy EHDRs and request EHDRs of such frames SHOULD be processed, but they MAY be ignored also.

10.3.3 Fragmentation example

10.3.3.1 Single packet fragmentation

Refer to Figure 10-9. Assume that fragmentation has been enabled for a given SID.

1) (Requesting state) – CM wants to transmit a 1018-byte packet. CM calculates how much physical layer overhead (POH) is required and requests the appropriate number of mini-slots. CM makes a request in a contention region. Go to step 2.

2) (Waiting for grant) – CM monitors MAPs for a grant or grant pending for this SID. If the CM's ack time expires before the CM receives a grant or grant pending, the CM retries requesting for the packet until the retry count is exhausted – then the CM gives up on that packet. Go to step 3.

3) (First fragment) – Prior to giving up in step 2, the CM sees a grant for this SID that is less than the requested number of mini-slots. The CM calculates how much MAC information can be sent in the granted number of mini-slots using the specified burst profile. In the example in Figure 10-9, the first grant can hold 900 bytes after subtracting the POH. Since the fragment overhead (FRAG HDR, FHCS and FCRC) is 16 bytes, 884 bytes of the original packet can be carried in the fragment. The CM creates a fragment composed of the FRAG HDR, FHCS, 884 bytes of the original packet, and an FCRC. The CM marks the fragment as first and prepares to send the fragment. Go to step 4.

4) (First fragment, multiple grant mode) – CM looks to see if there are any other grants or grant pendings enqueued for this SID. If so, the CM sends the fragment with the piggyback field in the FRAG HDR set to zero and awaits the time of the subsequent grant to roll around. Go to step 6. If there are not any grants or grant pendings, go to step 5.

5) (First fragment, piggyback mode) – If there are no other grants or grant pendings for this SID in this MAP, the CM calculates how many mini-slots are required to send the remainder of the fragmented packet, including the fragmentation overhead, and physical layer overhead, and inserts this amount into the piggyback field of the FRAG HDR. The CM then sends the fragment and starts its ack timer for the piggyback request. In the example in Figure 10-9, the CM sends up a request for enough mini-slots to hold the POH plus 150 bytes (1018 – 884 + 16). Go to step 6.

6) (Waiting for grant) – The CM is now waiting for a grant for the next fragment. If the CM's ack timer expires while waiting on this grant, the CM should send up a request for enough mini-slots to send the remainder of the fragmented packet, including the fragmentation overhead, and physical layer overhead. Go to step 7.

7) (Receives next fragment grant) – Prior to giving up in step 6, the CM sees another grant for this SID. The CM checks to see if the grant size is large enough to hold the remainder of the fragmented packet, including the fragmentation overhead and physical layer overhead. If so, go to step 10. If not, go to step 8.

8) (Middle fragment, multiple grant mode) – Since the remainder of the packet (plus overhead) will not fit in the grant, the CM calculates what portion will fit. The CM encapsulates this portion of the packet as a middle fragment. The CM then looks for any other grants or grant pendings enqueued for this SID. If either are present, the CM sends the fragment with the piggyback field in the FRAG HDR set to zero and awaits the time of the subsequent grant to roll around. Go to step 6. If there are not any grants or grant pendings, go to step 9.

9) (Middle fragment, piggyback mode) – The CM calculates how many mini-slots are required to send the remainder of the fragmented packet, including the fragmentation overhead and physical layer overhead, and inserts this amount into the piggyback field of the FRAG HDR. The CM then sends the fragment and starts its ack timer for the piggyback request. Go to step 6.

10) (Last fragment) – The CM encapsulates the remainder of the packet as a last fragment. If there is no other packet enqueued or there is a another grant or a grant pending enqueued for this SID, the CM places a zero in the REQ field of the FRAG HDR. If there is another packet enqueued with no grant of grant pending, the CM calculates the number of mini-slots required to send the next packet and places this number in the REQ field in the FRAG HDR. The CM then transmits the packet. Go to step 11. In the example in Figure 10-9, the grant is large enough to hold the remaining 150 bytes plus POH.

11) (Normal operation) – The CM then returns the normal operation of waiting for grants and requesting for packets. If at any time fragmentation is enabled and a grant arrives that is smaller than the request, the fragmentation process starts again as in step 2.

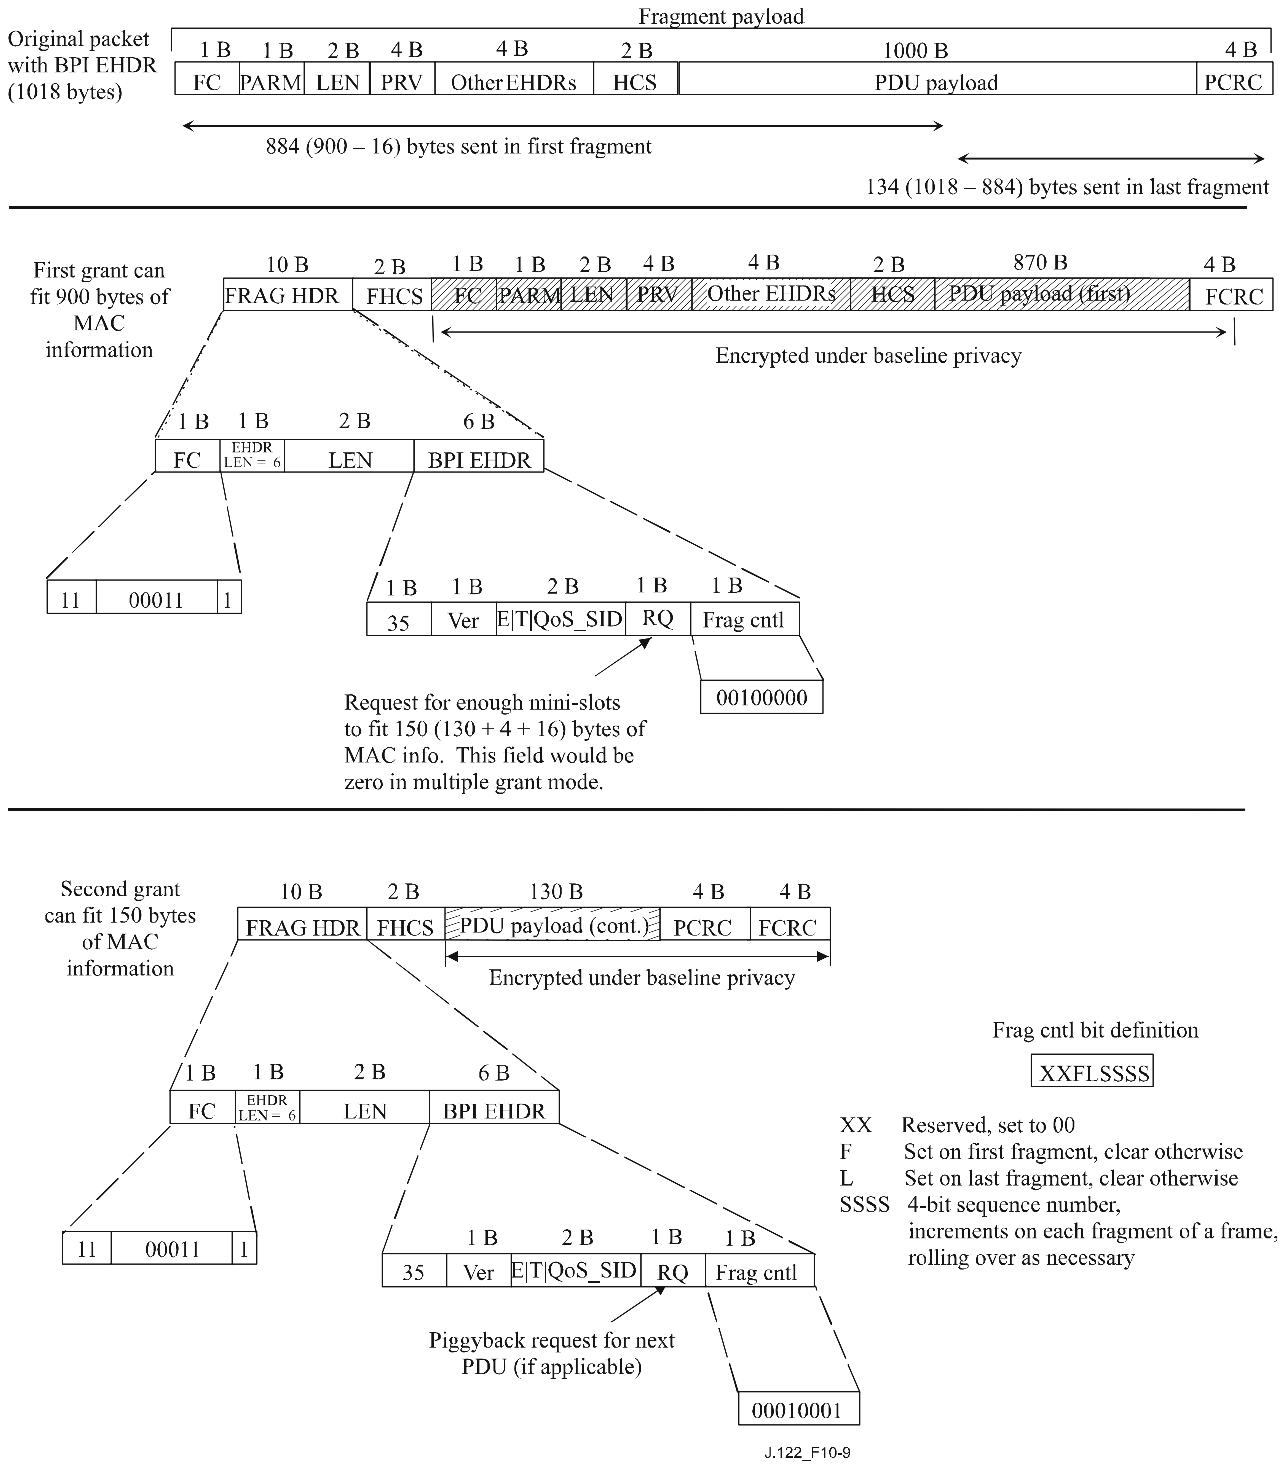

Figure 10-9 – Example of fragmenting a single packet

10.3.3.2 Concatenated packet fragmentation

After the CM creates the concatenated packet, the CM treats the concatenated packet as a single PDU. Figure 10-10 shows an example of a concatenated packet broken into 3 fragments. Note that the packet is fragmented without regard to the packet boundaries within the concatenated packet.

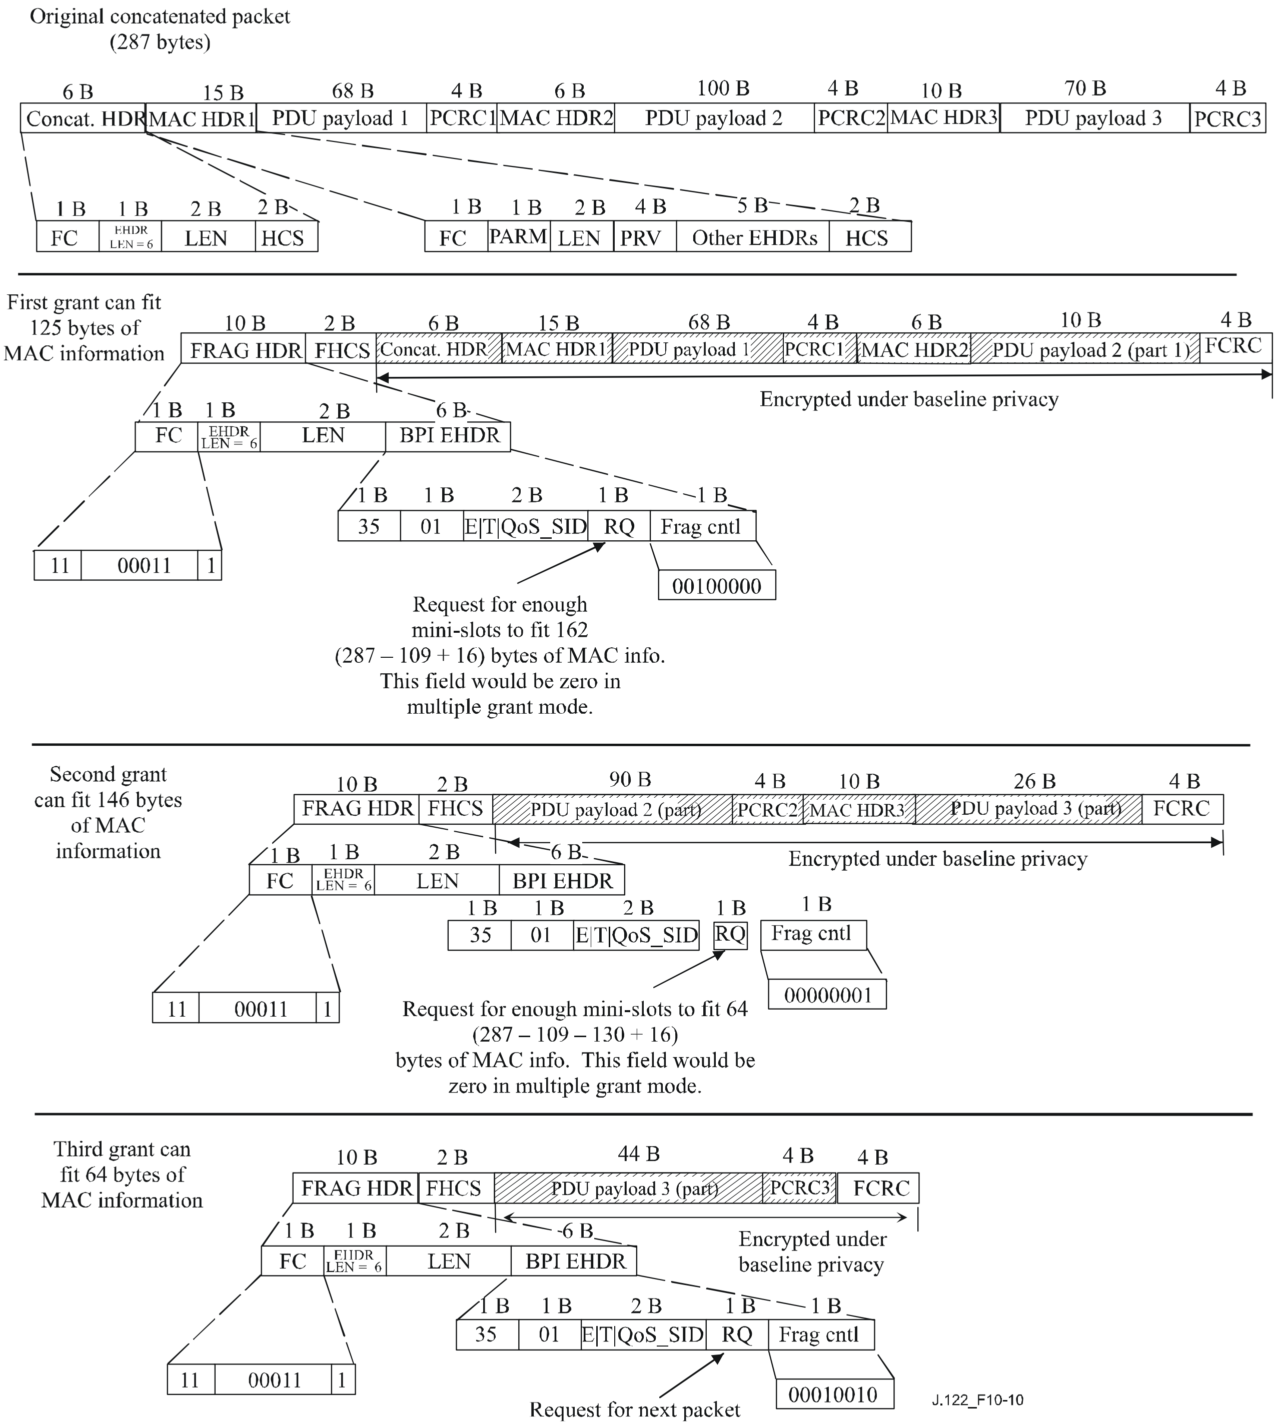

Figure 10-10 – Fragmented concatenated packet example

10.4 Payload header suppression

This overview clause explains the principles of payload header suppression. The subsequent clauses explain the signalling for initialization, operation and termination. Finally, specific upstream and downstream examples are given. The definitions shown in Table 10-5 are used.

Table 10-5 – Payload header suppression definitions

PHS	Payload header suppression	Suppressing an initial byte string at the sender and restoring the byte string at the receiver
PHS rule	Payload header suppression rule	A set of TLVs that apply to a specific PHS index
PHSF	Payload header suppression field	A string of bytes representing the header portion of a PDU in which one or more bytes will be suppressed (i.e., a snapshot of the uncompressed PDU header inclusive of suppressed and unsuppressed bytes)
PHSI	Payload header suppression index	An 8-bit value that references the suppressed byte string
PHSM	Payload header suppression mask	A bit mask that indicates which bytes in the PHSF to suppress and which bytes to not suppress
PHSS	Payload header suppression size	The length of the suppressed field in bytes. This value is equivalent to the number of bytes in the PHSF and also the number of valid bits in the PHSM
PHSV	Payload header suppression verify	A flag that tells the sending entity to verify all bytes that are to be suppressed

10.4.1 Overview

In payload header suppression, a repetitive portion of the payload headers following the extended header field is suppressed by the sending entity and restored by the receiving entity. In the upstream, the sending entity is the CM and the receiving entity is the CMTS. In the downstream, the sending entity is the CMTS and the receiving entity is the CM. The MAC extended header contains a payload header suppression index (PHSI) that references the payload header suppression field (PHSF).

Although PHS may be used with any service flow type, it has been designed for use with the unsolicited grant service (UGS) scheduling type. UGS works most efficiently with packets of a fixed length. PHS works well with UGS because, unlike other header compression schemes sometimes used with IP data, PHS always suppresses the same number of bytes in each packet. PHS will always produce a fixed length compressed packet header.

The sending entity uses classifiers to map packets into a service flow. The classifier uniquely maps packets to its associated payload header suppression rule. The receiving entity uses the service identifier (SID) and the PHSI to restore the PHSR.

Once the PHSF and PHSS fields of a rule are known, the rule is considered "fully defined" and none of its fields can be changed. If modified PHS operation is desired for packets classified to the flow, the old rule must be removed from the service flow and a new rule must be installed.

When a classifier is deleted, any associated PHS rule MUST also be deleted.

PHS has a PHSV option to verify or not verify the payload before suppressing it. PHS also has a PHSM option to allow select bytes not to be suppressed. This is used for sending bytes that change, such as IP sequence numbers, and still suppressing bytes that do not change.

PHS rules are consistent for all scheduling service types. Requests and grants of bandwidth are specified after suppression has been accounted for. For unsolicited grant services, the grant size is chosen with the unsolicited grant size TLV. The packet with its header suppressed may be equal to or less than the grant size.

The CMTS MUST assign all PHSI values just as it assigns all SID values. Either the sending or the receiving entity MAY specify the PHSF and PHSS. This provision allows for pre-configured headers, or for higher level signalling protocols outside the scope of this Recommendation to establish cache entries. PHS is intended for unicast service, and is not defined for multicast service.

It is the responsibility of the higher-layer service entity to generate a PHS rule that uniquely identifies the suppressed header within the service flow. It is also the responsibility of the higher-layer service entity to guarantee that the byte strings being suppressed are constant from packet to packet for the duration of the active service flow.

10.4.2 Example applications

- A classifier on an upstream service flow that uniquely defines a voice over IP (VoIP) flow by specifying protocol type of UDP, IP SA, IP DA, UDP source port, UDP destination port, the service flow reference, and a PHS Size of 42 bytes. A PHS rule references this classifier providing a PHSI value that identifies this VoIP media flow. For the upstream case, 42 bytes of payload header are verified and suppressed, and a 2-byte extended header containing the PHSI is added to every packet in that media flow.
- A classifier that identifies the packets in a service flow, of which 90% match the PHSR. Verification is enabled. This may apply in a packet compression situation where every so often compression resets are done and the header varies. In this example, the scheduling algorithm would allow variable bandwidth, and only 90% of the packets might get their headers suppressed. Since the existence of the PHSI extended header will indicate the choice made, the simple SID/PHSI lookup at the receiving entity will always yield the correct result.
- A classifier on an upstream service flow that identifies all IP packets by specifying ethertype of IP, the service flow ID, a PHSS of 14 bytes, and no verification by the sending entity. In this example, the CMTS has decided to route the packet, and knows that it will not require the first 14 bytes of the Ethernet header, even though some parts such as the source address or destination address may vary. The CM removes 14 bytes from each upstream frame (Ethernet header) without verifying their contents and forwards the frame to the service flow.

10.4.3 Operation

To clarify operational packet flow, this clause describes one potential implementation. CM and CMTS implementations are free to implement payload header suppression in any manner as long as the protocol specified in this clause is followed. Figure 10-11 illustrates the following procedure.

A packet is submitted to the CM MAC service layer. The CM applies its list of classifier rules. A match of the rule will result in an upstream service flow, SID and a PHS rule. The PHS rule provides PHSF, PHSI, PHSM, PHSS and PHSV. If PHSV is set to zero, or is not present, the CM will compare the bytes in the packet header with the bytes in the PHSF that are to be suppressed as indicated by the PHSM. If they match, the CM will suppress all the bytes in the upstream suppression field except the bytes masked by PHSM. The CM will then insert the PHSI into the PHS_Parm field of the service flow EH element, and queue the packet on the upstream service flow.

When the packet is received by the CMTS, the CMTS will determine the associated SID either by internal means or from other extended headers elements such as the BPI extended header. The CMTS uses the SID and the PHSI to look up PHSF, PHSM and PHSS. The CMTS reassembles the packet and then proceeds with normal packet processing. The reassembled packet will contain bytes from the PHSF. If verification was enabled, then the PHSF bytes will equal the original header byes. If verification was not enabled, then there is no guarantee that the PHSF bytes will match the original header bytes.

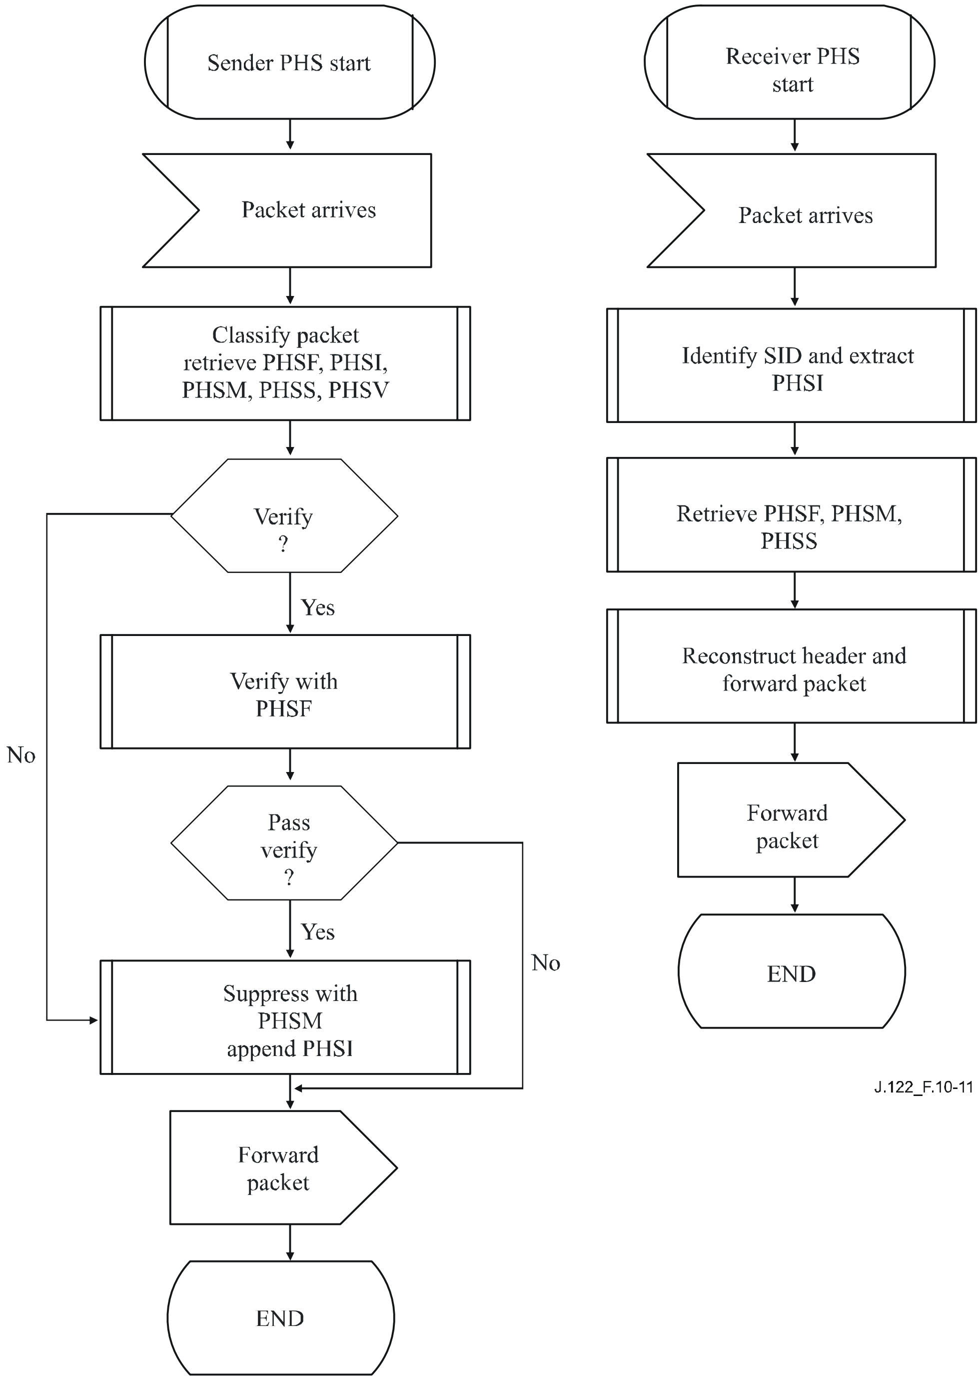

Figure 10-11 – Payload header suppression operation

A similar operation occurs in the downstream. The CMTS applies its list of classifiers. A match of the classifier will result in a downstream service flow and a PHS rule. The PHS rule provides PHSF, PHSI, PHSM, PHSS and PHSV. If PHSV is set to zero, or is not present, the CMTS will verify the downstream suppression field in the packet with the PHSF. If they match, the CMTS will suppress all the bytes in the downstream suppression field except the bytes masked by PHSM. The CMTS will then insert the PHSI into the PHS_Parm field of the service flow EH element, and queue the packet on the downstream service flow.

The CM will receive the packet based upon the Ethernet destination address filtering. The CM then uses the PHSI to lookup PHSF, PHSM and PHSS. The CM reassembles the packet and then proceeds with normal packet processing.

Figure 10-12 demonstrates packet suppression and restoration when using PHS masking. Masking allows only bytes that do not change to be suppressed. Note that the PHSF and PHSM span the entire suppression field, including suppressed and unsuppressed bytes.

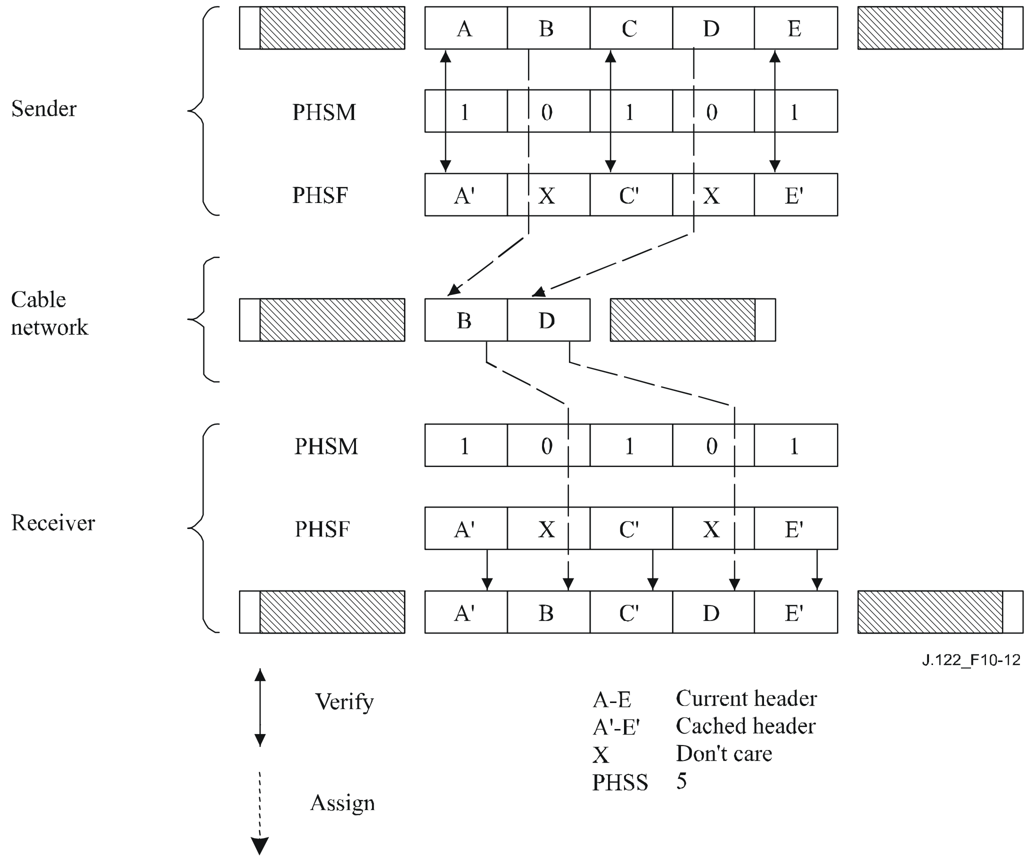

Figure 10-12 – Payload header suppression with masking

10.4.4 Signalling

Payload header suppression requires the creation of three objects:

- service flow;
- classifier;
- payload header suppression rule.

These three objects MAY be created in separate message flows, or MAY be created simultaneously.

PHS rules are created with registration, DSA or DSC messages. The CMTS MUST define the PHSI when the PHS rule is created. PHS rules are deleted with the DSC or DSD messages. The CM or CMTS MAY define the PHSS and PHSF.

Figure 10-13 shows the two ways to signal the creation of a PHS Rule.

It is possible to partially define a PHS rule (in particular, the size of the rule) at the time a service flow is created.

As an example, it is likely that when a service flow is first provisioned, the size of the header field to be suppressed will be known. The values of some items within the field (e.g., IP addresses, UDP port numbers, etc.) may not be known and would be provided in a subsequent DSC as part of the activation of the service flow (using the "set PHS rule" DSC action).

A PHS rule is partially defined when the PHSF and PHSS field values are not both known. Once both PHSF and PHSS are known, the rule is considered fully defined, and MUST NOT be modified via DSC signalling. PHSV and PHSM fields have default values, thus are not required to fully define a PHS rule. If PHSV and PHSM are not known when the rule becomes fully defined, their default values are used, and MUST NOT be modified via DSC signalling.

Each step of the PHS rule definition, whether it is a registration request, DSA or a DSC, MUST contain service flow ID (or reference), classifier ID (or reference) to uniquely identify the PHS rule being defined. A PHS index and service ID pair is used to uniquely identify the PHS rule during upstream packet transfer. A PHS index is enough to uniquely identify the PHS rule used in downstream packet transfer.

10.4.5 Payload header suppression examples

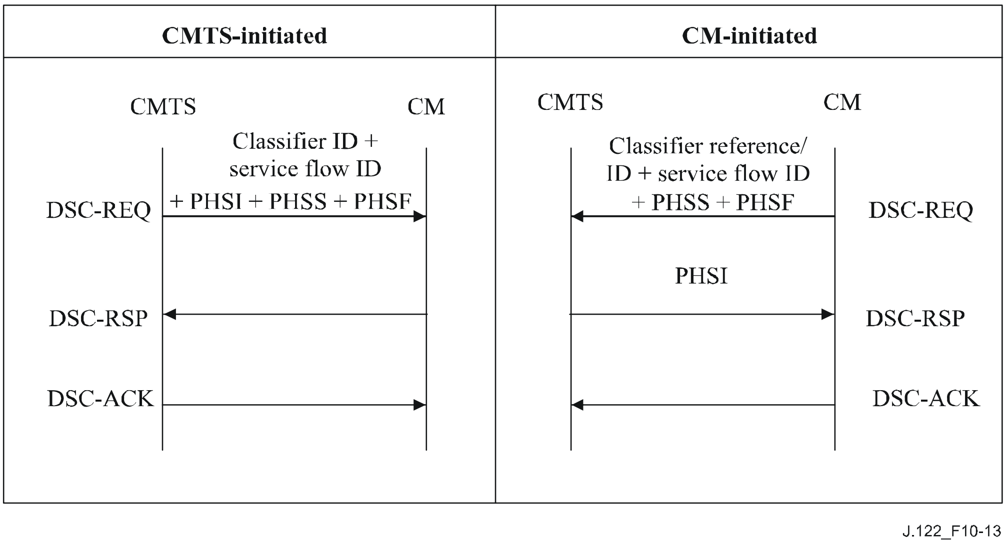

Figure 10-13 – Payload header suppression signalling example

10.4.5.1 Upstream example

A service class with the service class name of "G711-US-UGS-HS-42" is established, which is intended for G.711 VoIP traffic in the upstream with unsolicited grant service. When classifiers are added to the flow, a PHSS value of 42 is included, which explicitly states that the first 42 bytes following the MAC extended header on all packets in that flow must be verified, suppressed and restored. In this example, the service class is configured such that any packet that does not verify correctly will not have its header suppressed and will be discarded since it will exceed the unsolicited grant size (refer to clause C.2.2.6.3).

Figure 10-14 shows the encapsulation used in the upstream with and without payload header suppression. An RTP voice over IP payload without IPsec is used as a specific example to demonstrate efficiency.

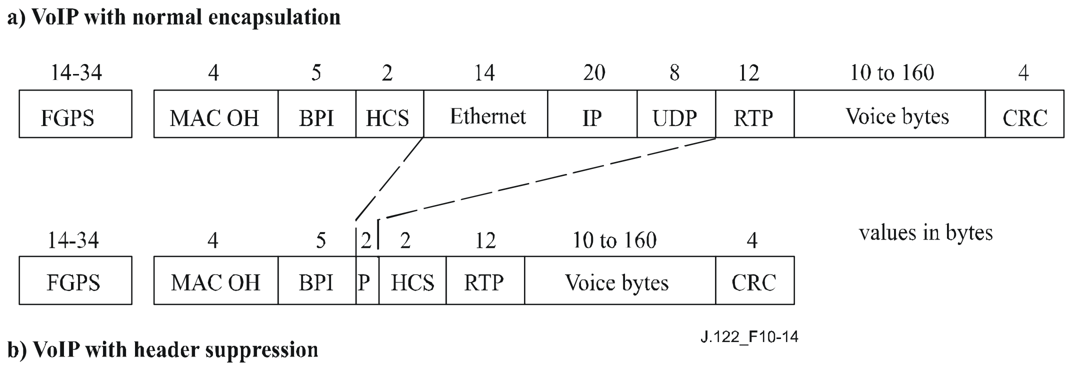

Figure 10-14 – Upstream payload header suppression example

Figure 10-14 a) shows a normal RTP packet carried on an upstream channel. The beginning of the frame represents the physical layer overhead (FGPS) of FEC, guard time, preamble and stuffing bytes. Stuffing bytes occur in the last code word and when mapping blocks to mini-slots. Next is the MAC layer overhead including the 6-byte MAC header with a 5-byte BPI extended header, the 14-byte Ethernet header, and the 4-byte Ethernet CRC trailer. The VoIP payload uses a 20-byte IP header, an 8-byte UDP header, and a 12-byte RTP header. The voice payload is variable and depends upon the sample time and the compression algorithm used.

Figure 10-14 b) shows the same payload with payload header suppression enabled. In the upstream, payload header suppression begins with the first byte after the MAC header checksum. The 14-byte Ethernet header, the 20-byte IP header, and the 8-byte UDP header have been suppressed, and a 2-byte PHS extended header element has been added, for a net reduction of 40 bytes. In this example of an established VoIP connection, these fields remain constant from packet to packet, and are otherwise redundant.

10.4.5.2 Downstream example

A service class with the service class name of "G711-DS-HS-30" is established, which is intended for G.711 VoIP traffic in the downstream. When classifiers are added to the service flow, a PHSS value of 30 is included which explicitly indicates that 30 bytes of the payload header on all packets must be processed for suppression and restoration according to the PHSM. Any packet that does not verify correctly will not have its header suppressed but will be transmitted subject to the traffic shaping rules in place for that service flow.

Figure 10-15 shows the encapsulation used in the downstream with and without payload header suppression. An RTP voice over IP payload without IPsec is used as a specific example to demonstrate efficiency.

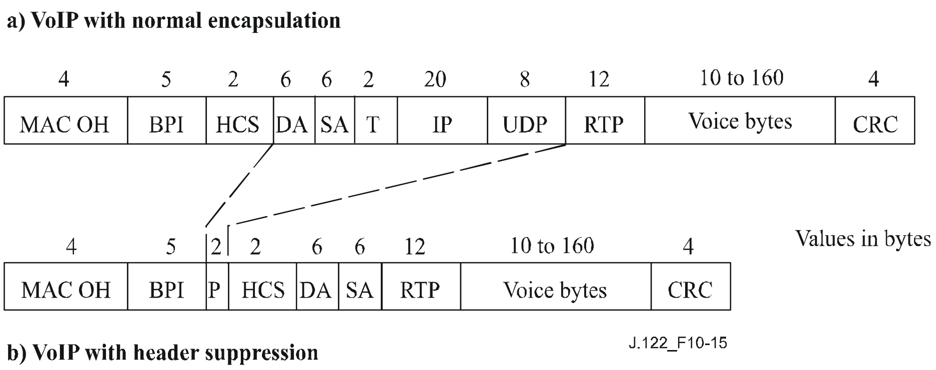

Figure 10-15 – Downstream payload header suppression example

Figure 10-15 a) shows a normal RTP packet carried on a downstream channel. The layer 2 overhead includes the 6-byte MAC header with a 5-byte BPI extended header, the 14-byte Ethernet header (6-byte destination address, 6-byte source address, and 2-byte ethertype field), and the 4-byte Ethernet CRC trailer. The layer 3 VoIP payload uses a 20-byte IP header, an 8-byte UDP header, and a 12-byte RTP header. The voice payload is variable and depends upon the sample time and the compression algorithm used.

Figure 10-15 b) shows the same payload with payload header suppression enabled. In the downstream, payload header suppression begins with the thirteenth byte after the MAC header checksum. This retains the Ethernet destination address and source address, which is required so that the CM may filter and receive the packet. The remaining 2 bytes of the Ethernet header, the 20-byte IP header, and the 8-byte UDP header have been suppressed, and a 2-byte PHS extended header element has been added, for a net reduction of 28 bytes. In this example of an established VoIP connection, these fields remain constant from packet to packet, and are thus redundant.

11 Cable modem – CMTS interaction

This clause covers the key requirements for the interaction between a CM and a CMTS. It only covers interaction between a CM and CMTS that are both compliant with this version of the Recommendation, and only in the case where the CM is using a configuration file with QoS parameters compliant with this version of the Recommendation. Issues involving interoperability with equipment and configuration files compliant with previous versions of this Recommendation (DOCS 1.x) are discussed in Annex G. The interaction can be broken down into five basic categories: initialization, authentication, configuration, authorization and signalling.

11.1 CMTS initialization

The mechanism utilized for CMTS initialization (local terminal, file download, SNMP, etc.) is described in [DOCS5]. It MUST meet the following criteria for system interoperability:

- The CMTS MUST be able to reboot and operate in a stand-alone mode using configuration data retained in non-volatile storage.
- If valid parameters are not available from non-volatile storage or via another mechanism such as the spectrum management system (see [b-VECCHI]), the CMTS MUST NOT generate any downstream messages (including SYNC). This will prevent CMs from transmitting.
- The CMTS MUST provide the information defined in clause 8 to CMs for each upstream channel.

11.2 Cable modem initialization

The procedure for initialization of a cable modem MUST be as shown in Figure 11-1. This figure shows the overall flow between the stages of initialization in a CM. This shows no error paths, and is simply to provide an overview of the process. The more detailed finite state machine representations of the individual sections (including error paths) are shown in the subsequent figures. Timeout values are defined in Annex B.

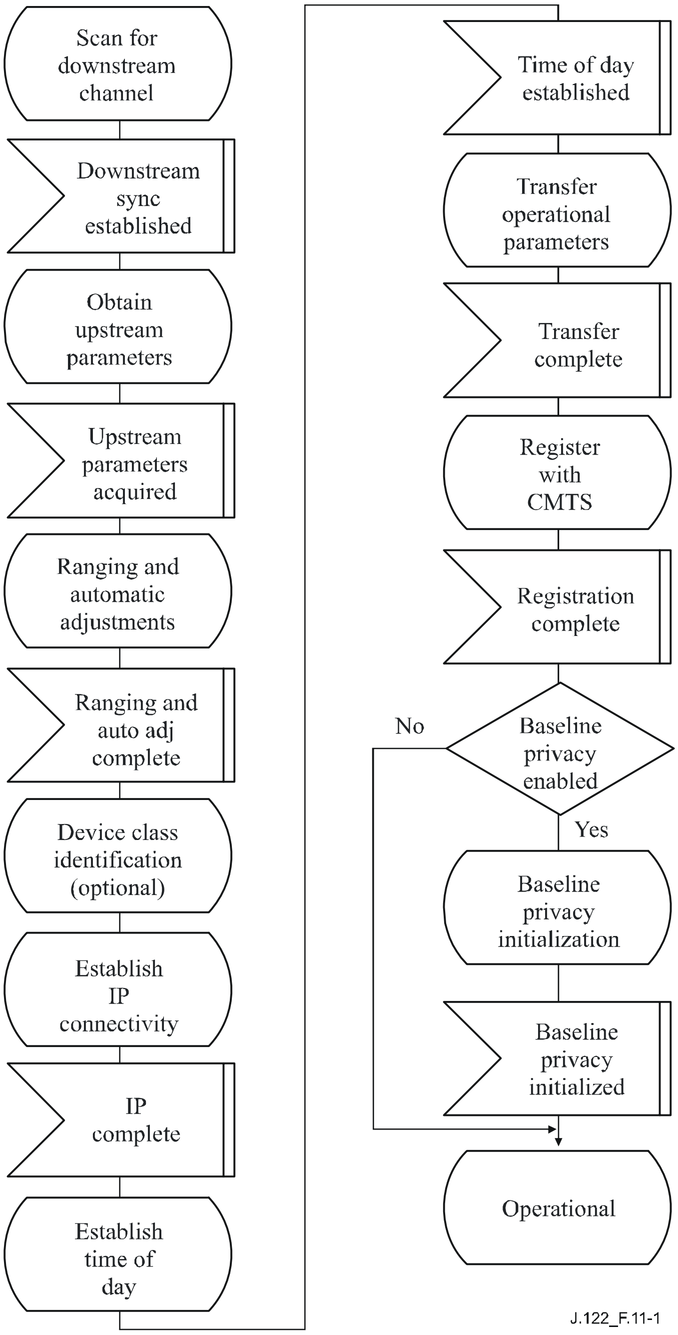

Figure 11-1 – CM initialization overview

The procedure for initializing a cable modem and for a CM to re-initialize its MAC can be divided into the following phases:

- scanning and synchronization to downstream;
- obtaining upstream parameters;
- ranging and automatic adjustments;
- device class identification (optional);
- establishing IP connectivity;
- establishing time of day;
- transferring operational parameters;
- registration;

- baseline privacy initialization, if CM is provisioned to run baseline privacy.

Each CM contains the following information when shipped from the manufacturer:

- A unique IEEE 802 48-bit MAC address that is assigned during the manufacturing process. This is used to identify the modem to the various provisioning servers during initialization.
- Security information as defined in [DOCS8] (e.g., X.509 certificate) used to authenticate the CM to the security server and authenticate the responses from the security and provisioning servers.

The SDL (specification and description language) notation used in the following figures is shown in Figure 11-2 (refer to [ITU-T Z.100]).

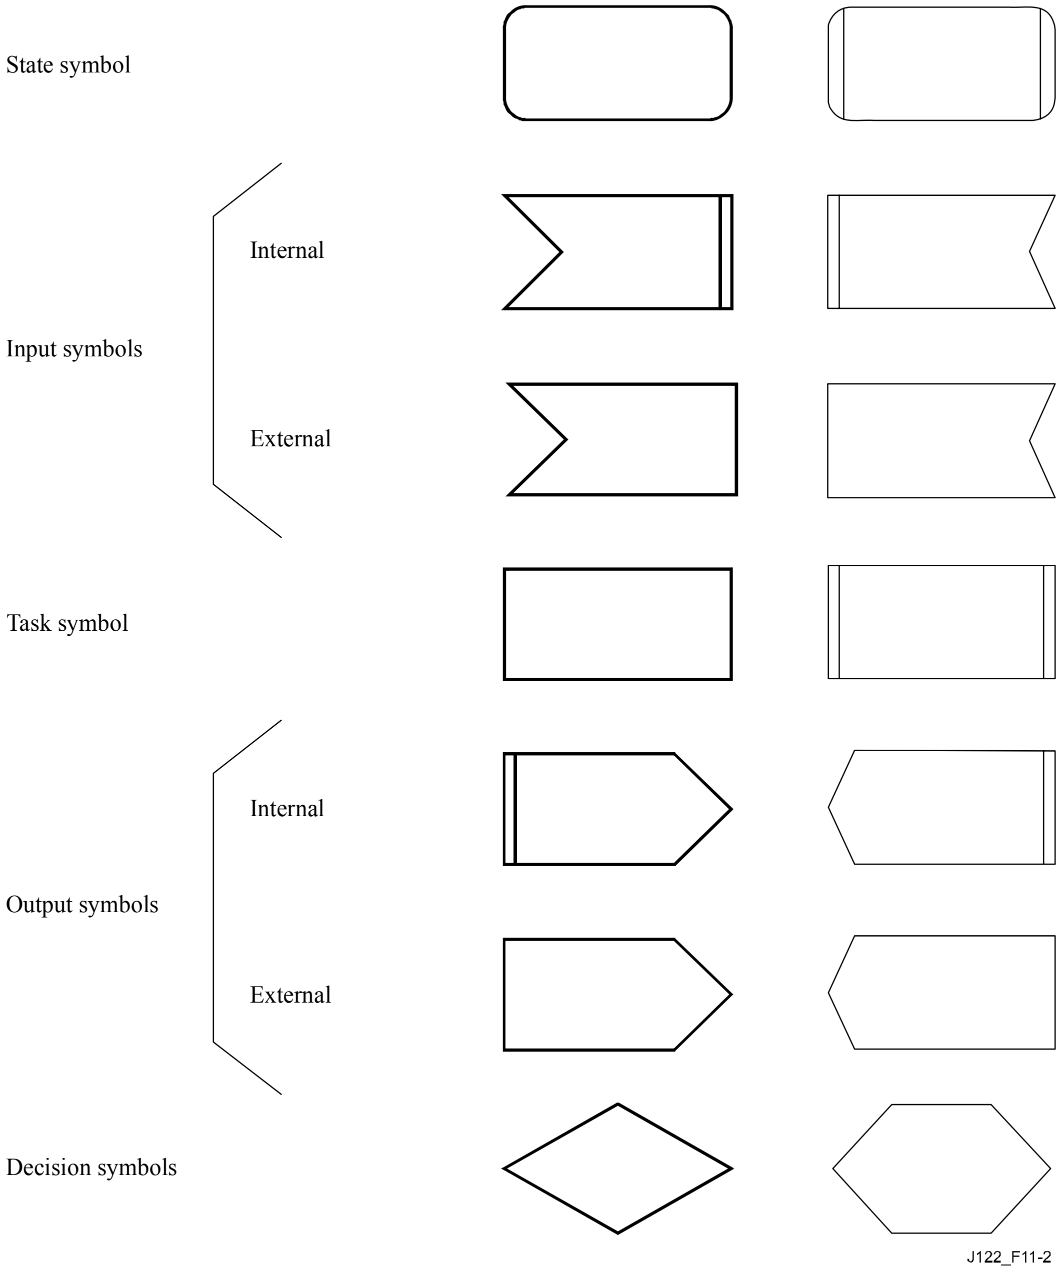

Figure 11-2 – SDL notation

11.2.1 Scanning and synchronization to downstream

On initialization or a "re-initialize MAC" operation, the cable modem MUST acquire a downstream channel. The CM MUST have non-volatile storage in which the last operational parameters are stored and MUST first try to re-acquire this downstream channel. If this fails, it MUST begin to continuously scan the 6-MHz channels of the downstream frequency band of operation until it finds a valid downstream signal.

A downstream signal is considered to be valid when the modem has achieved the following steps:

- synchronization of the QAM symbol timing;
- synchronization of the FEC framing;
- synchronization of the MPEG packetization;
- recognition of SYNC downstream MAC messages.

While scanning, it is desirable to give an indication to the user that the CM is doing so.

In order to support redundant CMTS architectures, when a CM in the operational state detects that the downstream signal is invalid (i.e., does not meet the four criteria above), the CM MUST NOT immediately perform a re-initialize MAC operation. It must instead attempt to re-establish synchronization on the current downstream channel (see clause 11.5). Such re-establishment attempts MUST continue until the operation of periodic ranging as specified in Figure 11-17 calls for a "re-initialize MAC" operation after the expiration of timeout T4 or 16 expirations of timeout T3. Figure 11-17 shows the procedure that MUST be followed by a cable modem during standard operation.

An interruption of downstream signal occurs when the following conditions are met:

- The interruption occurs on a downstream that is valid (per clauses 6.3.5, 6.3.6, 6.3.7, 6.3.8, 6.3.9 and 6.3.10) before and after the loss.
- The interruption is defined as an instantaneous loss of signal and, after a predetermined delay, an instantaneous return to the original signal fidelity.
- The restored downstream signal is the original signal transmitted from the original source.
- The carrier frequency, physical plant and path delays remain the same before and after the interruption.
- There are no changes in any downstream signalling parameter, including the modulation and the M/N ratio, from before to after the interruption.

When a CM in the operational state in S-CDMA mode receives an interruption of downstream signal less than or equal to 5 ms, the CM MUST recover from the outage such that its fixed timing error is not greater than 0.02 of the nominal modulation interval (in addition to the allowed jitter defined in clause 6.2.21.8). When a CM in the operational state in TDMA mode receives an interruption of downstream signal less than or equal to 5 ms, the CM MUST recover from the outage such that the first upstream transmission after the CM resumes normal operation is performed within an accuracy of ±250 nanoseconds plus ±0.5 symbols (refer to clause 6.2.19.1). In either mode, the CM MUST continue with normal operation within 2 s from the end of the interruption. The CM is not required to continue normal operation if it receives a second interruption of downstream signal prior to the first receipt of a RNG-RSP with status "success".

When a CM in the operational state receives an interruption of downstream signal greater than 5 ms but less than the lost sync interval (see Annex B), the CM MAY continue with normal operation as long as it recovers within 2 s:

- with a fixed timing error not greater than 0.02 of the nominal modulation interval (in addition to the allowed jitter defined in clause 6.2.21.8), if the CM is in S-CDMA mode;
- within the accuracy of clause 6.2.19.1, if the CM is in TDMA mode.

If the CM cannot recover according to the preceding recovery time, timing and jitter specifications, the CM MUST re-acquire upstream timing to an accuracy of at least 1 µs and be ready to respond to a ranging opportunity within 2 s, and MUST receive a RNG-RSP message with status "success" before resuming its upstream transmission. For the ranging process, the CM MAY use broadcast or unicast initial maintenance intervals, or station maintenance intervals, although a CM in S-CDMA mode MUST NOT use spreader-on station maintenance. For the ranging process, the CM MUST use its primary SID in the INIT-RNG-REQ or RNG-REQ message and MUST use its known timing offset. After completing this ranging process by receiving RNG-RSP with the ranging status set to "success", the CM MUST return to normal operation. A CMTS MUST process INIT-RNG-REQ messages with a primary SID from any CM that is in normal operation. If the primary SID used by the CM in INIT-RNG-REQ is no longer valid, the CMTS SHOULD send a RNG-RSP message to the CM with the ranging status set to "abort".

In all cases, after the first successful ranging opportunity subsequent to the interruption, the CM MUST meet the timing requirements specified in clauses 6.2.19.1 and 6.2.21.8.

11.2.2 Obtain upstream parameters

Refer to Figure 11-3. After synchronization, the CM MUST wait for an upstream channel descriptor message (UCD) from the CMTS in order to retrieve a set of transmission parameters for a possible upstream channel. These messages are transmitted periodically from the CMTS for all available upstream channels and are addressed to the MAC broadcast address. The CM MUST determine whether it can use the upstream channel from the channel description parameters.

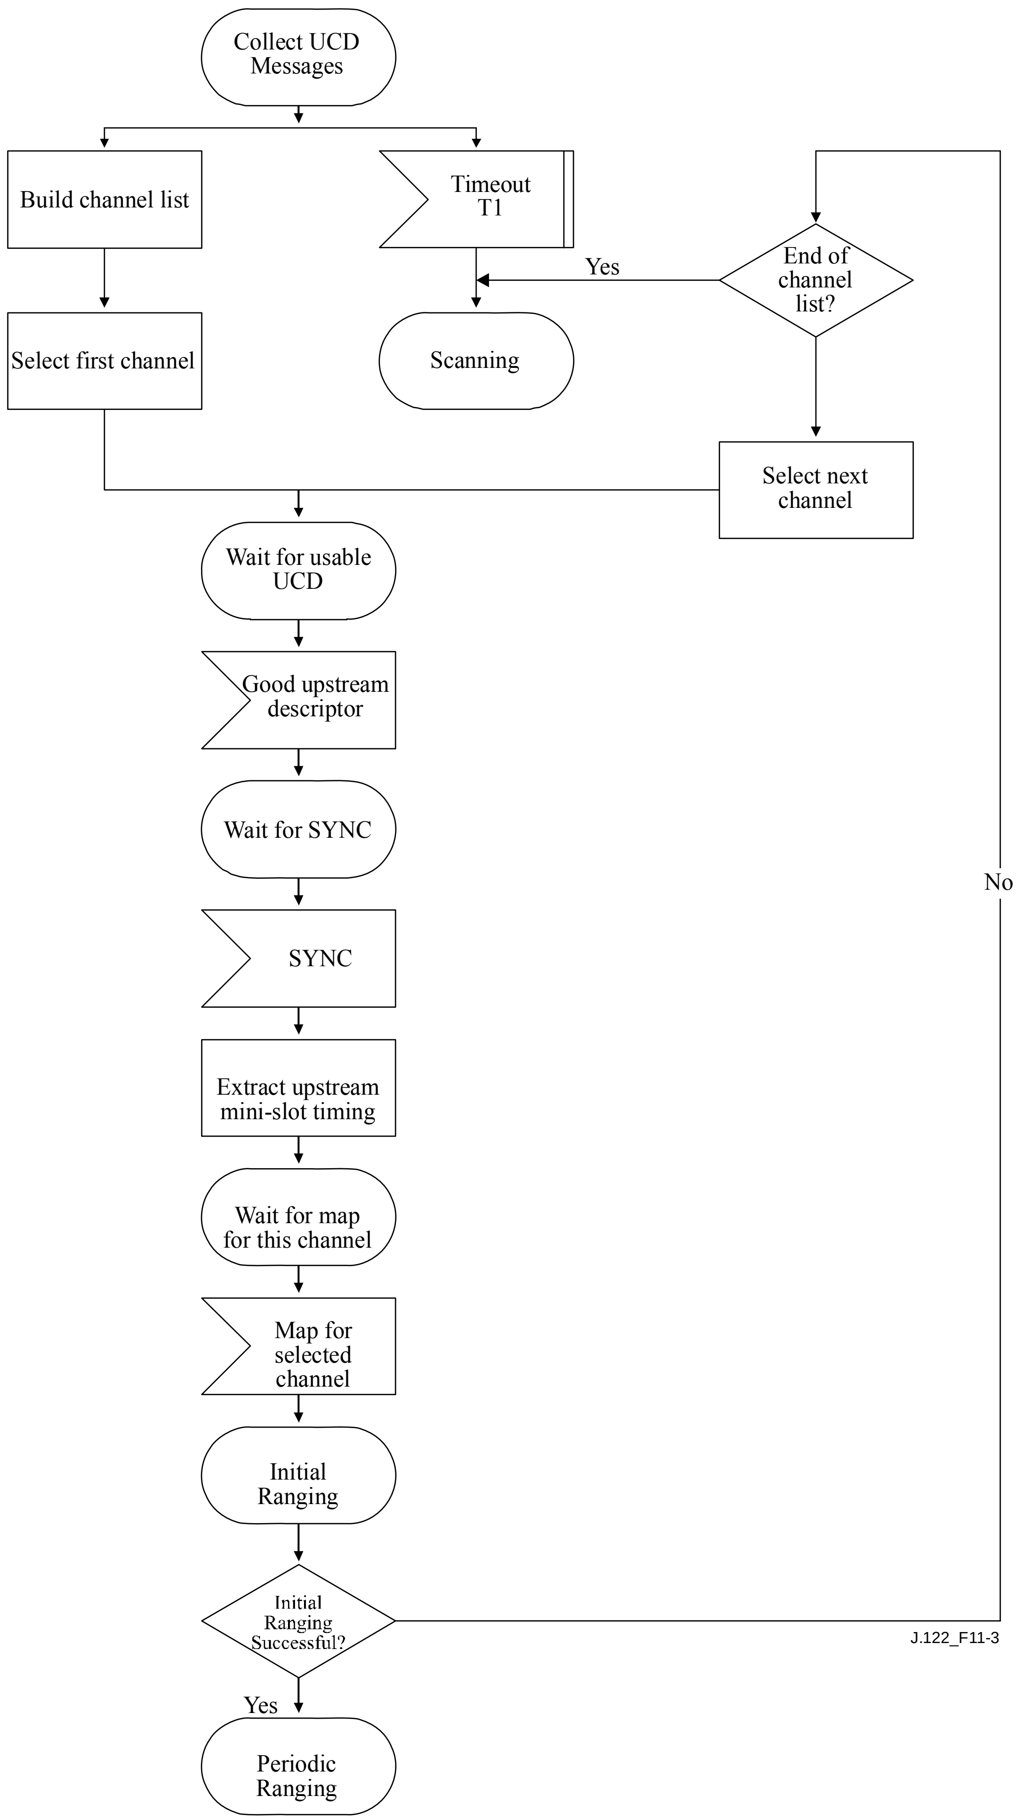

Figure 11-3 – Obtaining upstream parameters

The CM MUST collect all UCDs with different channel ID fields to build a set of usable channel IDs. If no channel can be found after a suitable timeout period, the CM MUST continue scanning to find another downstream channel.

The CM MUST determine whether it can use the upstream channel from the channel description parameters. If the channel is not suitable, the CM MUST try other channels until it finds a usable channel.

Before attempting initial ranging on an upstream, the CM categorizes the available upstreams into the following types based on the UCD for each channel:

1) With a UCD (MAC management message type 2) offering DOCS 1.x burst descriptors only.
2) With a UCD (MAC management message type 2) offering both DOCS 2.0 TDMA and DOCS 1.x burst descriptors.
3) With a UCD (MAC management message type 29) for a DOCS 2.0-only upstream.

The CM MUST have non-volatile storage in which channel ID of the last upstream on which the CM successfully completed registration is stored. If multiple upstreams are available, the CM MUST attempt to use the one that matches this stored channel ID. If none of the available upstreams match that stored ID, or if the CM is unable to successfully complete initial ranging on the matching channel, then the CM MUST preferentially select upstream channels in the following order. Type 3 channels are first, followed by type 2 channels, and type 1 channels are last. The CM MUST NOT begin initial ranging on a type 1 or type 2 upstream until it has allowed sufficient time, at least the UCD interval (refer to Annex B), to determine if a type 3 upstream is available. If initial ranging fails on a type 3 upstream, the CM MUST ensure that it has allowed sufficient time to detect any other type 3 upstreams that are available before moving on to a type 2 or type 1 upstream. Of course, once the CM has waited enough time to ensure that it knows about any available type 3 upstreams, it will also know about any available type 2 upstreams and it MUST try them in preference to any type 1 upstreams.

If the channel is suitable, the CM MUST extract the parameters for this upstream from the UCD. If the UCD message contains a type 19 TLV, the CM MUST (except as noted below) perform a bitwise AND of its ranging class ID with the TLV 19 Value. If the result of the bitwise AND is zero, the CM MUST consider the channel unusable and try other channels until it finds a usable channel.

If the CM is obtaining upstream parameters prior to performing initial ranging with an assigned SID or in preparation for ranging in response to a DCC-REQ, it MUST ignore TLV-19.

If the CM is not responding to a DCC-REQ and is going to be ranging with the initialization SID and if the UCD contains a type 18 TLV, the CM performs a bitwise AND of its ranging class ID with the TLV 18 value. If the result of the bitwise AND is equal to the CM's ranging class ID, the CM MUST inhibit initial ranging and start the T17 timer. If the UCD change count in the UCD message for the channel the CM is using, is incremented while the T17 timer is active, the CM will re-inspect the TLV-18 value and re-start the T17 timer if necessary. If the T17 timer expires or the TLV-18 value is updated to permit ranging for the CMs ranging class, the CM will resume normal ranging. If the CM should undergo a lost SYNC event while waiting for T17, it MUST re-initialize the MAC. A CM that is performing initial ranging with an assigned SID MUST ignore TLV-18.

After having transmitted at least one initial maintenance RNG-REQ message, the CM MUST ignore TLV-18 or TLV-19 values in any new UCD message for the channel even if the new UCD contains a ranging required TLV.

The CM MUST wait for the next SYNC message and extract the upstream mini-slot timestamp from this message[26]. The CM then MUST wait for a bandwidth allocation map for the selected channel. It may begin transmitting upstream in accordance with the MAC operation and the bandwidth allocation mechanism.

The CM MUST perform initial ranging at least once, in accordance with Figure 11-6. If initial ranging is not successful, the next channel ID is selected, and the procedure restarted from UCD extraction. When there are no more channel IDs to try, the CM MUST continue scanning to find another downstream channel.

11.2.3 Message flows during scanning and upstream parameter acquisition

The CMTS MUST generate SYNC and UCD messages on the downstream at periodic intervals within the ranges defined in Annex B. These messages are addressed to all CMs. Refer to Figure 11-4.

[26] Alternatively, since the SYNC message applies to all upstream channels, the CM may have already acquired a time reference from previous SYNC messages. If so, it need not wait for a new SYNC.

CMTS		CM
Clock time to send SYNC	——— SYNC ——→	|
Clock time to send UCD	——— UCD ——→	|
		|
Clock time to send SYNC	——— SYNC ——→	|
		| Example of a UCD cycle prior to
		| CM power-on
		|
Clock time to send SYNC	——— SYNC ——→	|
		|
		|
		|
Clock time to send SYNC	——— SYNC ——→	|
		|
		|
Clock time to send SYNC	——— SYNC ——→	
Clock time to send UCD	——— UCD ——→	
Clock time to send SYNC	——— SYNC ——→	
		Power-on sequence complete
Clock time to send SYNC	——— SYNC ——→	
		Establish PHY synchronization and wait for UCD
Clock time to send SYNC	——— SYNC ——→	
Clock time to send SYNC	——— SYNC ——→	
Clock time to send UCD	——— UCD ——→	
		Obtain parameters for this upstream channel to use for initialization
Clock time to send SYNC	——— SYNC ——→	
		Extract slot info for upstream and wait for transmit opportunity to perform ranging
Clock time to send SYNC	——— SYNC ——→	
Clock time to send MAP	——— MAP ——→	
		Start ranging process

Figure 11-4 – Message flows during scanning and upstream parameter acquisition

11.2.4 Ranging and automatic adjustments

The ranging and adjustment process is fully defined in clause 8 and in the following clauses. The message sequence chart and the finite state machines on the following pages define the ranging and adjustment process which MUST be followed by compliant CMs and CMTSs. Refer to Figures 11-5 through 11-8.

NOTE – MAPs are transmitted as described in clause 8.

CMTS		CM
[Time to send the initial maintenance opportunity]		
Send MAP containing initial maintenance information element with a broadcast/multicast service ID	—— MAP ——→	
	←—— RNG-REQ or INIT-RNG-REQ ——	Transmit ranging packet in contention mode with service ID parameter = 0
[Receive recognizable ranging packet]		
Allocate temporary service ID		
Send ranging response	—— RNG-RSP ——→	
Add temporary service ID to poll list		Store temporary service ID and adjust other parameters
[Time to send the next MAP]		
Send MAP with station maintenance information element or unicast initial maintenance element to modem using temporary SID	—— MAP ——→	Recognize own temporary service ID in map
	←—— RNG-REQ ——	Reply to station maintenance opportunity poll or unicast initial maintenance opportunity poll
Send ranging response	—— RNG-RSP ——→	
		Adjust local parameters
[Time to send an initial maintenance opportunity]		
Send MAP containing initial maintenance information element with a broadcast/multicast service ID		
Send periodic transmit opportunity to broadcast address	—— MAP ——→	

NOTE – The CMTS MUST allow the CM sufficient time to have processed the previous RNG-RSP (i.e., to modify the transmitter parameters) before sending the CM a specific ranging opportunity. This is defined as CM ranging response time in Annex B.

Figure 11-5 – Ranging and automatic adjustments procedure

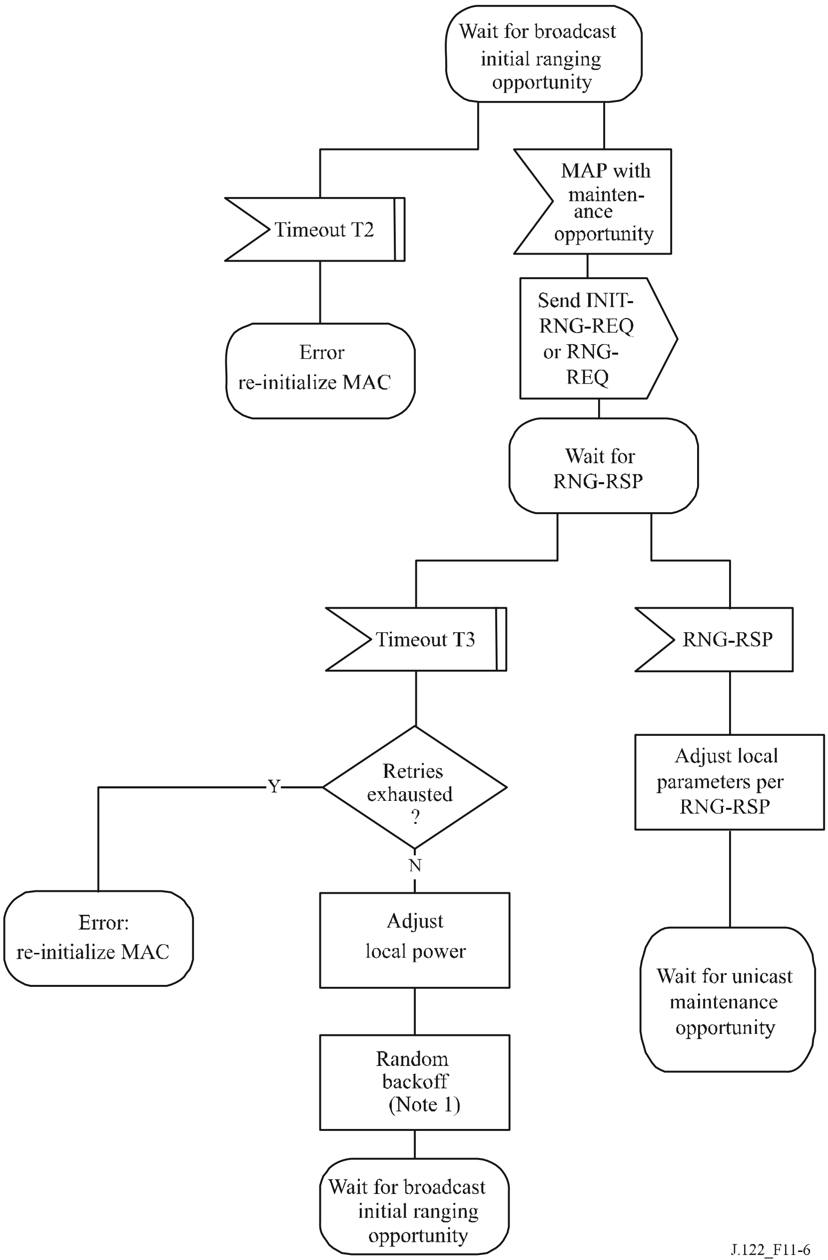

NOTE 1 – Timeout T3 may occur because the RNG-REQs from multiple modems collided. To avoid these modems repeating the loop in lockstep, a random backoff is required. This is a backoff over the ranging window specified in the MAP. T3 timeouts can also occur during multi-channel operation. On a system with multiple upstream channels, the CM MUST attempt initial ranging on every suitable upstream channel before moving to the next available downstream channel.
NOTE 2 – A unicast maintenance opportunity can be either a station maintenance or a unicast initial maintenance opportunity.

Figure 11-6 – Broadcast initial ranging – CM

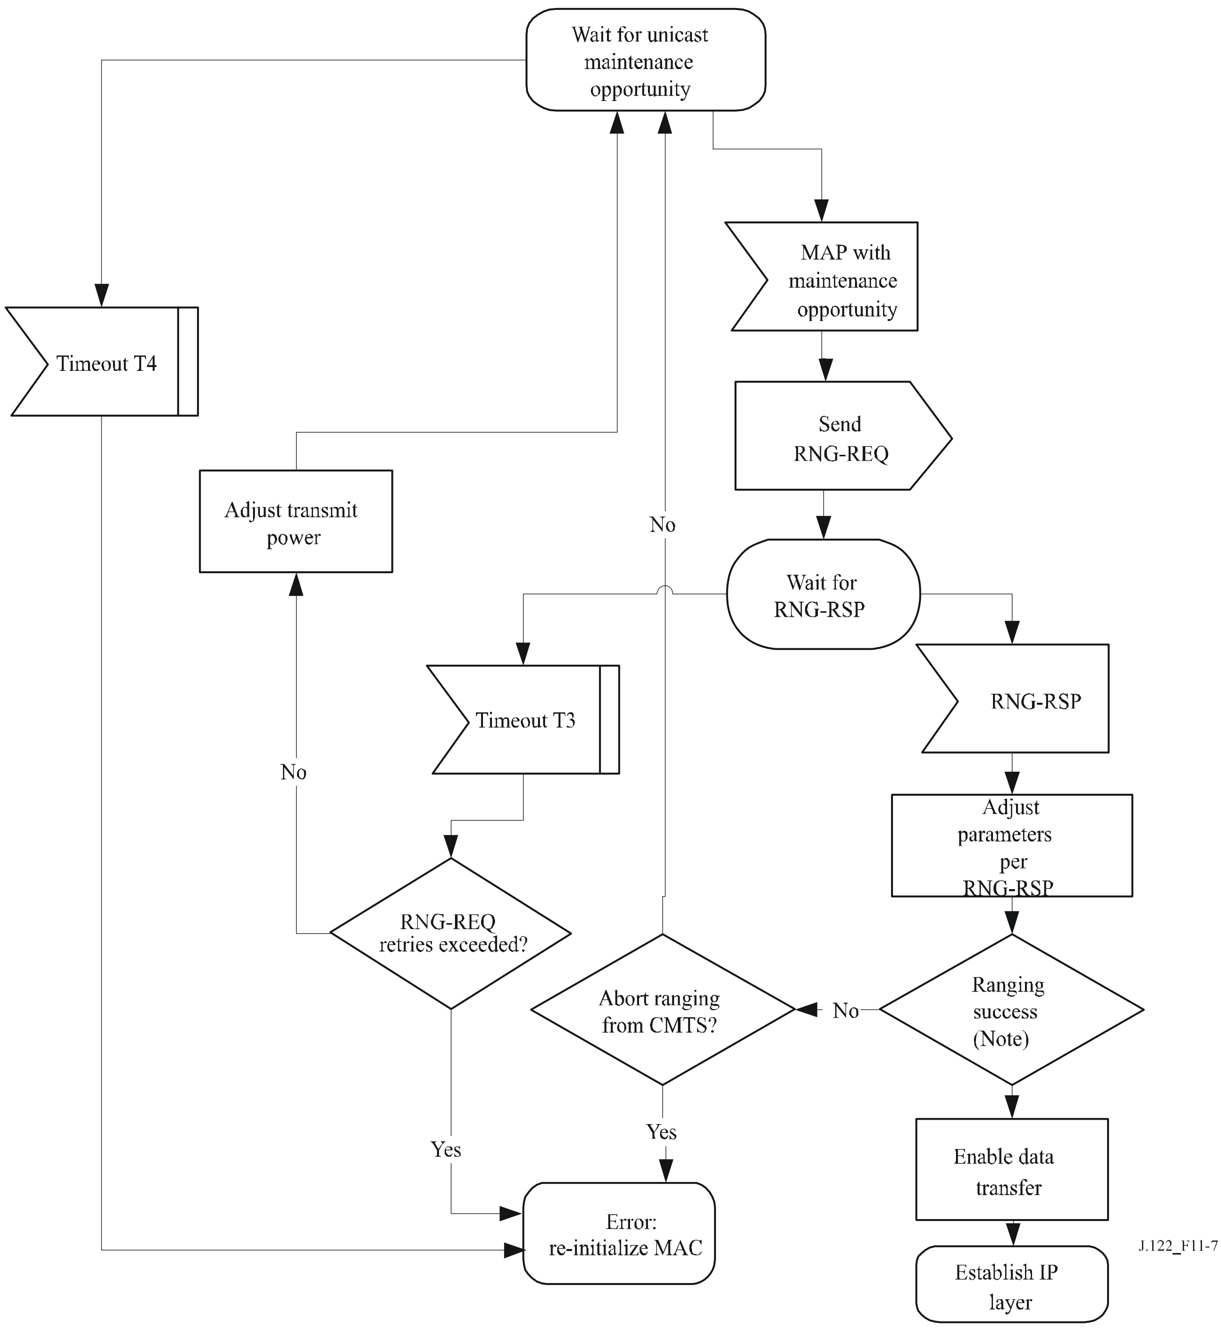

NOTE – Ranging request is within the tolerance of the CMTS.

Figure 11-7 – Unicast initial ranging – CM

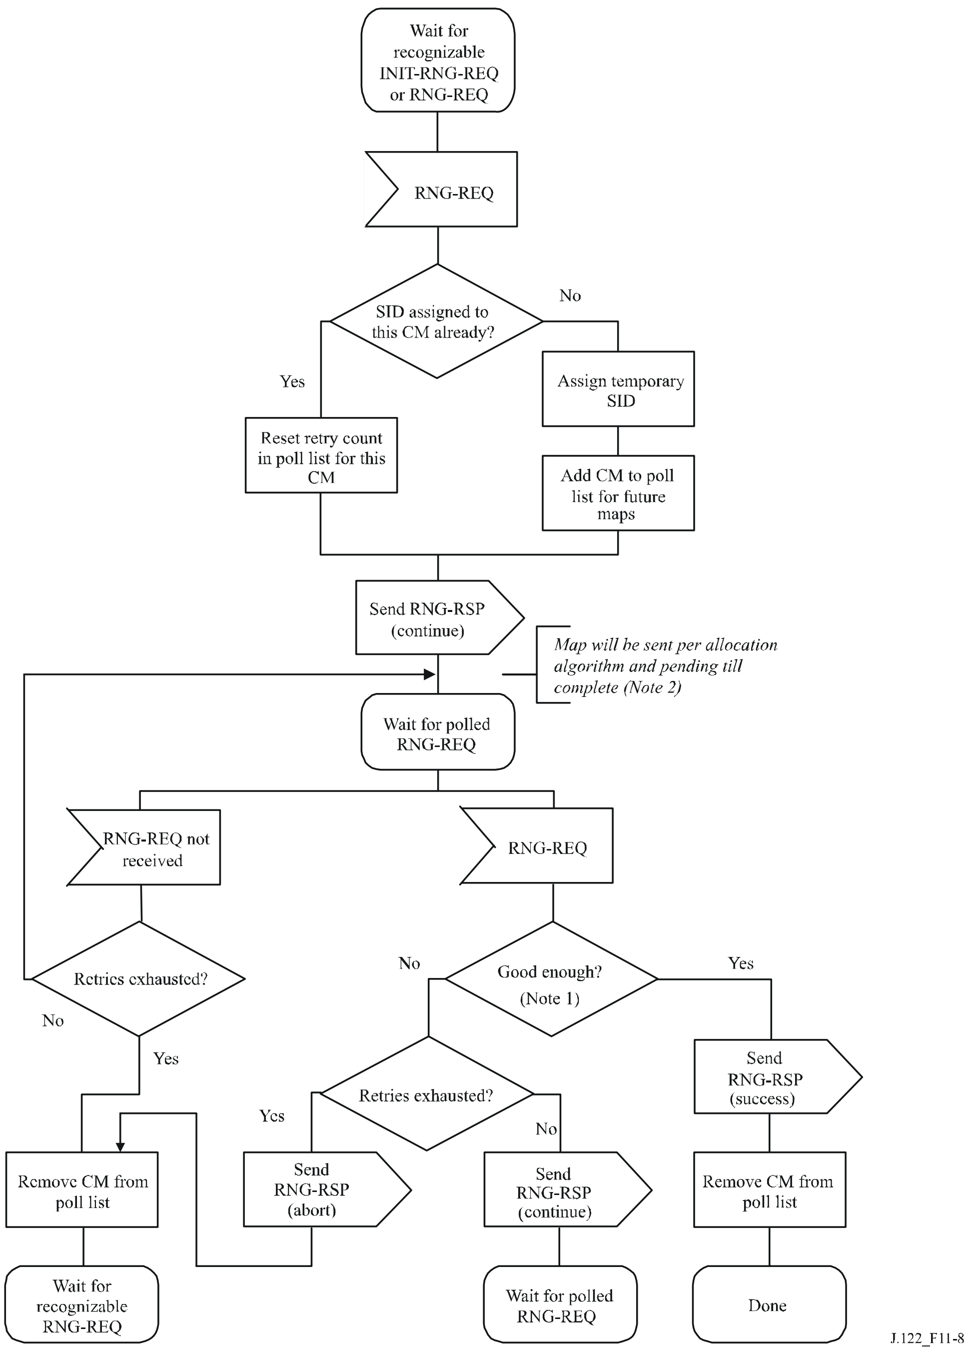

NOTE 1 – "Good enough" means ranging is within the tolerable limits of the CMTS.
NOTE 2 – RNG-REQ pending-till-complete was non-zero, the CMTS SHOULD hold off the station maintenance opportunity accordingly unless needed, for example, to adjust the CMs power level. If opportunities are offered prior to the pending-till-complete expiry, the "good-enough" test that follows receipt of a RNG-RSP MUST NOT judge the CMs transmit equalization until pending-till-complete expires.

Figure 11-8 – Initial ranging – CMTS

11.2.4.1 Ranging parameter adjustment

Adjustment of local parameters (e.g., transmit power) in a CM as a result of the receipt (or non-receipt) of an RNG-RSP is considered to be implementation-dependent with the following restrictions (refer to clause 8.3.6):

- All parameters MUST be within the approved range at all times.
- Power adjustment MUST start from the minimum value unless a valid power is available from non-volatile storage, in which case this MUST be used as a starting point.
- Power adjustment MUST be capable of being reduced or increased by the specified amount in response to RNG-RSP messages.
- If, during initialization, power is increased to the maximum value (without a response from the CMTS) it MUST wrap back to the minimum.
- For multi-channel support, the CM MUST attempt initial ranging on every suitable upstream channel before moving to the next available downstream channel.
- For multi-channel support, the CM MUST use the upstream channel ID of the range response as specified in clause 8.3.6 and in Appendix III.

11.2.5 Device class identification

After ranging is complete and before establishing IP connectivity, the CM MAY identify itself to the CMTS for use in provisioning. Refer to Figure 11-9.

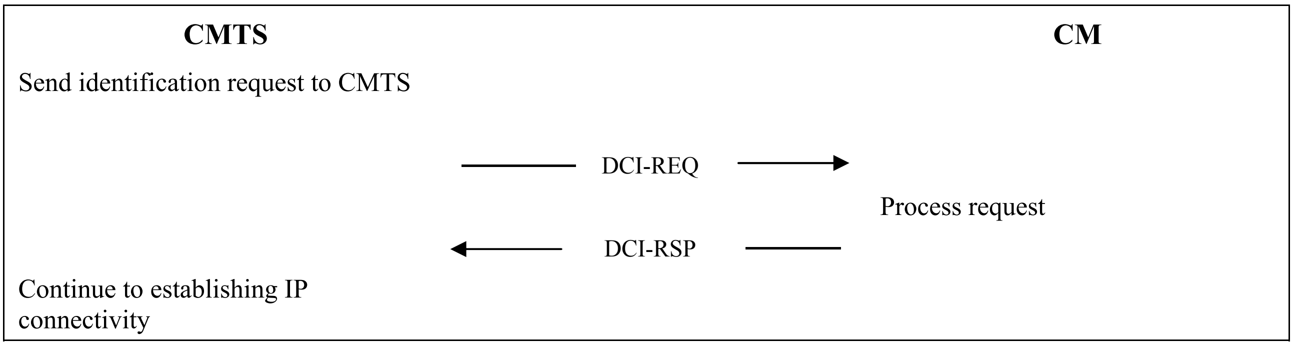

Figure 11-9 – Device class identification

If implemented, the CM MUST use an adaptive timeout for device class identification based on binary exponential backoff, similar to that used for TFTP. Refer to clause 11.2.8 for details.

11.2.6 Establish IP connectivity

At this point, the CM MUST invoke DHCP mechanisms [RFC 2131] in order to obtain an IP address and any other parameters needed to establish IP connectivity (refer to Annex D). The DHCP response MUST contain the name of a file that contains further configuration parameters. Refer to Figure 11-10.

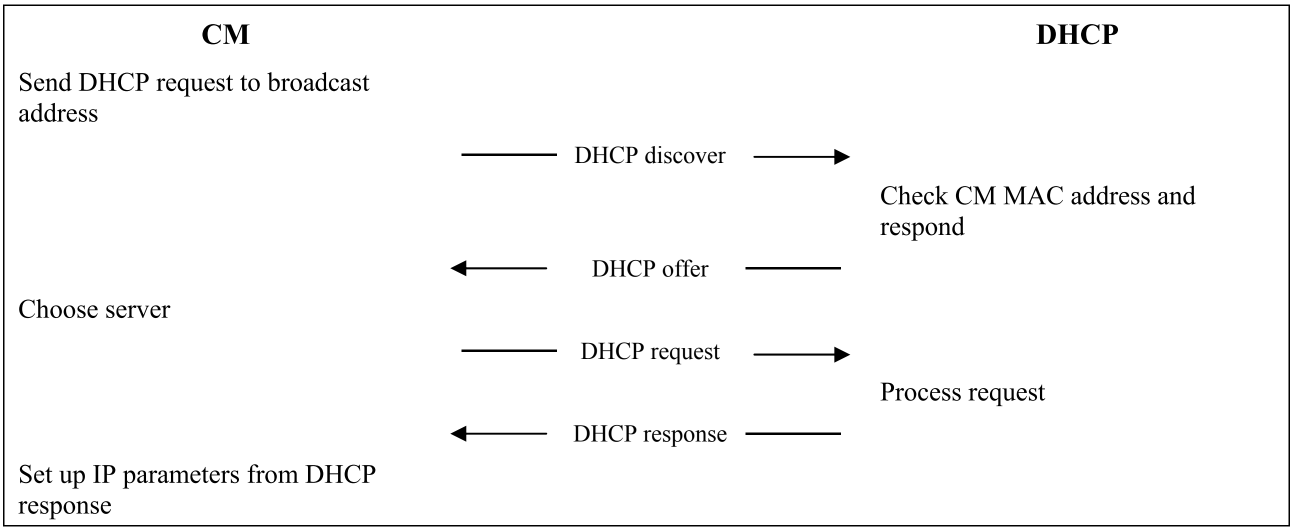

Figure 11-10 – Establishing IP connectivity

As required in [RFC 2131], the CM MUST adopt a retransmission strategy that incorporates a randomized exponential backoff algorithm to determine the delay between retransmissions. The backoff algorithm SHOULD use a backoff start value of 4 seconds, a backoff end value of 64 seconds and randomized value of ±1 second. The CM MUST limit the number of retransmissions. The CM SHOULD limit the number of retransmissions (excluding the first transmission) to a maximum retries value of 5 but MAY implement less than the maximum retries value. The CM SHOULD also implement a different retransmission strategy for the RENEWING and REBINDING states, as recommended in [RFC 2131], which is based on one-half of the remaining lease time. Upon failure of obtaining or rebinding an IP address, the CM MUST attempt to establish communication on another upstream channel (see clause 11.2.2) or resume scanning to find another downstream channel (see clause 11.2.1).

11.2.7 Establish time of day

The CM and CMTS need to have the current date and time. This is required for time-stamping logged events which can be retrieved by the management system. This need not be authenticated and need only be accurate to the nearest second.

The protocol by which the time of day MUST be retrieved is defined in [RFC 868]. Refer to Figure 11-11. The request and response MUST be transferred using UDP. The time retrieved from the server (UTC) MUST be combined with the time offset received from the DHCP response to create the current local time.

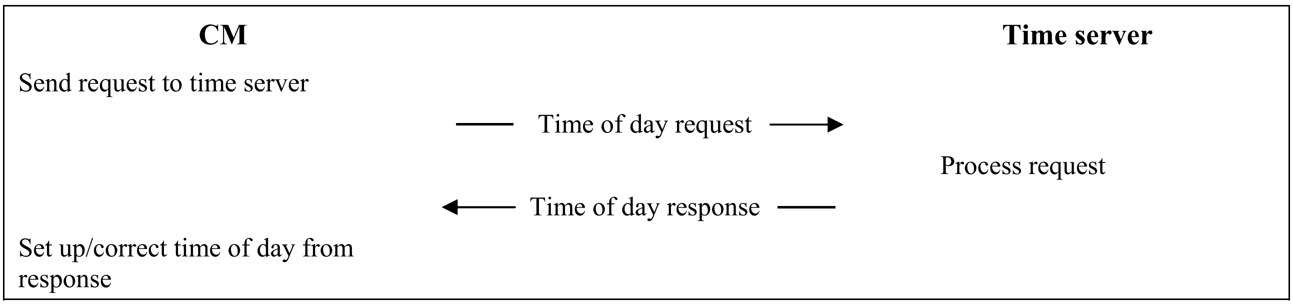

Figure 11-11 – Establishing time of day

The DHCP server may offer a CM multiple time of day server IP addresses to attempt. The CM MUST attempt all time of day servers included in the DHCP offer until local time is established.

Successfully acquiring the time of day is not mandatory for a successful registration, but it is necessary for ongoing operation. If a CM is unable to establish time of day before registration it MUST log the failure, generate an alert to management facilities, then proceed to an operational state and retry periodically.

The specific timeout for time of day Requests is implementation-dependent. However, for each server defined the CM MUST NOT exceed more than 3 time of day requests in any 5-minute period. At minimum, the CM MUST issue at least 1 time of day request per 5-minute period for each server specified until local time is established.

11.2.8 Transfer operational parameters

After DHCP is successful, the modem MUST download the parameter file using TFTP, as shown in Figure 11-12. The TFTP configuration parameter server is specified by the "siaddr" field of the DHCP response. The CM MUST use an adaptive timeout for TFTP based on binary exponential backoff. Refer to [RFC 1123] and [RFC 2349].

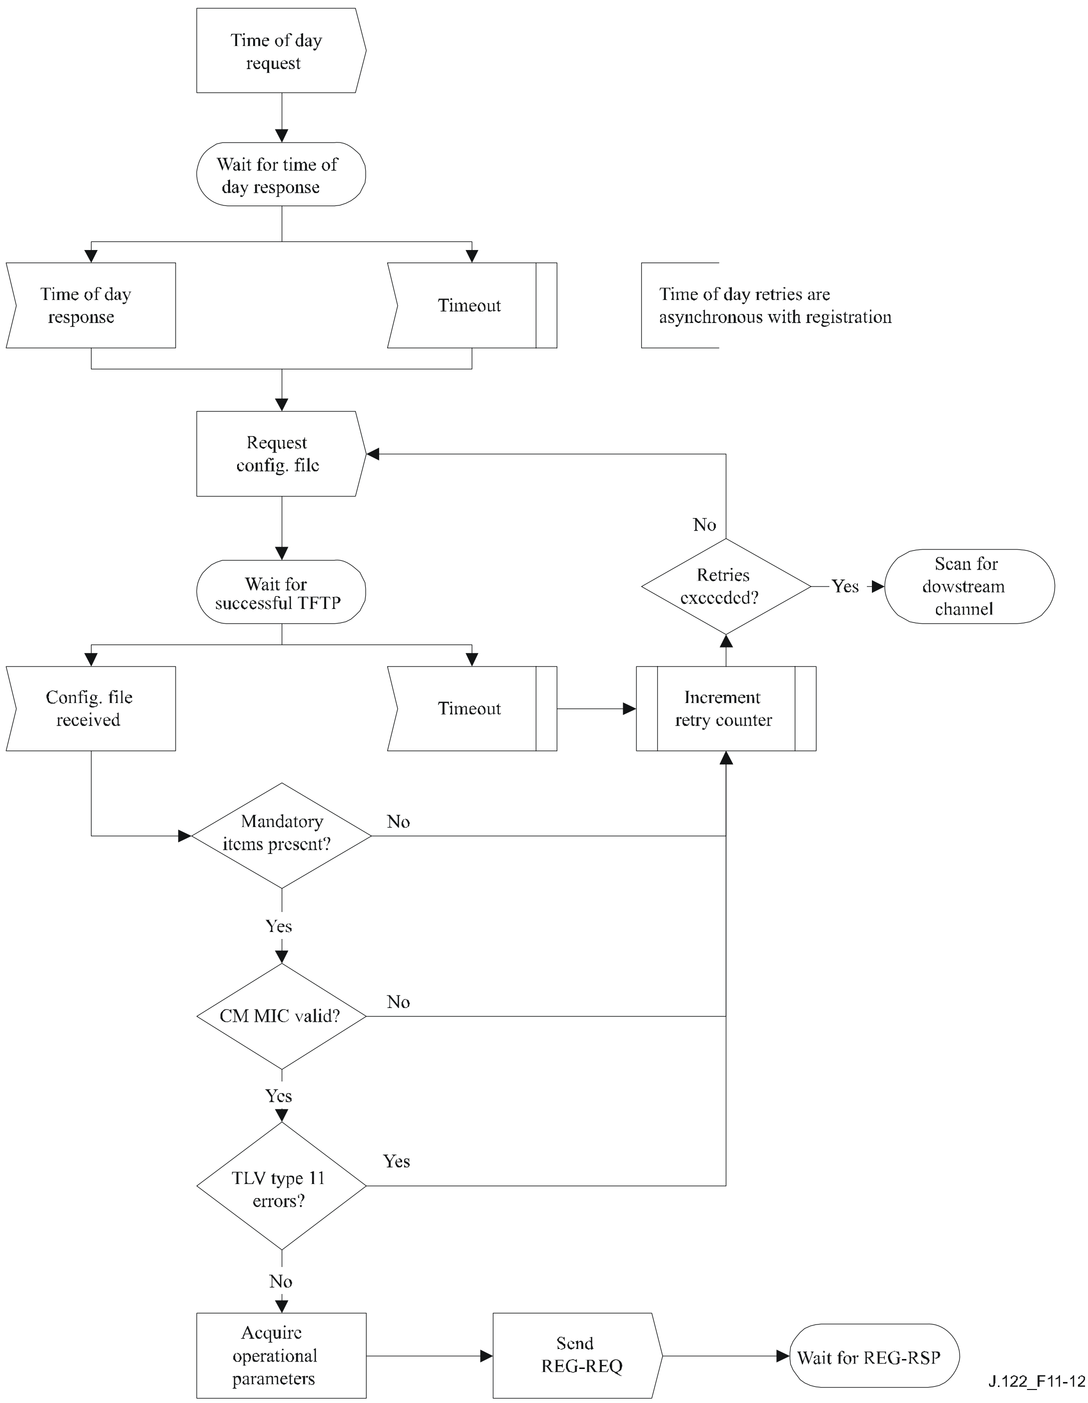

Figure 11-12 – Registration – CM

The parameter fields required in the DHCP response, and the format and content of the configuration file MUST be as defined in Annex D. Note that these fields are the minimum required for interoperability.

If a modem downloads a configuration file containing an upstream channel ID configuration setting (clause C.1.1.2) different from what the modem is currently using, the modem MUST NOT send a registration request message to the CMTS. Likewise, if a modem downloads a configuration file containing a single downstream channel frequency (clause C.1.1.21.1.2) and/or downstream frequency range (clause C.1.1.21.2) that does not include the downstream frequency the modem is currently using, or a downstream frequency configuration setting (clause C.1.1.1) different from what the modem is currently using and no downstream channel list, the modem MUST NOT send a registration request message to the CMTS. In either case, the modem MUST redo initial ranging using the configured upstream channel and/or downstream frequency per clause 8.3.6.3.

If a modem downloads a configuration file containing the enable 2.0 mode TLV set to disable (see clause C.1.1.19), then it MUST NOT operate in 2.0 mode until it registers again and downloads a configuration file which does not have this TLV set to disable. This is true no matter what type of upstream the CM is using at the time it is attempting registration. If it is using a type 3 channel (as described in clause 11.2.2), it MUST NOT send a registration request message to the CMTS. The modem MUST redo initial ranging using a type 1 or type 2 channel. If no such upstream channel is available (or if the modem is unable to range successfully on one) the modem MUST scan for a new downstream, at which point 2.0 mode is no longer disabled. If the modem downloads a configuration file that does not disable 2.0 mode, then, regardless of what kind of upstream it is using at the time of registration, it will continue to operate with 2.0 mode enabled until it is no longer registered. This means that if a modem registers on a type 1 channel with a configuration file that enables 2.0 mode operation, and then ends up on a type 2 or type 3 upstream without going through re-registration, the CM would immediately start operating in 2.0 mode.

11.2.9 Registration

A CM MUST be authorized to forward traffic into the network once it is initialized and configured. The CM is authorized to forward traffic into the network via registration. To register with a CMTS, the CM MUST forward its configured class of service and any other operational parameters in the configuration file (refer to clause 8.3.7) to the CMTS as part of a registration request. The CM MUST perform the following operations before registration (refer to Figure 11-12):

- Check configuration file mandatory items (refer to clause D.2.2). The CM MUST reject a configuration file if it lacks any mandatory items.
- Calculate a MIC per clause D.2.3.1 and compare it to the CM MIC included in the configuration file. If the MIC is invalid, the CM MUST reject the configuration file.
- If the configuration file contains TLV-11 encoding, the CM MUST follow the configuration process defined in clause 3.4 of [DOCS5]. The CM MUST reject a configuration file in the case of TLV-11 processing failure.

Figure 11-12 shows the procedure that MUST be followed by the CM.

The configuration parameters downloaded to the CM MUST include a network access control object (see clause C.1.1.3). If this is set to "no forwarding," the CM MUST NOT forward data from attached CPE to the network, yet the CM MUST respond to network management requests. This allows the CM to be configured in a mode in which it is manageable but will not forward data.

Once the CM has sent a registration request to the CMTS, it MUST wait for a registration response to authorize it to forward traffic to the network. Figure 11-13 shows the waiting procedure that MUST be followed by the CM.

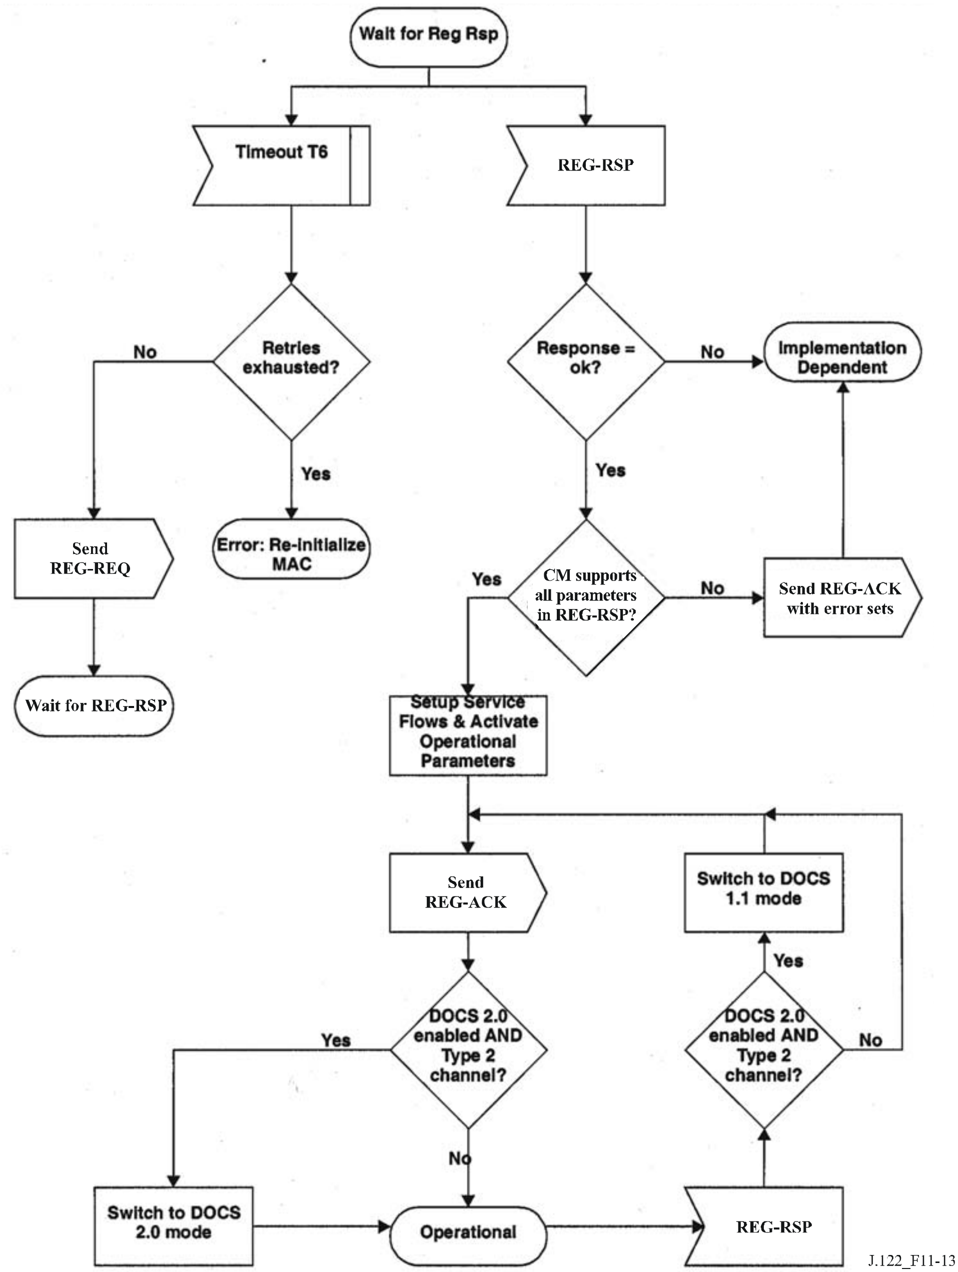

Figure 11-13 – Wait for registration response – CM[27]

NOTE – In Figures 11-12 and 11-13, a type 2 upstream supports both DOCS 1.1 and DOCS 2.0 TDMA bursts (see 11.2.2). On such an upstream, it is necessary that the CMTS knows whether the CM has calculated a request to transmit data based on the DOCS 2.0 TDMA data IUCs (9, 10) or on the DOCS 1.x IUCs (5, 6). On these upstreams, the CM calculates all requests up to and including the request to transmit the REG-ACK based on the DOCS 1.x IUCs. If the CM is enabled to operate in 2.0 mode (see clause C.1.1.19), all requests to transmit data on such an upstream, subsequent to the transmission of the REG-ACK, are calculated based on the DOCS 2.0 TDMA data IUCs. If the CMTS indicates that it did not receive the REG-ACK by retransmitting the REG-RSP, a CM using such an upstream MUST revert to the DOCS 1.x data IUCs to request bandwidth for retransmitting the REG-ACK.

[27] Refer to clause 11.2.2 for channel type definitions.

The CMTS MUST perform the following operations to confirm the CM authorization (refer to Figures 11-14 and 11-15):

- Calculate a MIC per clause D.3.1 and compare it to the CMTS MIC included in the registration request. If the MIC is invalid, the CMTS MUST respond with an authentication failure.
- If present, check the TFTP server timestamp field. If the CMTS detects that the time is different from its local time by more than CM configuration processing time (refer to Annex B), the CMTS MUST indicate authentication failure in the REG-RSP. The CMTS SHOULD also make a log entry stating the CM MAC address from the message.
- If present, check the TFTP server provisioned modem address field. If the provisioned modem address does not match the requesting modem's actual address, the CMTS MUST indicate authentication failure in the REG-RSP. The CMTS SHOULD also make a log entry stating the CM MAC address from the message.
- If the registration request contains DOCS 1.0 class-of-service encodings, verify the availability of the class(es) of service requested. If unable to provide the class(es) of service, the CMTS MUST respond with a class-of-service failure and the appropriate service not available response code(s) (refer to clause C.1.3.4).
- If the registration request contains service flow encodings, verify the availability of the quality of service requested in the provisioned service flow(s). If unable to provide the service flow(s), the CMTS MUST respond with either a reject-temporary or a reject-permanent (see clause C.4) and the appropriate service flow response(s).
- If the registration request contains DOCS 1.0 class-of-service encodings and service flow encodings, the CMTS MUST respond with a class-of-service failure and a service not available response code set to 'reject-permanent' for all DOCS 1.0 classes and service flows requested.
- Verify the availability of any modem capabilities requested. If unable or unwilling to provide the modem capability requested, the CMTS MUST turn that modem capability 'off' (refer to clause 8.3.8.1.1).
- Assign a service flow ID for each class of service supported.
- Reply to the modem in a registration response.
- If the registration request contains service flow encoding and the REG-RSP is sent with a confirmation code of ok (0), the CMTS MUST wait for a registration acknowledgment as shown in Figure 11-14. If the registration request contains DOCS 1.0 class-of-service encodings, the CMTS MUST NOT wait for a registration acknowledgment.
- In a channel that supports both DOCS 1.x and DOCS 2.0 TDMA burst types (see clause 11.2), the CMTS MUST change the operational state of the DOCS 2.0 enabled CM (see clause C.1.1.19) to DOCS 2.0 TDMA after the registration acknowledgement message is received from the CM.
- If timer T9 expires, the CMTS MUST both de-assign the temporary SID from that CM and make some provision for aging out that SID.

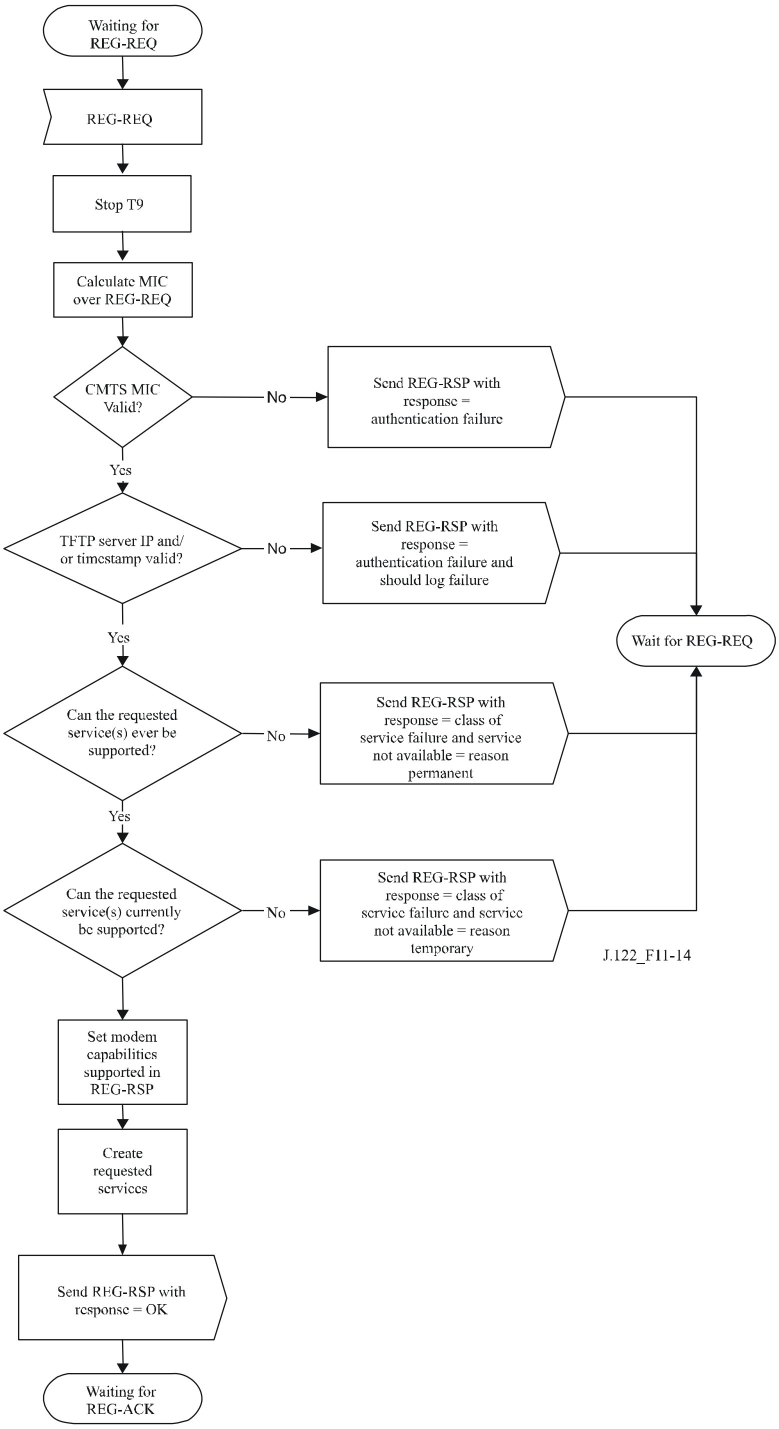

Figure 11-14 – Registration – CMTS

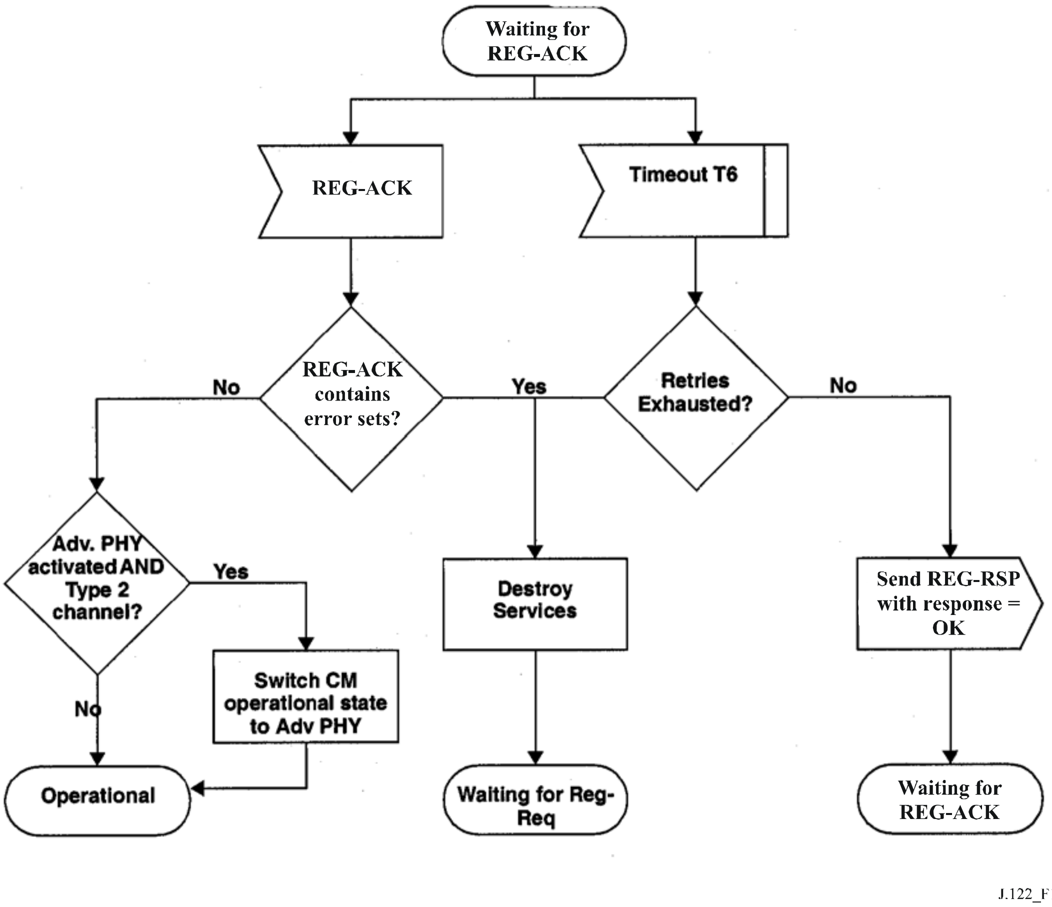

Figure 11-15 – Registration acknowledgment – CMTS

11.2.10 Baseline privacy initialization

Following registration, if the CM is provisioned to run baseline privacy, the CM MUST initialize baseline privacy operations as described in [DOCS8]. A CM is provisioned to run baseline privacy if the privacy enable TLV (see clause C.1.1.16) in the DOCS 1.1 style configuration file is explicitly/implicitly set to enable or the baseline privacy configuration setting (see clause C.3.2) is contained in the DOCS 1.0 style configuration file as specified in clauses A.1.1 and C.2 of the BPI+ Recommendation [DOCS8]. Note that the secure software download is required regardless of whether the CM is provisioned to run baseline privacy or not as specified in Appendix D of the BPI+ Recommendation [DOCS8].

When assigning provisioned SFIDs on receiving a registration request, the CMTS may re-use the temporary SID, assigning it to one of the service flows requested. If so, it MUST continue to allow initialization messages on that SID, since the registration response could be lost in transit. If the CMTS assigns all-new SIDs for class-of-service provisioning, it MUST age out the temporary SID. The aging-out MUST allow sufficient time to complete the registration process in case the registration response is lost in transit.

11.2.11 Service IDs during CM initialization

After completion of the registration process (see clause 11.2.9), the CM will have been assigned service flow IDs (SFIDs) to match its provisioning. However, the CM must complete a number of protocol transactions prior to that time (e.g., ranging, DHCP, etc.), and requires a temporary service ID in order to complete those steps.

On reception of an initial ranging request, the CMTS MUST allocate a temporary SID and assign it to the CM for initialization use. The CMTS MAY monitor use of this SID and restrict traffic to that needed for initialization. It MUST inform the CM of this assignment in the ranging response.

The temporary SID assigned by the CMTS MUST be in the unicast SID space (see clause 9.1.2). DOCSIS 2.0 CMs MUST support the capability to transmit unicast traffic on the expanded SID space (see clause C.1.3.1.3). If a CM supports the above capability the CMTS MAY assign SID numbers from the expanded unicast SID space in the registration response.

On receiving a ranging response addressed to it, the CM MUST use the assigned temporary SID for further initialization transmission requests until the registration response is received.

On receiving a ranging response instruction to move to a new downstream frequency and/or upstream channel ID, the CM MUST consider any previously assigned temporary SID to be de-assigned, and must obtain a new temporary SID via initial ranging.

It is possible that the ranging response may be lost after transmission by the CMTS. The CM MUST recover by timing out and re-issuing its initial ranging request. Since the CM is uniquely identified by the source MAC address in the ranging request, the CMTS MAY immediately reuse the temporary SID previously assigned. If the CMTS assigns a new temporary SID, it MUST make some provision for aging out the old SID that went unused (see clause 8.3.8).

When assigning provisioned SFIDs on receiving a registration request, the CMTS may reuse the temporary SID, assigning it to one of the service flows requested. If so, it MUST continue to allow initialization messages on that SID, since the registration response could be lost in transit. If the CMTS assigns all-new SIDs for class-of-service provisioning, it MUST age out the temporary SID. The aging-out MUST allow sufficient time to complete the registration process in case the registration response is lost in transit.

11.2.12 Multiple-channel support

In the event that more than one downstream signal is present in the system, the CM MUST operate using the first valid downstream signal that it encounters when scanning. It will be instructed via the parameters in the configuration file (see Annex C) to shift operation to different downstream and/or upstream frequencies if necessary.

Both upstream and downstream channels MUST be identified where required in MAC management messages using channel identifiers.

11.3 Standard operation

11.3.1 Periodic signal level adjustment

The CMTS MUST provide each CM a periodic ranging opportunity at least once every T4 seconds. The CMTS MUST send out periodic ranging opportunities at an interval sufficiently shorter than T4 that a MAP could be missed without the CM timing out. The size of this "subinterval" is CMTS-dependent.

The CM MUST re-initialize its MAC after T4 seconds have elapsed without receiving a periodic ranging opportunity.

Remote RF signal level adjustment at the CM is performed through a periodic ranging function using the RNG-REQ and RNG-RSP MAC messages. This is similar to initial ranging and is shown in Figures 11-16 and 11-17. On receiving a RNG-RSP, the CM MUST NOT transmit until the RF signal has been adjusted in accordance with the RNG-RSP and has stabilized (refer to clause 6). The CM MUST NOT transmit anything other than RNG-REQs, if it has been suspended by receiving a RNG-RSP with a ranging state of CONTINUE, until such time as it receives a RNG-RSP with a ranging state of SUCCESS. When operating in 1.x mode, the CM should continue normal operation and MUST NOT suspend on-going data transmission if it receives a RNG-RSP with a ranging status of CONTINUE.

The CMTS SHOULD NOT send a ranging status of CONTINUE in a RNG-RSP to any CMs in 2.0 mode unless the ranging parameters measured on the corresponding RNG-REQ are insufficient for the CMTS to guarantee proper reception of all burst types available to that CM. Additionally, upon sending a RNG-RSP with ranging status of CONTINUE to a CM in 2.0 mode, the CMTS SHOULD schedule another periodic ranging opportunity for the CM quickly so that the CM can return to a ranging status of SUCCESS as quickly as possible.

As described in clause 11.2.1, during normal operation in the S-CDMA mode, if a CM temporarily loses synchronization to the downstream signal, it is required to perform a ranging process before returning to normal operation. To facilitate this recovery, if the CMTS does not receive a RNG-REQ message from a CM during a station maintenance interval, the CMTS MAY schedule unicast initial maintenance opportunities, or it MAY temporarily reduce the time between unicast spreader-off station maintenance opportunities.

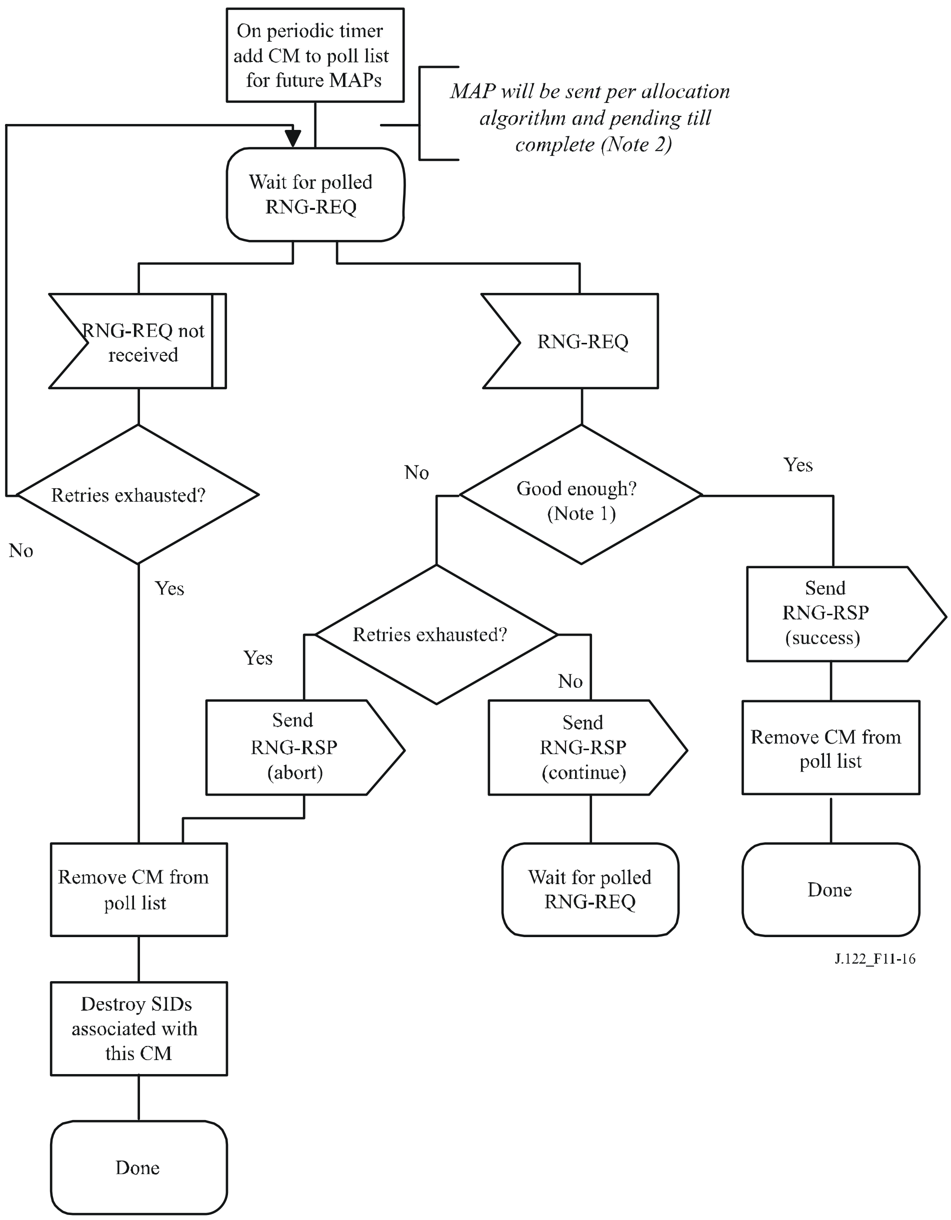

NOTE 1 – "Good enough" means ranging request is within the tolerance limits of the CMTS for power and transmit equalization (if supported).

NOTE 2 – RGN-REQ pending-till-complete was non-zero, the CMTS SHOULD hold off the station maintenance opportunity accordingly unless needed, for example, to adjust the CM's power level. If opportunities are offered prior to the pending-till-complete expiry, the "good-enough" test that follows receipt of RGN-RSP MUST NOT judge the CM's transmit equalization until pending-till-complete expires.

Figure 11-16 – Periodic ranging – CMTS

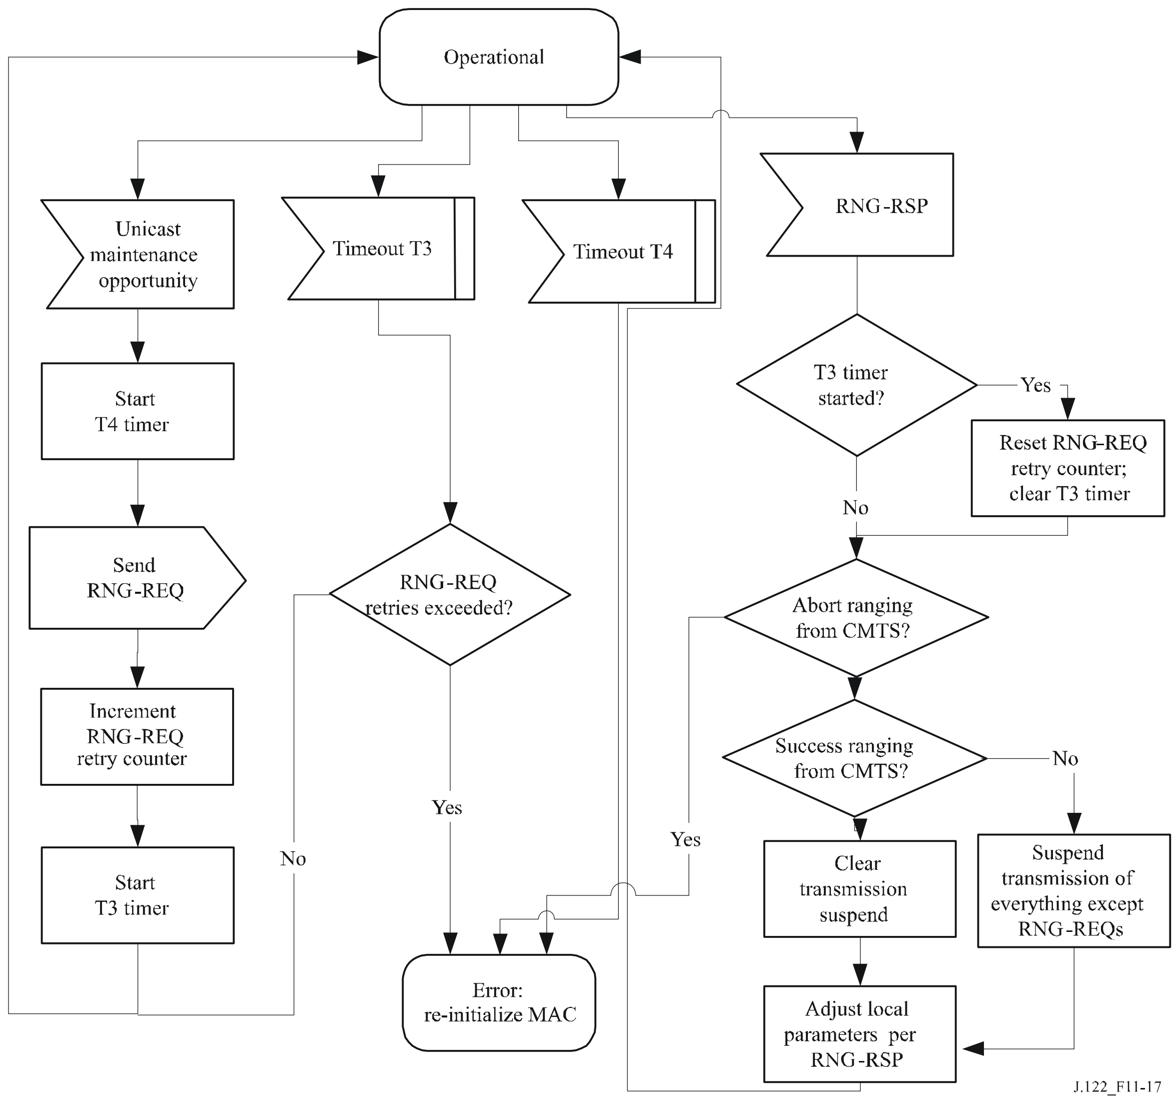

Figure 11-17 – Periodic ranging – CM view

11.3.2 Changing upstream channel descriptor message parameters

Whenever the CMTS is to change any of the upstream burst characteristics specified in the upstream channel descriptor (UCD) message (see clause 8.3.3), it must provide for an orderly transition from the old values to the new values by all CMs. Whenever the CMTS is to change any of the upstream characteristics, it MUST announce the new values in an UCD message, and the configuration change count field in that UCD message MUST be incremented to indicate that a value has changed.

After transmitting one or more UCD messages with the new value, the CMTS transmits a MAP message with a UCD change count matching the new configuration change count. The first interval in the MAP MUST be a data grant of at least 1 millisecond to the null service ID (zero). In the case of S-CDMA channels, that grant to the null service ID MUST be at least the longer of 1-ms or the duration of 2 S-CDMA frames to allow for the latency of the S-CDMA framing, and the start time of the MAP with the new UCD change count MUST correspond to the beginning of an S-CDMA frame. The CMTS MUST allow this time for cable modems to change their PMD sublayer parameters to match the new set. This time is independent of the lead time the CMTS needed to allow for in transmitting the MAP (see clause 9.1.5). The CMTS MUST transmit the new UCD message early enough that the CM receives the UCD message at least the UCD processing time (see Annex B) prior to the time the first MAP using the new UCD parameters arrives at the CM.

With the following exceptions, the CM MUST be able to transmit normally on the first grant following the grant to the NULL SID. The exceptions are the case where the new UCD message has changed the S-CDMA enable parameter, or the S-CDMA US ratio numerator or denominator. In these cases, the CM MAY redo initial ranging to establish timing synchronization for the new mode of operation before it resumes normal transmissions. If the CM was already registered with the CMTS, and it redoes initial ranging for this reason, it MUST use its primary SID instead of the initialization SID for the initial ranging process and it MUST NOT re-register. Using the ranging required parameter in the new UCD message, the CMTS can force the CM to perform ranging prior to making any other transmissions using the parameters in the new UCD message. In certain cases, channel-wide parameter changes (in particular, modulation rate or centre frequency) may invalidate pre-equalization and synchronization parameters, and normal operation may not be possible without re-ranging. If the CMTS changes the modulation rate or centre frequency on an S-CDMA channel, it MUST force re-ranging using the ranging required parameter.

In the case of an S-CDMA channel, the first UCD message with a new configuration count and any subsequent UCD messages that may be sent prior to the first MAP with the new UCD change count MUST have an updated timestamp snapshot corresponding to the start time of that first MAP with the new UCD change count. Also on an S-CDMA channel, the CMTS MUST maintain the continuity of the mini-slot and S-CDMA frame counters during the change in UCD parameters even if the size of a mini-slot is changed.

The CMTS MUST NOT transmit MAPs with the old UCD change count after transmitting the new UCD message.

The CM MUST use the parameters from the UCD message corresponding to the MAP's UCD change count for any transmissions it makes in response to that MAP. If the CM has, for any reason, not received the corresponding UCD message, it cannot transmit during the interval described by that MAP.

It is possible for the change in upstream parameters to cause the upstream to change from a type 1 upstream (see clause 11.2.2) to a type 2 upstream or a type 3 upstream. If this happens, and the CM registered with a configuration file that enables 2.0 mode (see clause 11.2.8), then the CM MUST operate in 2.0 mode. If the upstream has changed to a type 2 upstream this means that any request the CM transmits in an opportunity in the MAP with the new configuration change count or any subsequent MAP MUST be calculated in terms of IUCs 9 and 10, rather than IUCs 5 and 6, and the CMTS MUST issue grants using IUCs 9 and 10. However, if the CM registered with a configuration file that disabled 2.0 mode, then the CM MUST continue to calculate requests in terms of IUCs 5 and 6, and the CMTS MUST issue grants using IUCs 5 and 6. If the CM registered with a configuration file that disables 2.0 mode, and the new parameters have changed the upstream to a type 3 upstream, then the CM MUST immediately re-initialize the MAC layer and attempt registration. It should be understood by the network operator that changing a type 1 upstream to a type 3 upstream will cause a significant disruption of service for any CMs with 2.0 mode disabled as well as any DOCS 1.x CMs that are using the channel. Also such CMs will only be able to resume operation if there is a type 1 or type 2 upstream available to them.

11.3.3 Changing upstream channels

At any time after registration, the CMTS may direct the CM to change its upstream channel. This may be done for traffic balancing, noise avoidance or any of a number of other reasons which are beyond the scope of this Recommendation. Figure 11-18 shows the procedure that MUST be followed by the CMTS. Figures 11-19 and 11-20 show the corresponding procedure at the CM.

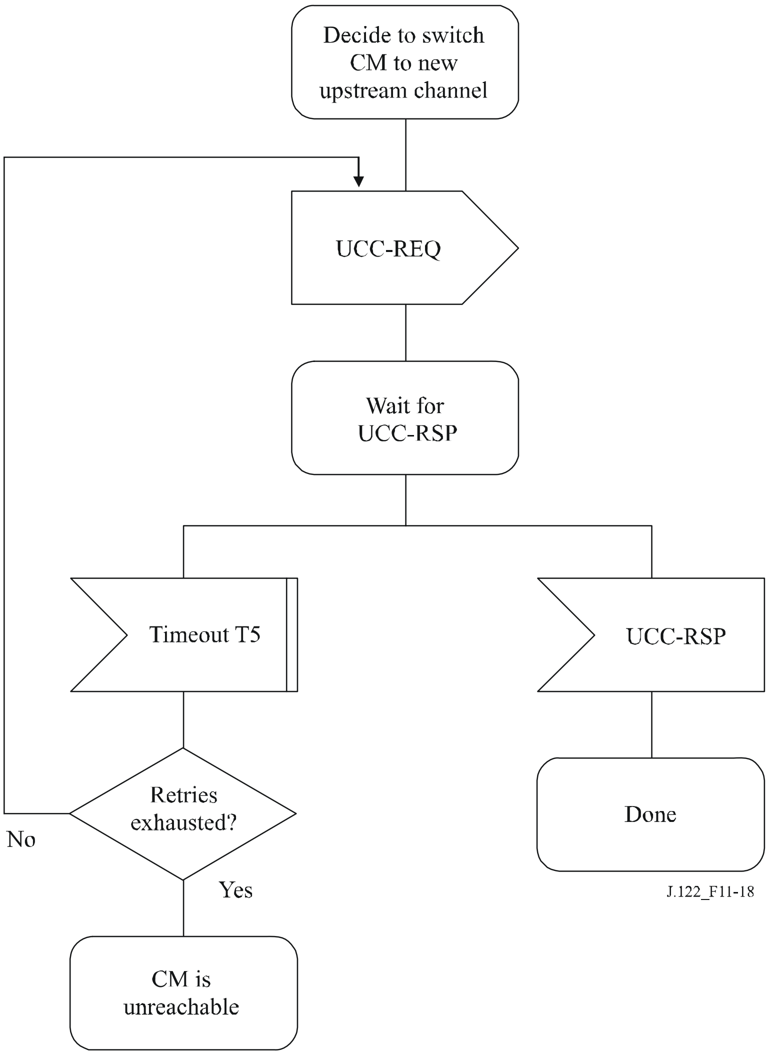

Figure 11-18 – Changing upstream channels: CMTS view

Note that if the CMTS retries the UCC-REQ, the CM may have already changed channels (if the UCC-RSP was lost in transit). Consequently, the CMTS MUST listen for the UCC-RSP on both the old and the new channels.

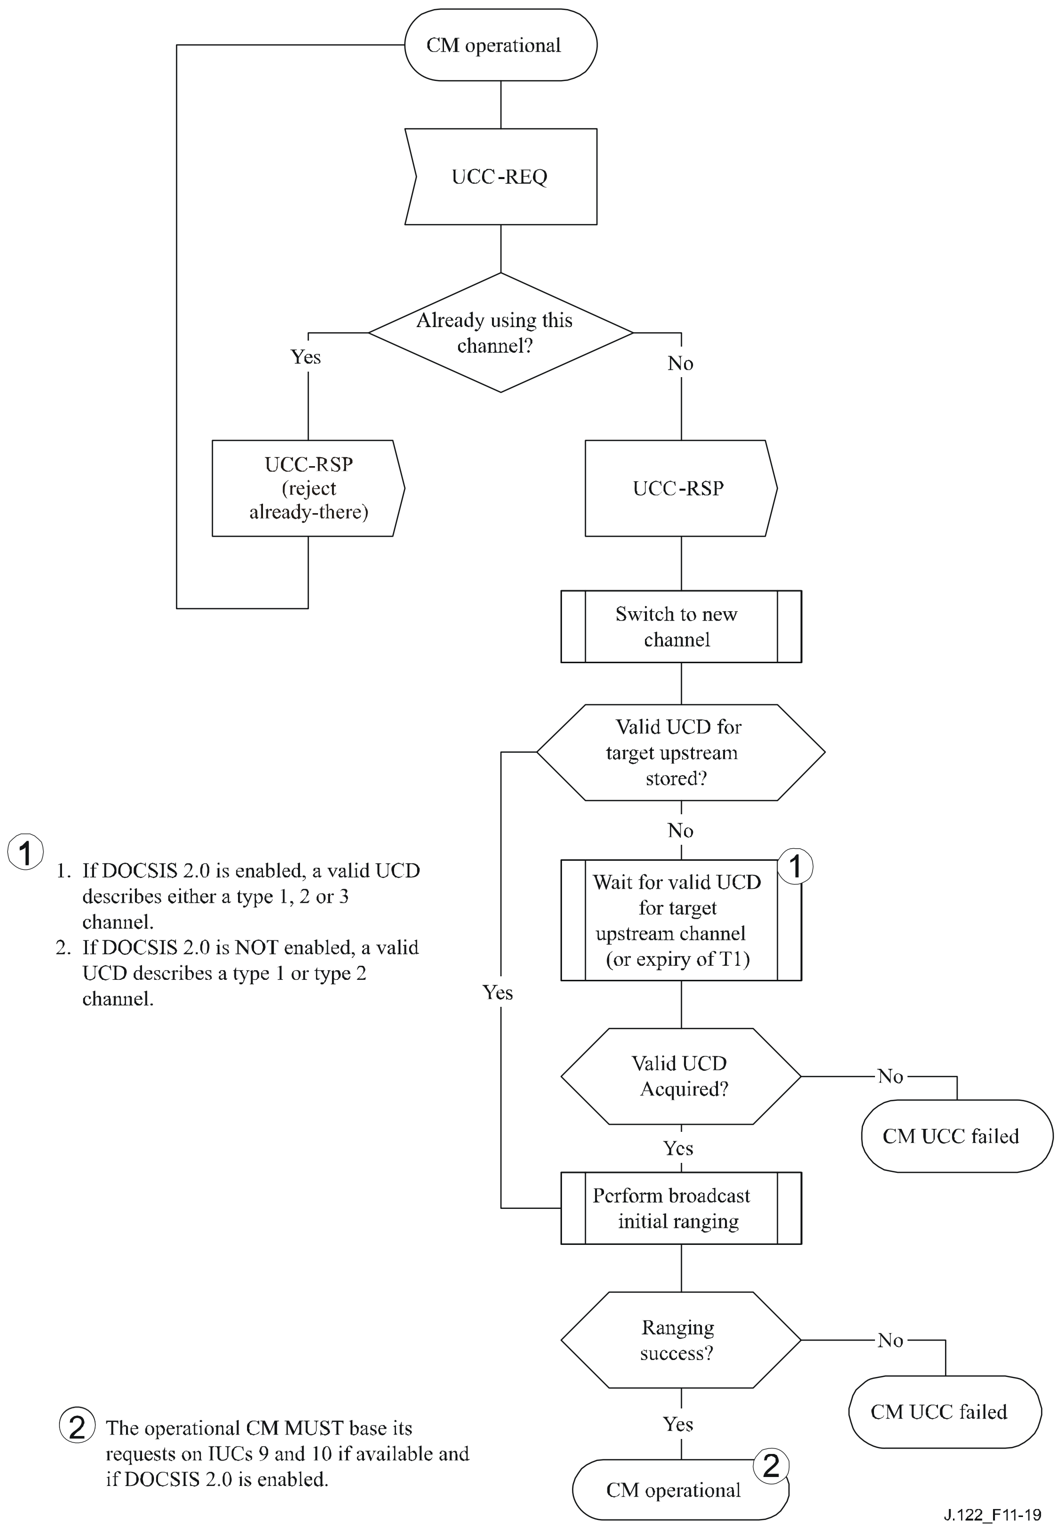

Figure 11-19 – Changing upstream channels: CM view, part 1

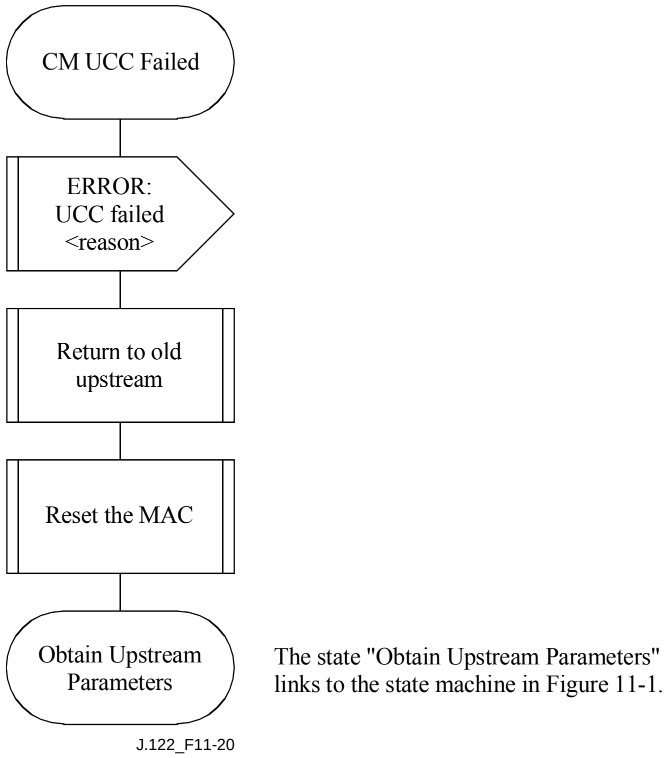

Figure 11-20 – Changing upstream channels: CM view, Part 2

Upon synchronizing with the new upstream channel, the CM MUST perform broadcast initial ranging on the new upstream channel.

If the CM has previously established ranging on the new channel, and if that ranging on that channel is still current (T4 has not elapsed since the last successful ranging), then the CM MAY use cached ranging information and omit ranging.

The CM SHOULD cache UCD information from multiple upstream channels to eliminate waiting for a UCD corresponding to the new upstream channel.

The CM MUST NOT perform re-registration, since its provisioning and MAC domain remain valid on the new channel.

If a DOCS 2.0-enabled CM is moved from a type 1 channel to a type 2 channel via a UCC, the CM MUST operate in 2.0 mode on the destination channel, basing its requests on IUCs 9 and 10 instead of IUCs 5 and 6. If a CM in which DOCS 2.0 is disabled in registration is moved from a type 1 channel to a type 2 channel via a UCC, the CM MUST continue to base its requests on the destination channel on IUCs 5 and 6.

11.4 Dynamic service

Service flows may be created, changed or deleted. This is accomplished through a series of MAC management messages referred to as dynamic service addition (DSA), dynamic service change (DSC) and dynamic service deletion (DSD). The DSA messages create a new service flow. The DSC messages change an existing service flow. The DSD messages delete a single existing upstream and/or a single existing downstream service flow. This is illustrated in Figure 11-21.

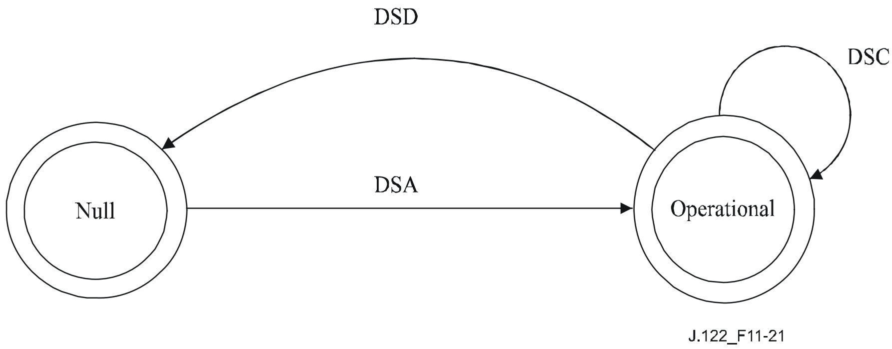

Figure 11-21 – Dynamic service flow overview

The null state implies that no service flow exists that matches the SFID and/or TransactionID in a message. Once the service flow exists, it is operational and has an assigned SFID. In steady state operation, a service flow resides in a nominal state. When dynamic service messaging is occurring, the service flow may transition through other states, but remains operational. Since multiple service flows may exist, there may be multiple state machines active, one for every service flow. Dynamic service messages only affect those state machines that match both the SFID and/or transaction ID or SFID only. A transaction ID which is reused for other SFID(s) indicates that the other side terminated the previous transaction. If a dynamic service request message is received that refers to the same transaction ID as one that has already been processed, but service flow(s) other than those locked in this transaction, the device MAY trigger a DSx ended input to the state machine(s) of SF(s) involved in the previous transaction. If privacy is enabled, both the CM and CMTS MUST verify the HMAC digest on all dynamic service messages before processing them, and discard any messages that fail.

Service flows created at registration time effectively enter the SF_operational state without a DSA transaction.

TransactionIDs are unique per transaction and are selected by the initiating device (CM or CMTS). To help prevent ambiguity and provide simple checking, the TransactionID number space is split between the CM and CMTS. The CM MUST select its TransactionIDs from the first half of the number space (0x0000 to 0x7FFF). The CMTS MUST select its TransactionIDs from the second half of the number space (0x8000 to 0xFFFF).

Each dynamic service message sequence is a unique transaction with an associated unique transaction identifier. To help support transaction identifier uniqueness between two devices in different states, the transaction initiating device SHOULD change the transaction identifier for each new initiated transaction. It MUST wait at least T10 to reuse the transaction identifier. The DSA/DSC transactions consist of a request/response/acknowledge sequence. The DSD transactions consist of a request/response sequence. The response messages MUST contain a confirmation code of okay unless some exception condition was detected. The acknowledge messages MUST include the confirmation code in the response unless a new exception condition arises. A more detailed state diagram, including transition states, is shown below. The detailed actions for each transaction will be given in the following clauses.

11.4.1 Dynamic service flow state transitions

The dynamic service flow state transition diagram is the top-level state diagram and controls the general service flow state. As needed, it creates transactions, each represented by a transaction state transition diagram, to provide the DSA, DSC and DSD signalling. Each transaction state transition diagram only communicates with the parent dynamic service flow state transition diagram. The top-level state transition diagram filters dynamic service messages and passes them to the appropriate transaction based on service flow identifier (SFID), service flow reference number and TransactionID.

If a single dynamic service message affects a pair of service flows, a single transaction is initiated which communicates with both parent dynamic service flow state transition diagrams. In this case, both service flows MUST remain locked in the same state until they receive the DSx succeeded or DSx failed input from the DSx transaction state transition diagram. During that "lock interval", if a message is received that refers to only one of the two service flows, it MUST be treated as if it refers to both service flows, so that both service flows stay in the same state. If a DSD-REQ message is received during the lock interval that refers to only one of the two service flows, the device MUST handle the event normally, by sending the SF Delete-Remote to the ongoing DSx transaction and by initiating a DSD-Remote transaction, and in addition, it MUST initiate a DSD-Local transaction to delete the second service flow of the locked pair.

If a DSC request is received that refers to two service flows locked in different transactions, and they are in different states, the device MUST reject the request without affecting the ongoing transactions.

There are six different types of transactions: locally initiated or remotely initiated for each of the DSA, DSC and DSD messages. Most transactions have three basic states: pending, holding and deleting. The pending state is typically entered after creation and is where the transaction is waiting for a reply. The holding state is typically entered once the reply is received. The purpose of this state is to allow for retransmissions in case of a lost message, even though the local entity has perceived that the transaction has completed. The deleting state is only entered if the service flow is being deleted while a transaction is being processed.

The flow diagrams provide a detailed representation of each of the states in the transaction state transition diagrams. All valid transitions are shown. Any inputs not shown should be handled as a severe error condition.

With one exception, these state diagrams apply equally to the CMTS and CM. In the dynamic service flow changing-local state, there is a subtle difference in the CM and CMTS behaviours. This is called out in the state transition and detailed flow diagrams.

NOTE – The 'Num Xacts' variable in the dynamic service flow state transition diagram is incremented every time the top-level state diagram creates a transaction and is decremented every time a transaction terminates. A dynamic service flow MUST NOT return to the null state until it is deleted and all transactions have terminated.

The inputs for the state diagrams are identified below.

Dynamic service flow state transition diagram inputs from unspecified local, higher-level entities:

- Add;
- Change;
- Delete.

Dynamic service flow state transition diagram inputs from DSx transaction state transition diagrams:

- DSA Succeeded;
- DSA Failed;
- DSA ACK Lost;
- DSA Erred;
- DSA Ended;
- DSC Succeeded;
- DSC Failed;

- DSC ACK Lost;
- DSC Erred;
- DSC Ended;
- DSD Succeeded;
- DSD Erred;
- DSD Ended.

DSx transaction state transition diagram inputs from the dynamic service flow state transition diagram:

- SF Add;
- SF Change;
- SF Delete;
- SF Abort Add;
- SF Change-Remote;
- SF Delete-Local;
- SF Delete-Remote;
- SF DSA-ACK Lost;
- SF-DSC-REQ Lost;
- SF-DSC-ACK Lost;
- SF DSD-REQ Lost;
- SF Changed;
- SF Deleted.

The creation of DSx transactions by the dynamic service flow state transition diagram is indicated by the notation:

DSx-[Local|Remote] (initial_input)

where initial_input may be SF Add, DSA-REQ, SF Change, DSC-REQ, SF Delete or DSD-REQ depending on the transaction type and initiator.

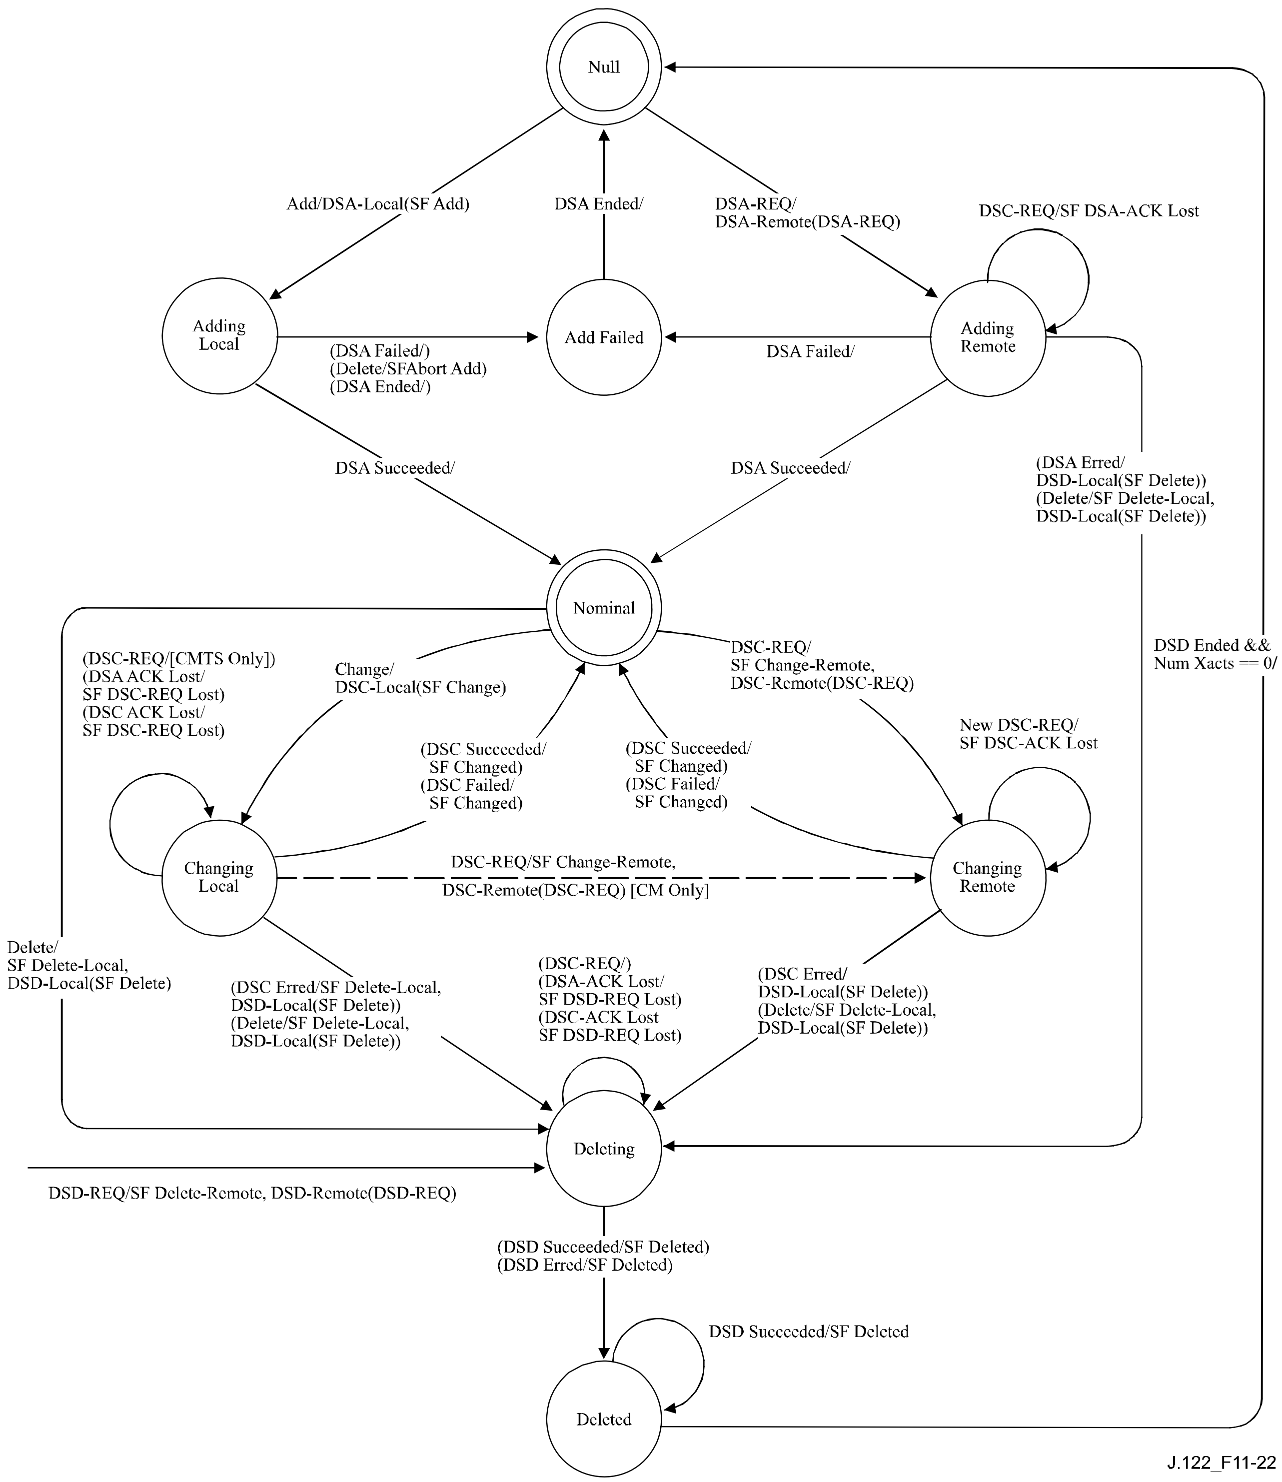

Figure 11-22 – Dynamic service flow state transition diagram

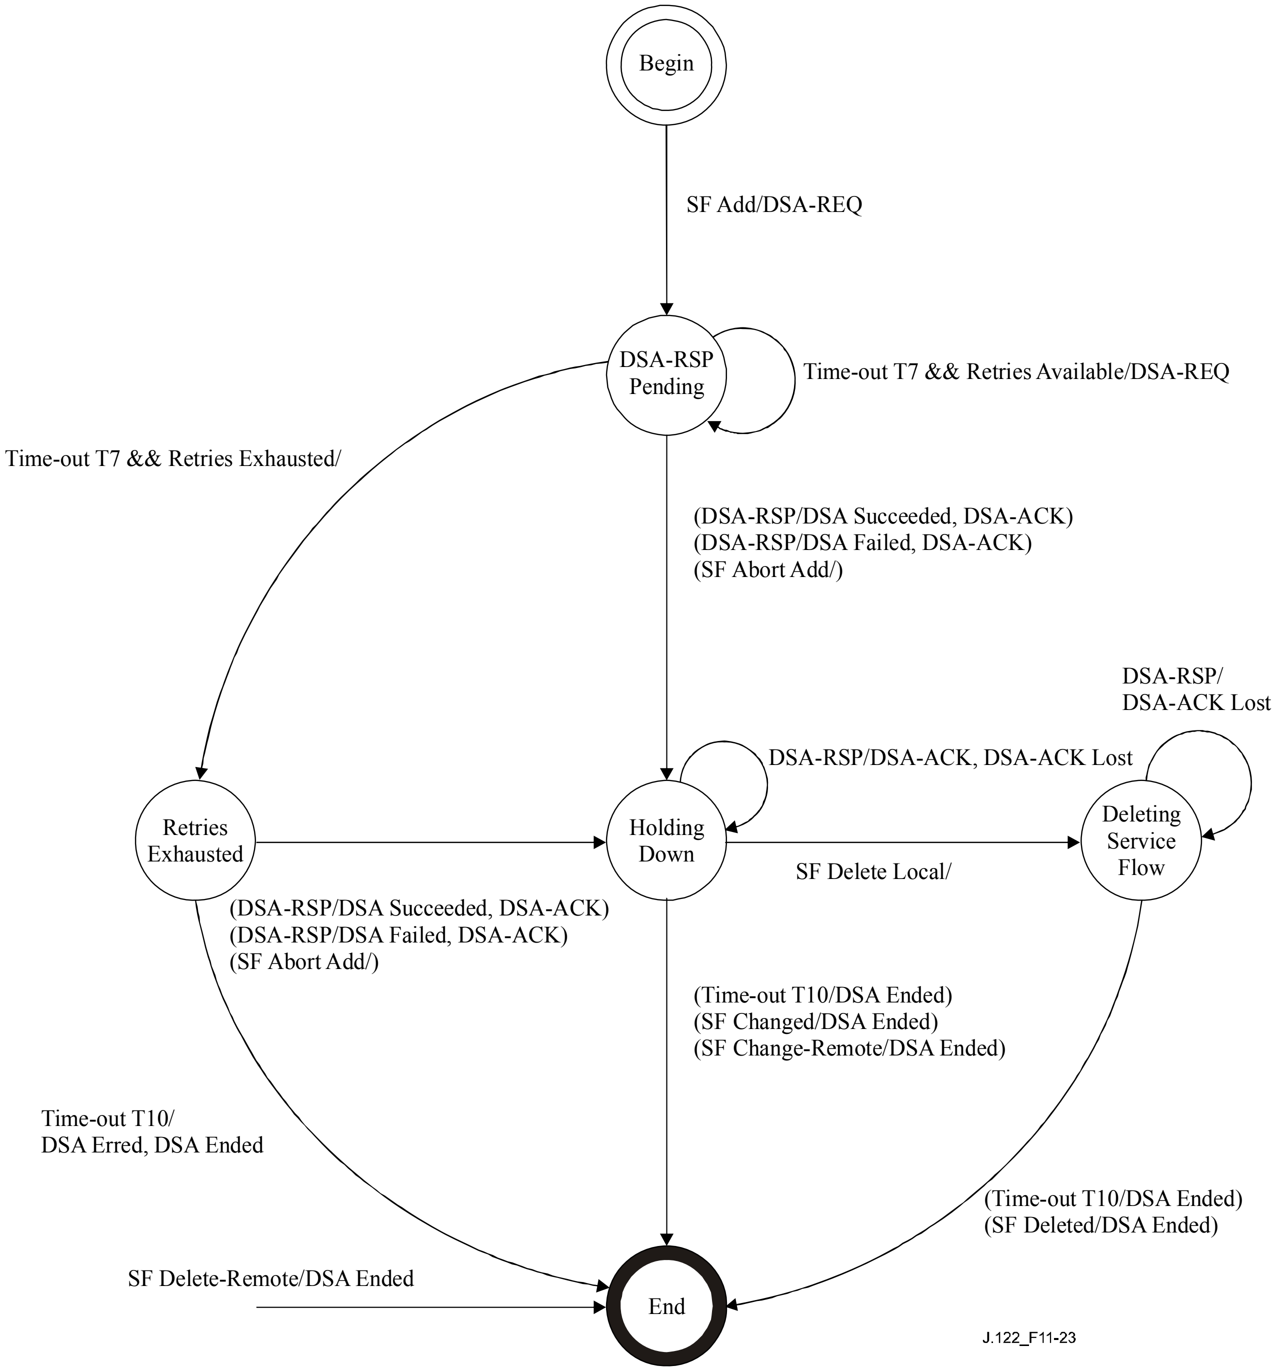

Figure 11-23 – DSA – Locally-initiated transaction state transition diagram

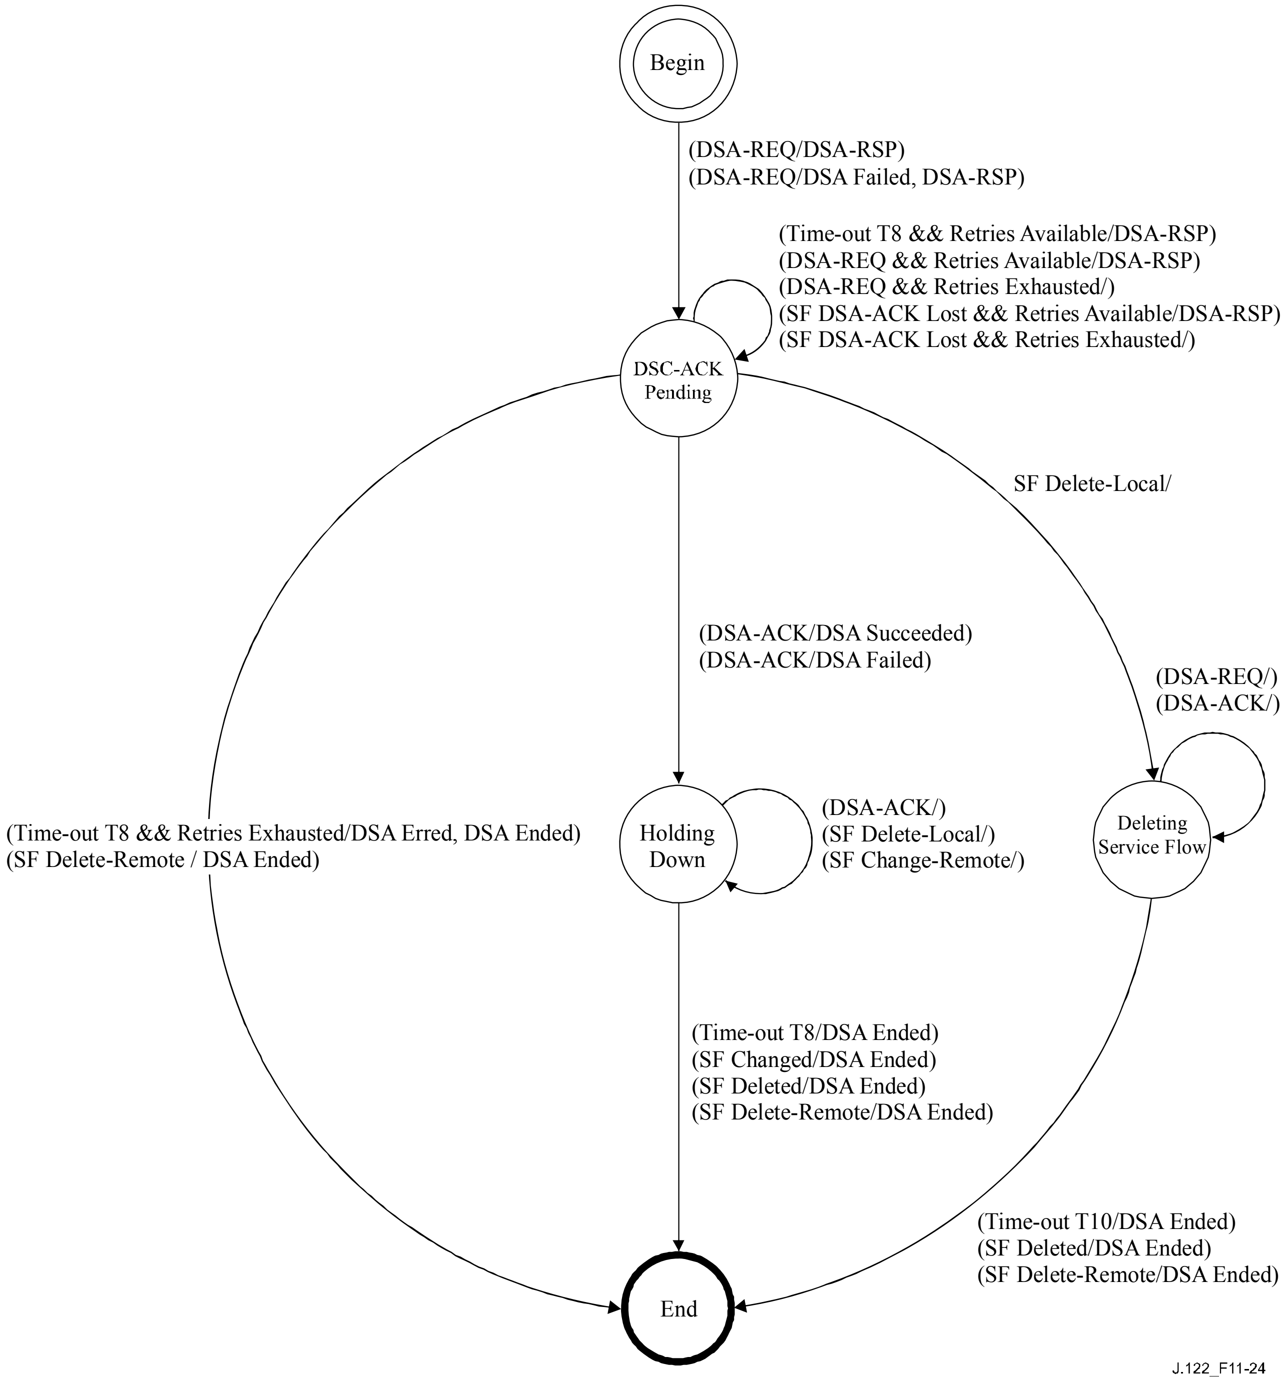

Figure 11-24 – DSA – Remotely-initiated transaction state transition diagram

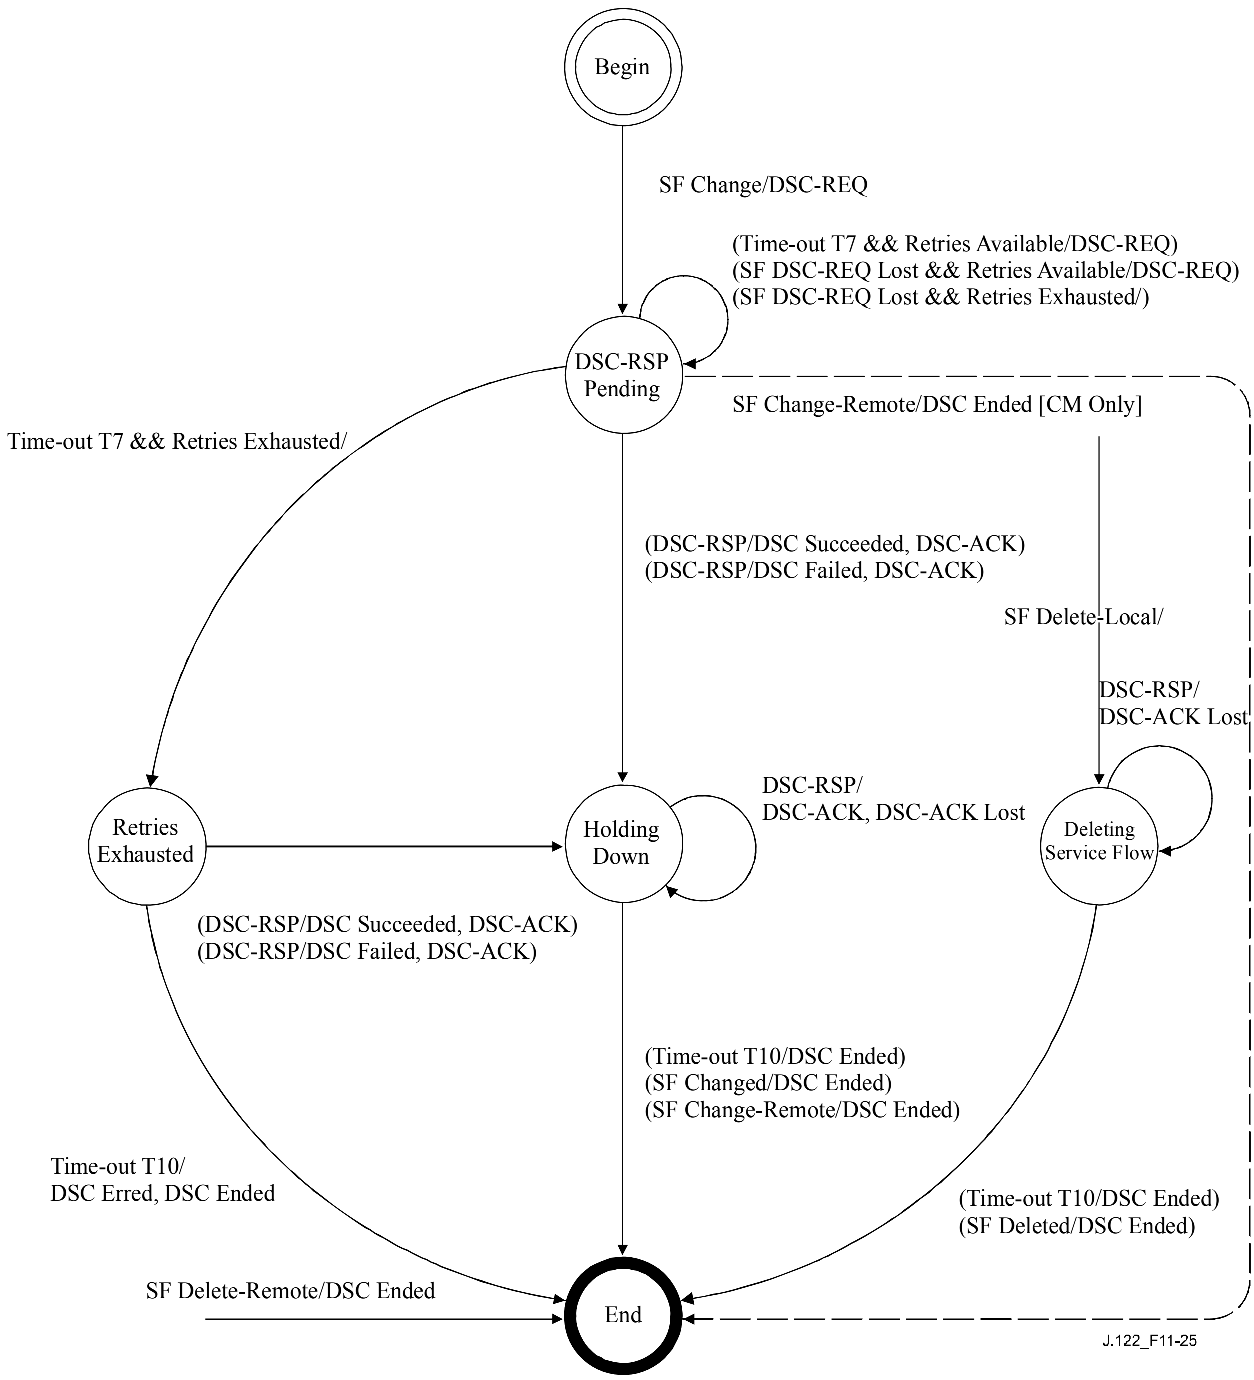

Figure 11-25 – DSC – Locally-initiated transaction state transition diagram

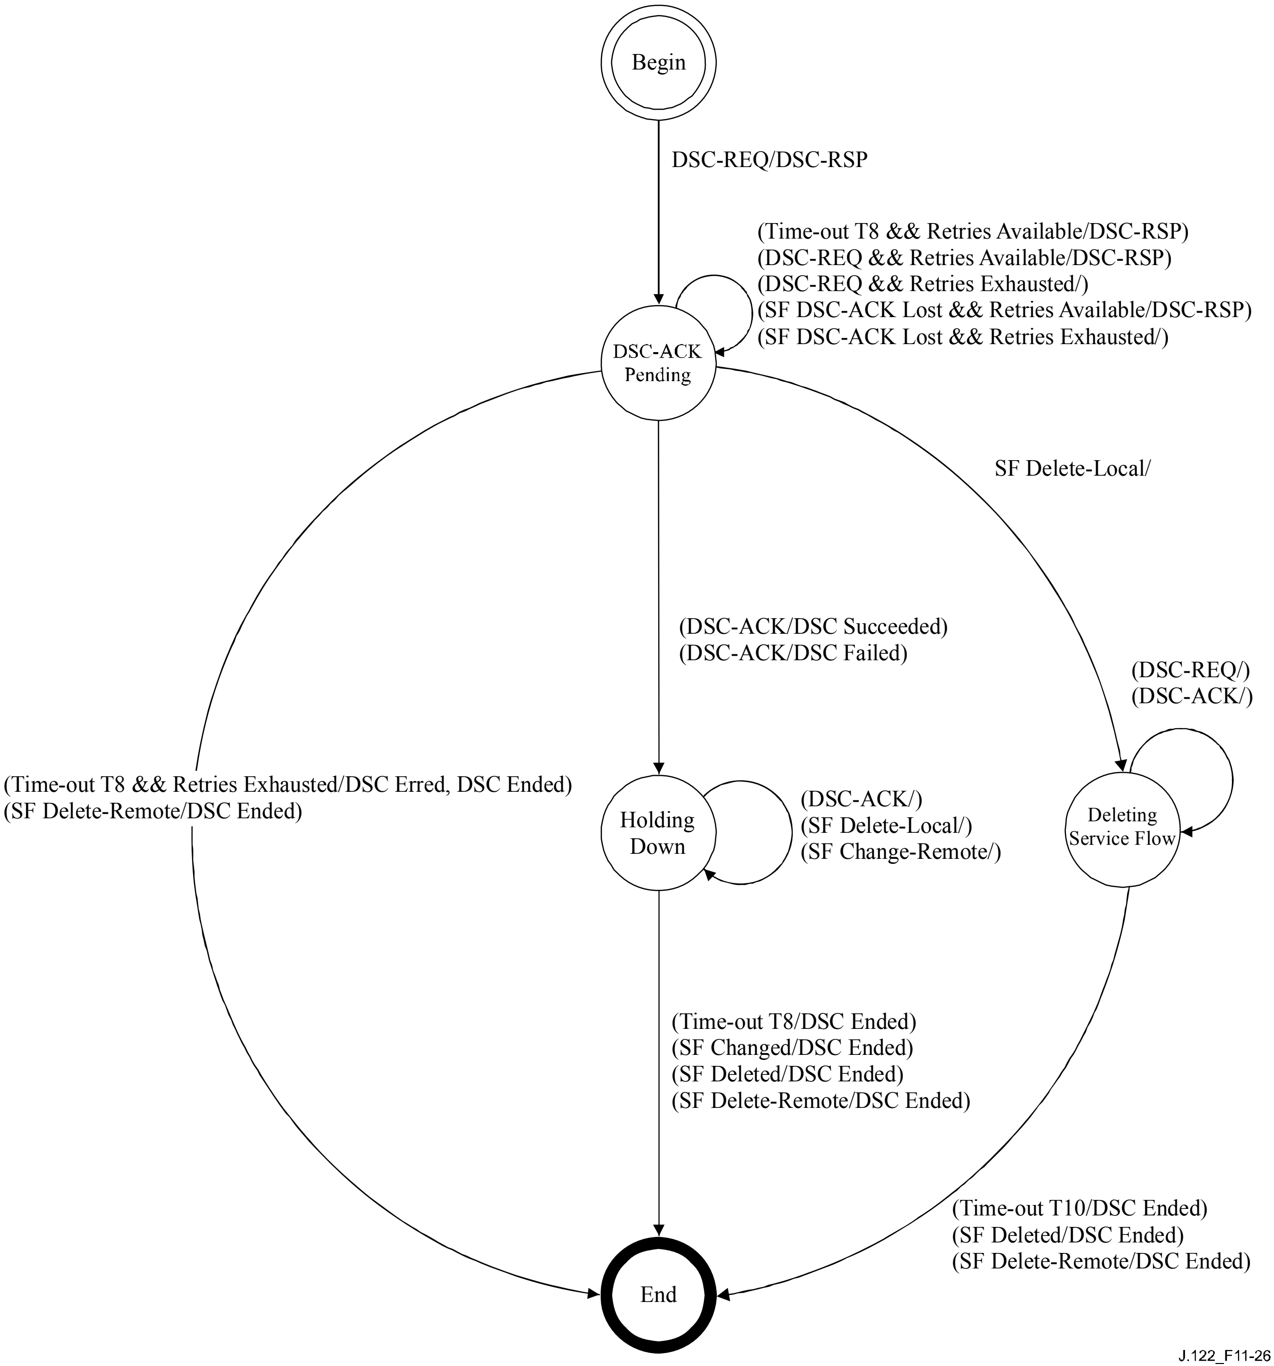

Figure 11-26 – DSC – Remotely-initiated transaction state transition diagram

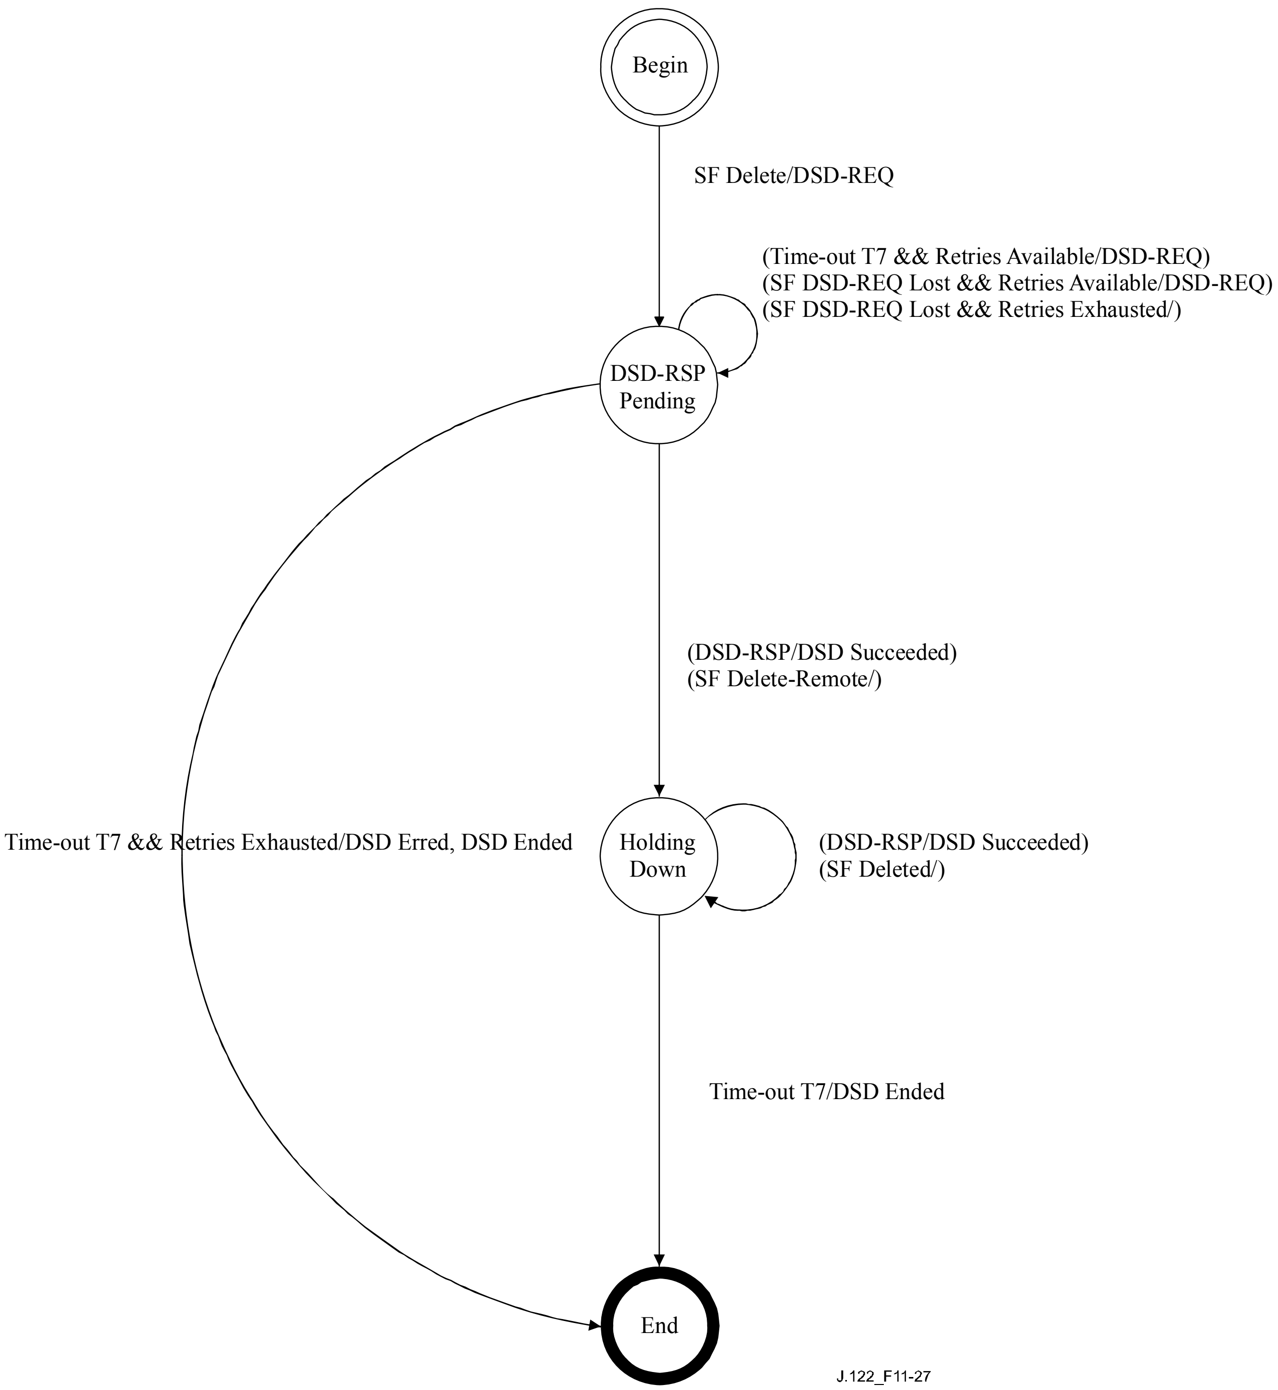

Figure 11-27 – DSD – Locally-initiated transaction state transition diagram

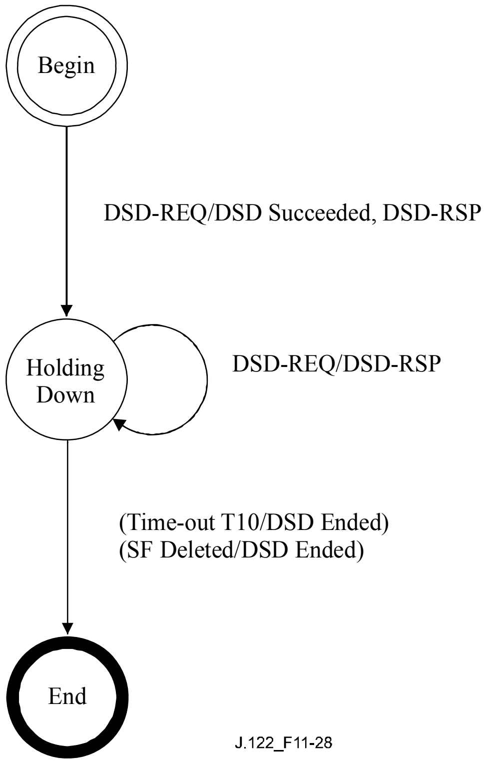

Figure 11-28 – Dynamic deletion (DSD) – Remotely-initiated transaction state transition diagram

11.4.2 Dynamic service addition

11.4.2.1 CM-initiated dynamic service addition

A CM wishing to create an upstream and/or a downstream service flow sends a request to the CMTS using a dynamic service addition request message (DSA-REQ). The CMTS checks the CM's authorization for the requested service(s) and whether the QoS requirements can be supported, and generates an appropriate response using a dynamic service addition response message (DSA-RSP). The CM concludes the transaction with an acknowledgment message (DSA-ACK).

In order to facilitate a common admission response, an upstream and a downstream service flow can be included in a single DSA-REQ. Both service flows are either accepted or rejected together.

CM		CMTS
New service flow(s) needed Check if resources are available		
Send DSA-REQ	—— DSA-REQ ——→	**Receive DSA-REQ**
		Check if CM authorized for service(s)[a)] Check service flow(s) QoS can be supported Create SFID(s) If upstream AdmittedQosParamSet is non-null, create SID If upstream ActiveQosParamSet is non-null, enable reception of data on new upstream service flow
Receive DSA-RSP	←—— DSA-RSP ——	**Send DSA-RSP**
If ActiveQosParamSet is non-null, enable transmission and/or reception of data on new service flow(s)		
Send DSA-ACK	—— DSA-ACK ——→	**Receive DSA-ACK**
		If downstream ActiveQosParamSet is non-null, enable transmission of data on new downstream service flow

a) Authorization can happen prior to the DSA-REQ being received by the CMTS. The details of CMTS signalling to anticipate a DSA-REQ are beyond the scope of this Recommendation.

Figure 11-29 – Dynamic service addition initiated from CM

11.4.2.2 CMTS-initiated dynamic service addition

A CMTS wishing to establish an upstream and/or a downstream dynamic service flow(s) with a CM performs the following operations. The CMTS checks the authorization of the destination CM for the requested class of service and whether the QoS requirements can be supported. If the service can be supported, the CMTS generates new SFID(s) with the required class of service and informs the CM using a dynamic service addition request message (DSA-REQ). The CM checks that it can support the service and responds using a dynamic service addition response message (DSA-RSP). The transaction completes with the CMTS sending the acknowledge message (DSA-ACK).

CM		CMTS
		New service flow(s) required for CM Check CM authorized for service(s) Check service flow(s) QoS can be supported Create SFID(s) If upstream AdmittedQosParamSet is non-null, create SID If upstream ActiveQosParamSet is non-null, enable reception of data on new upstream service flow
Receive DSA-REQ	←—— DSA-REQ ——	**Send DSA-REQ**
Confirm CM can support service flow(s) Add downstream SFID (if present) Enable reception on any new downstream service flow		
Send DSA-RSP	—— DSA-RSP ——→	**Receive DSA-RSP**
		Enable transmission and reception of data on new service flow(s)
Receive DSA-ACK	←—— DSA-ACK ——	**Send DSA-ACK**
Enable transmission on new upstream service flow		

Figure 11-30 – Dynamic service addition initiated from CMTS

11.4.2.3 Dynamic service addition state transition diagrams

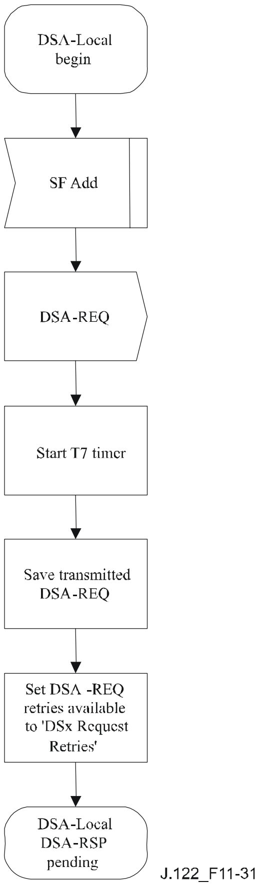

Figure 11-31 – DSA – Locally-initiated transaction begin state flow diagram

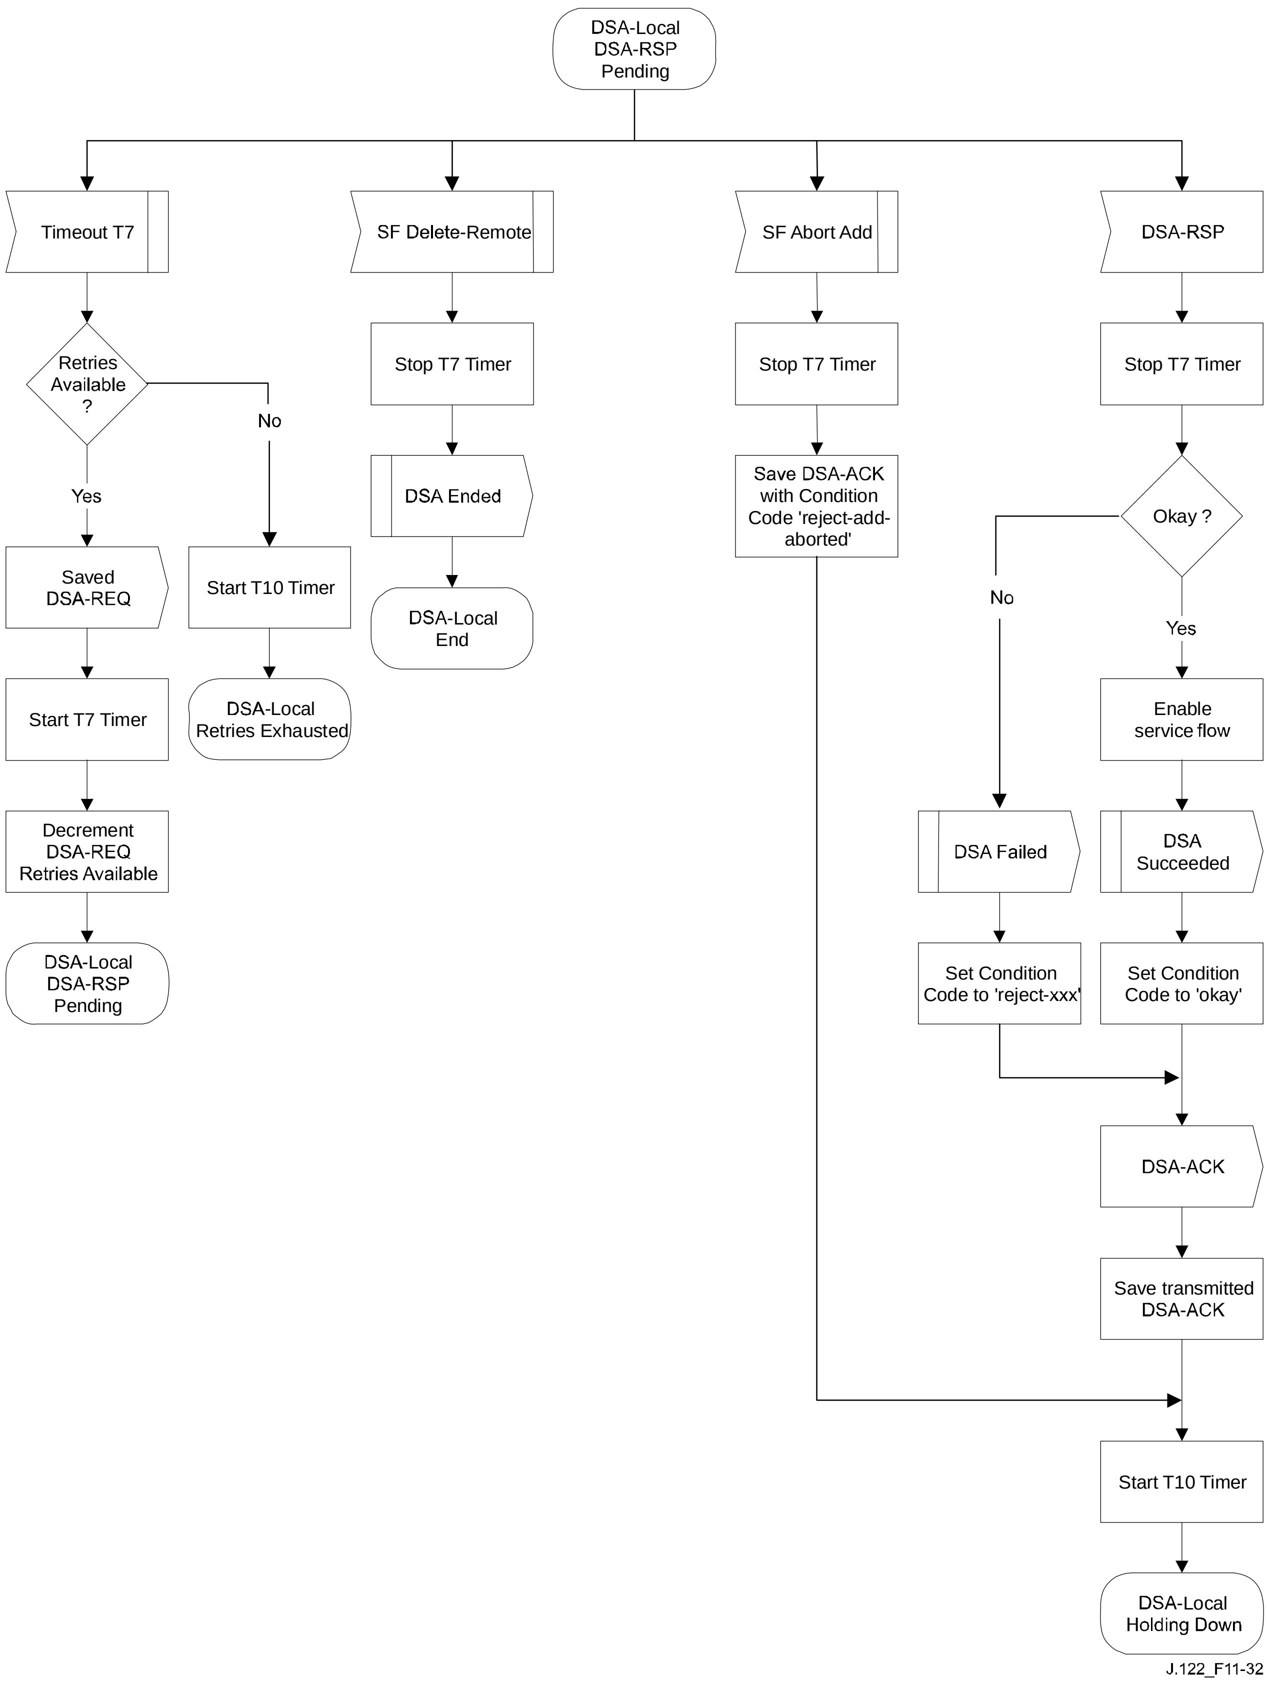

Figure 11-32 – DSA – Locally-initiated transaction DSA-RSP pending state flow diagram

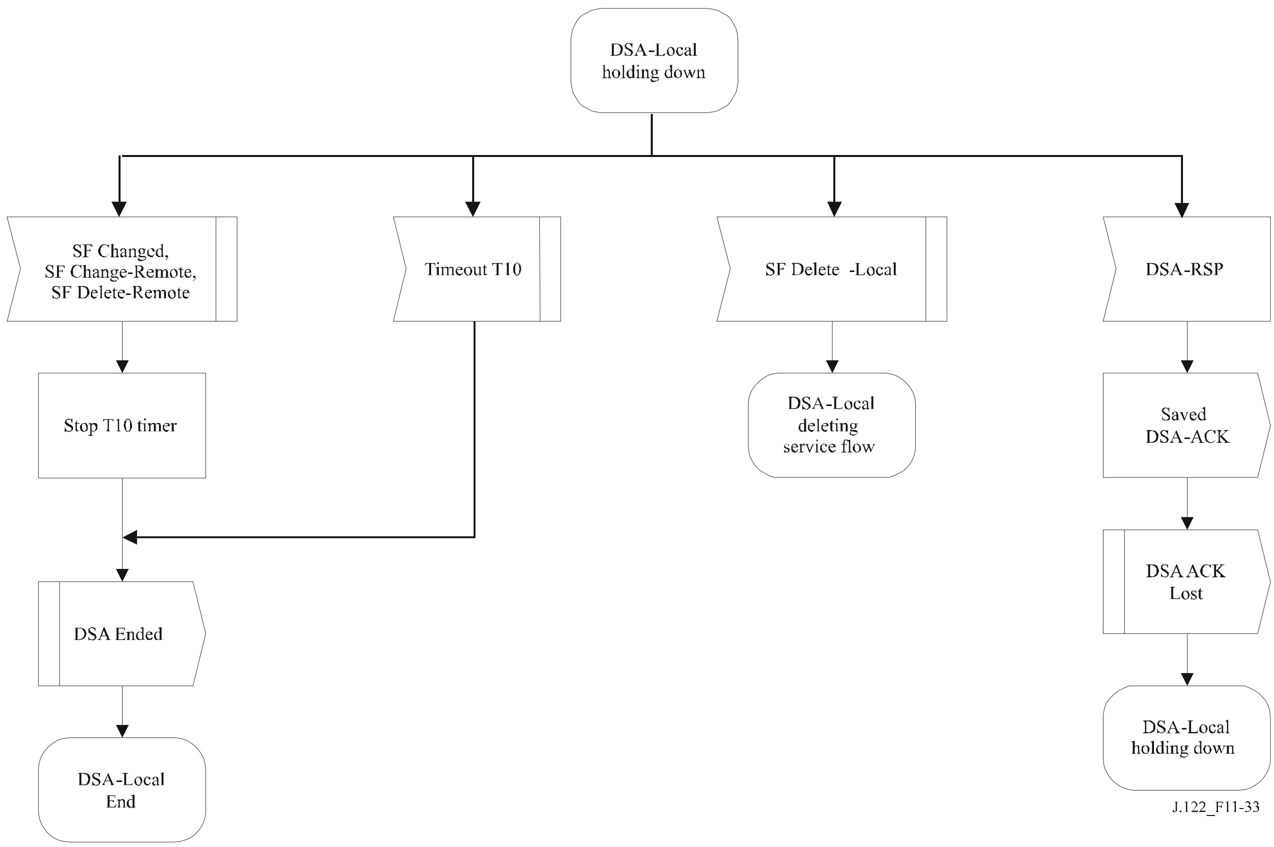

Figure 11-33 – DSA – Locally-initiated transaction holding state flow diagram

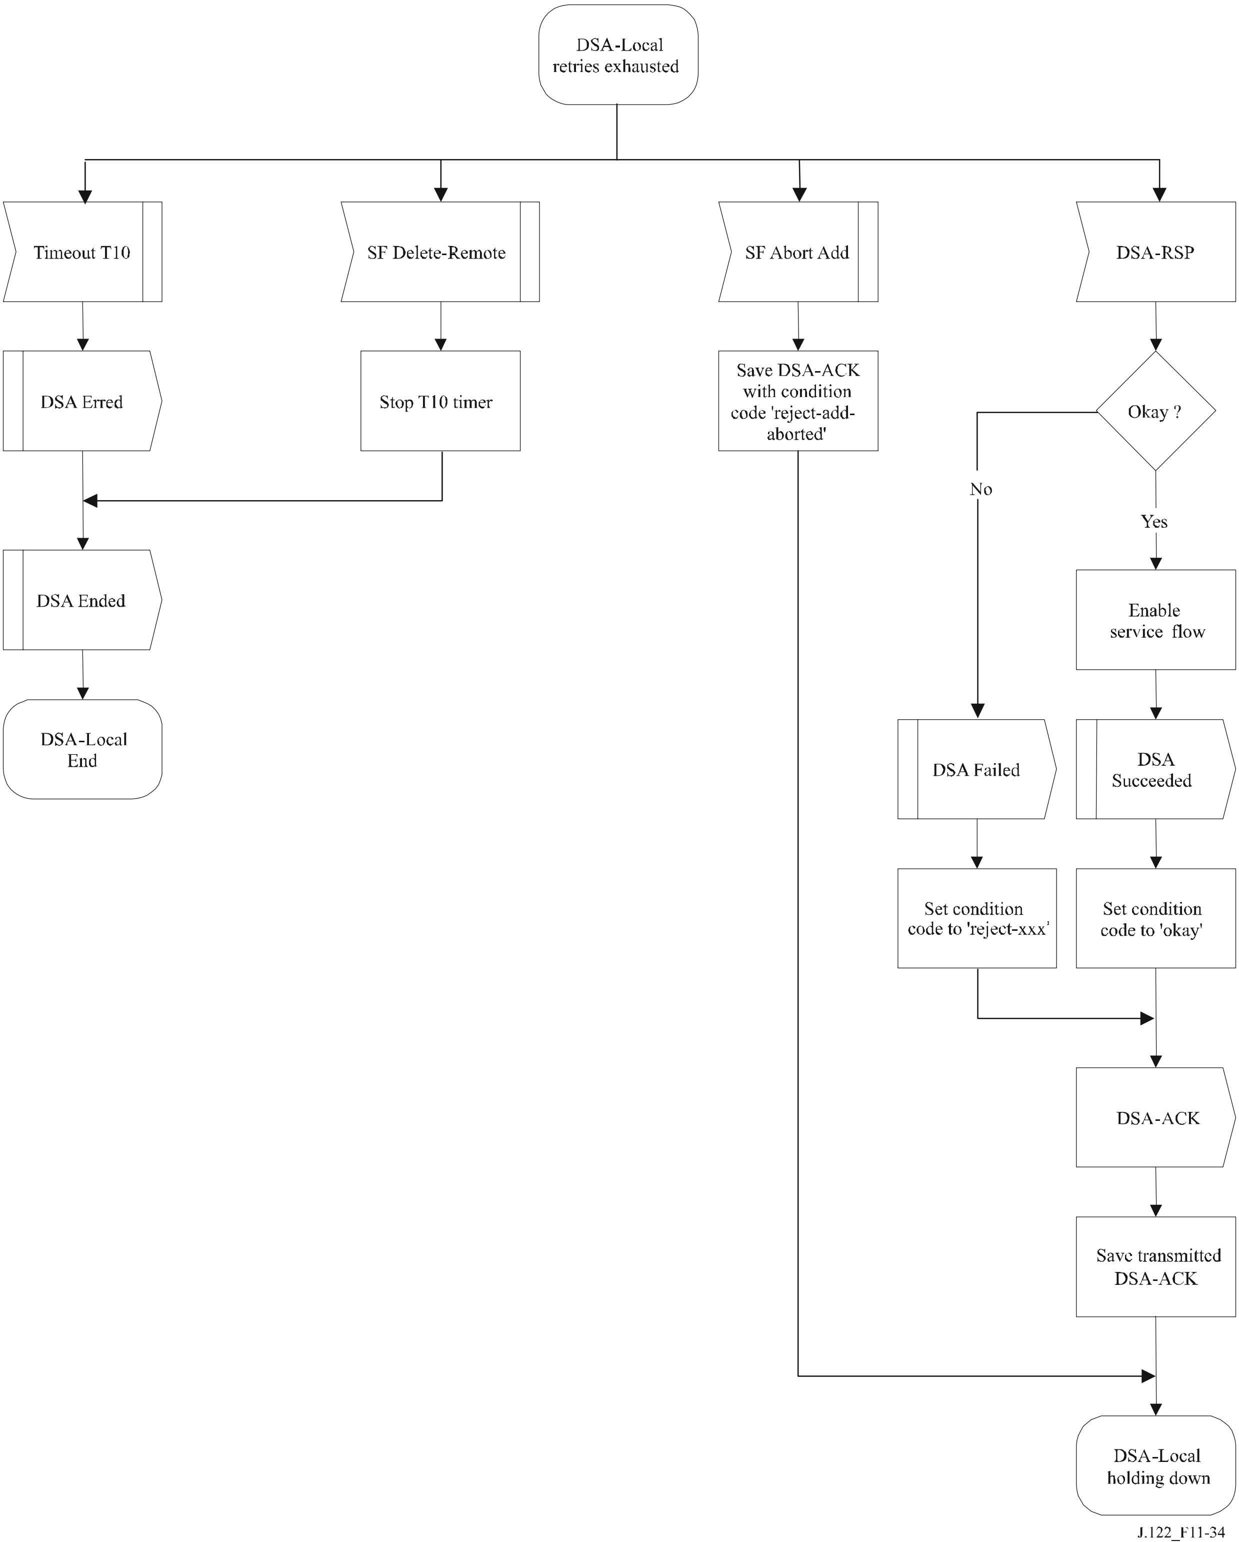

Figure 11-34 – DSA – Locally-initiated transaction retries exhausted state flow diagram

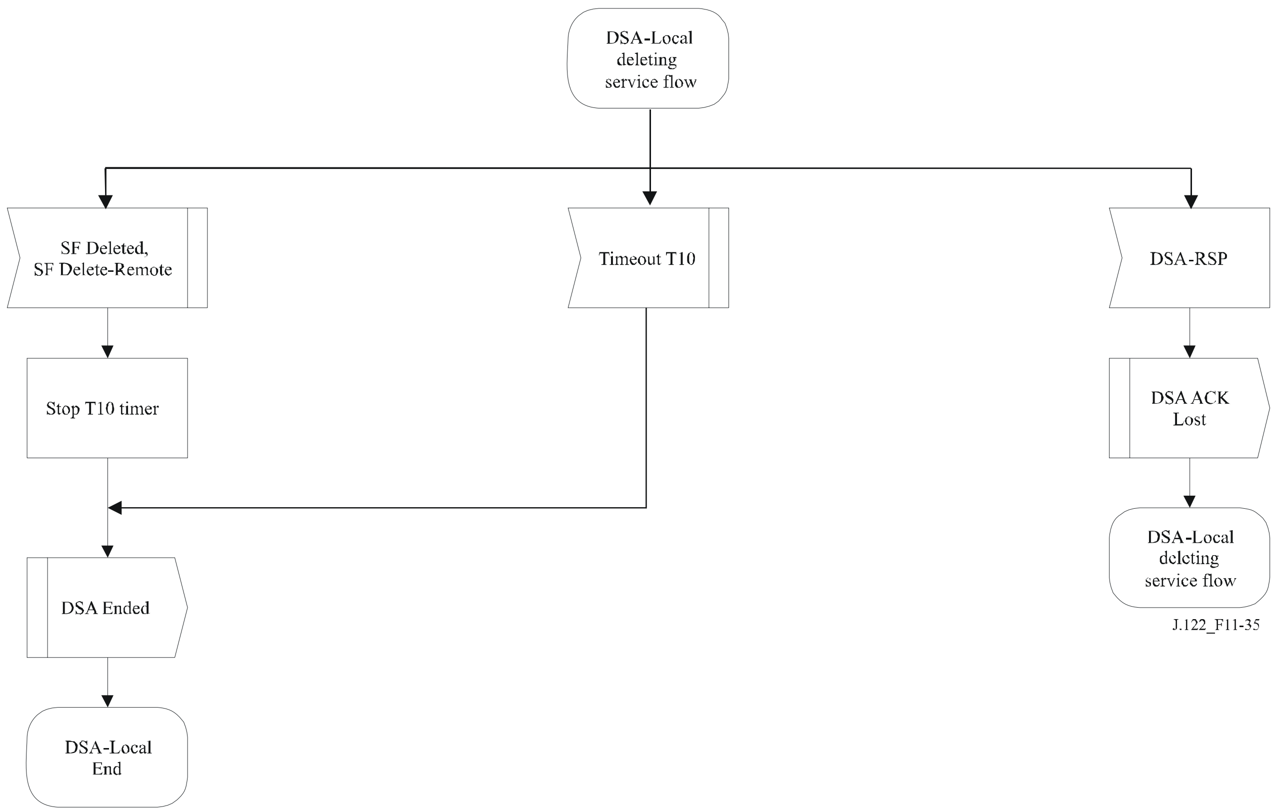

Figure 11-35 – DSA – Locally-initiated transaction deleting service flow state flow diagram

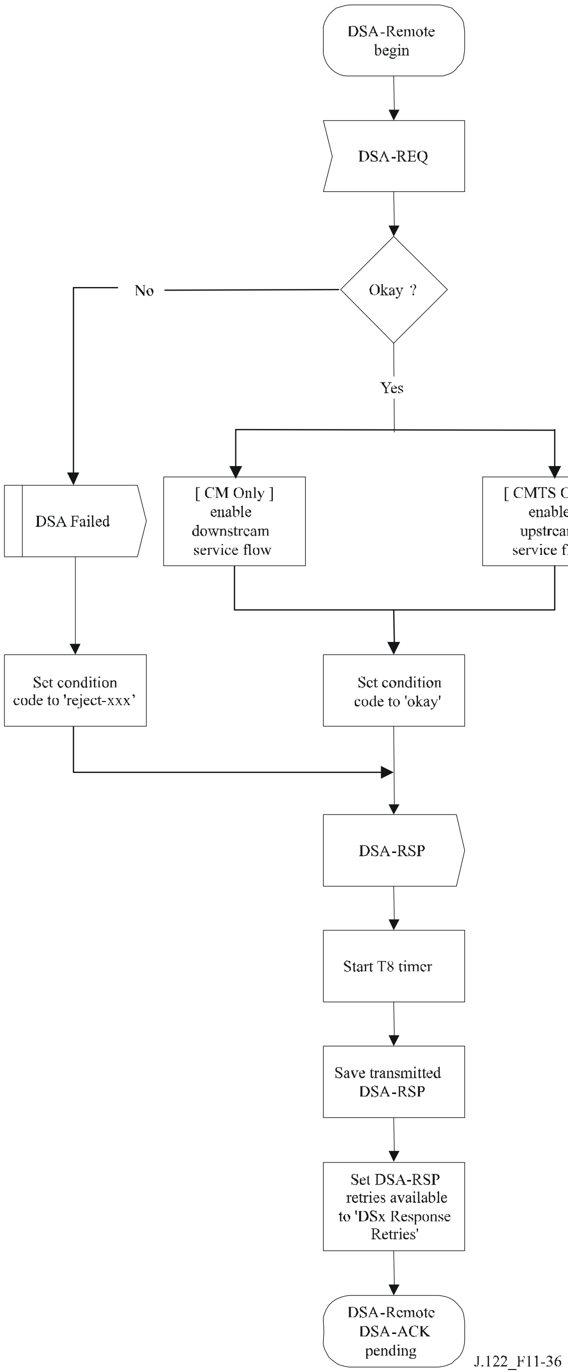

Figure 11-36 – DSA – Remotely-initiated transaction begin state flow diagram

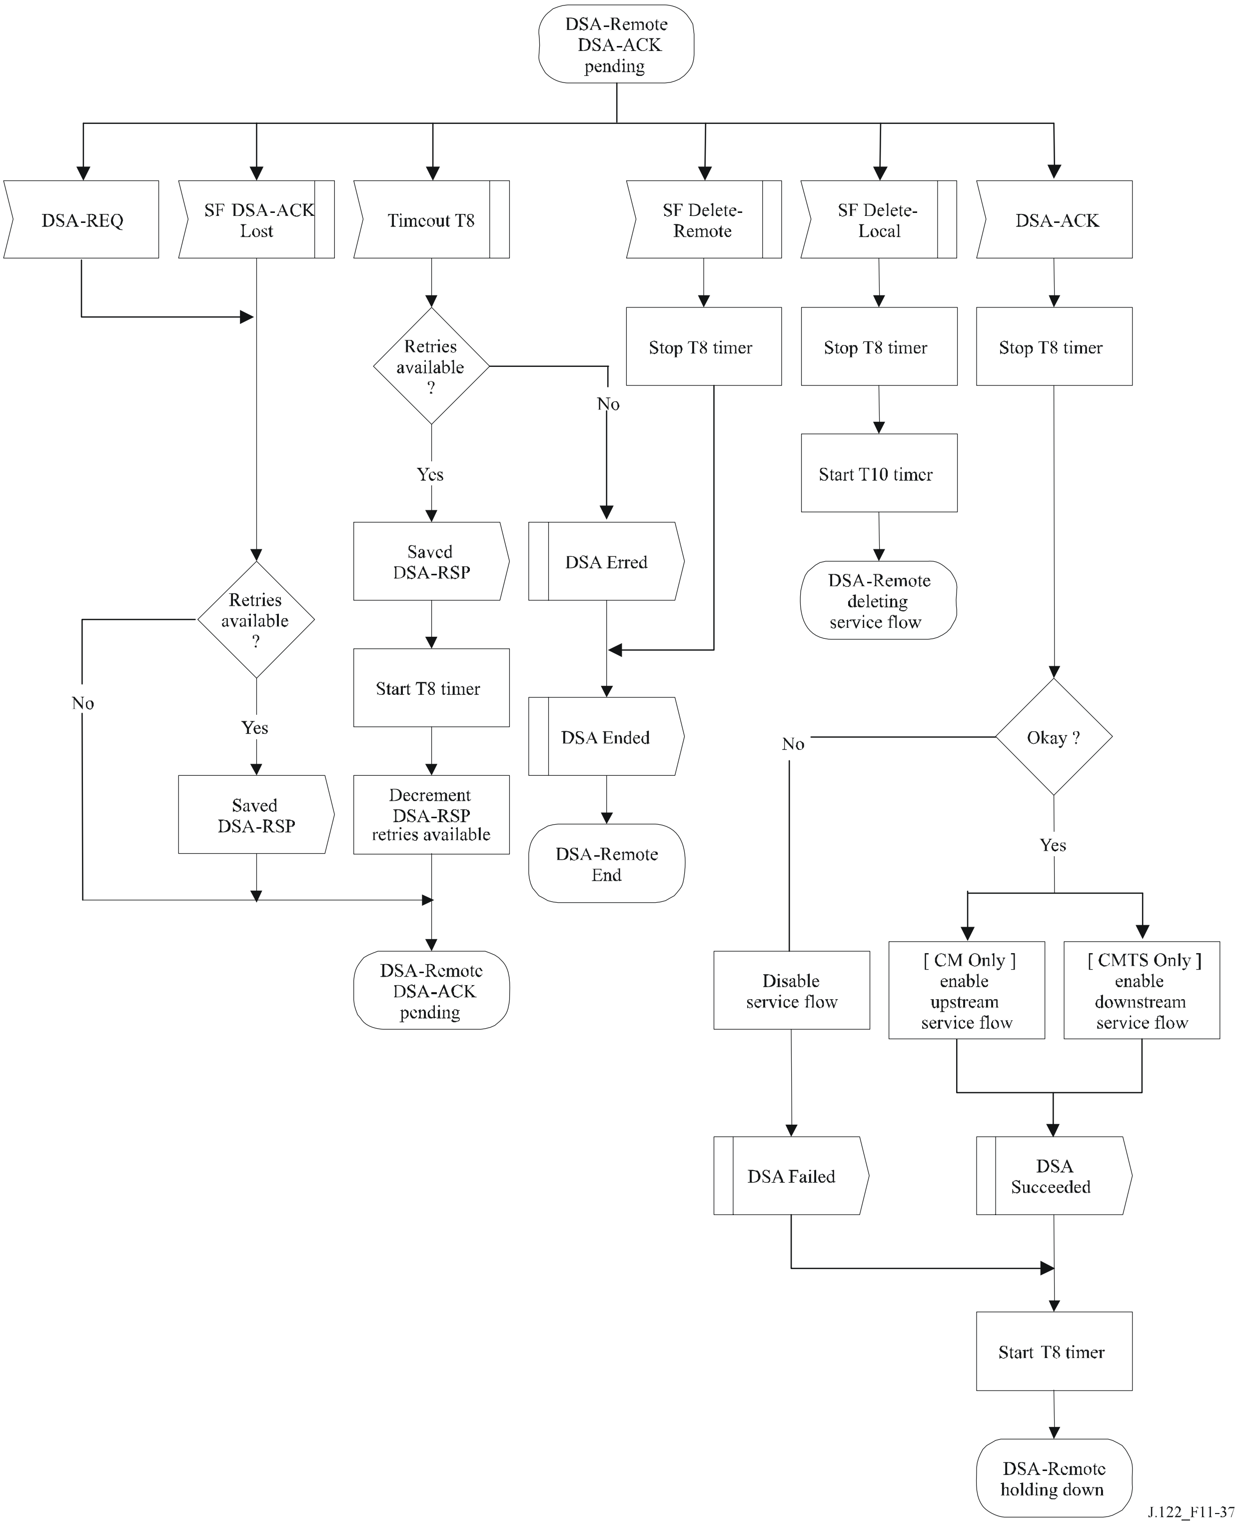

Figure 11-37 – DSA – Remotely-initiated transaction DSA-ACK pending state flow diagram

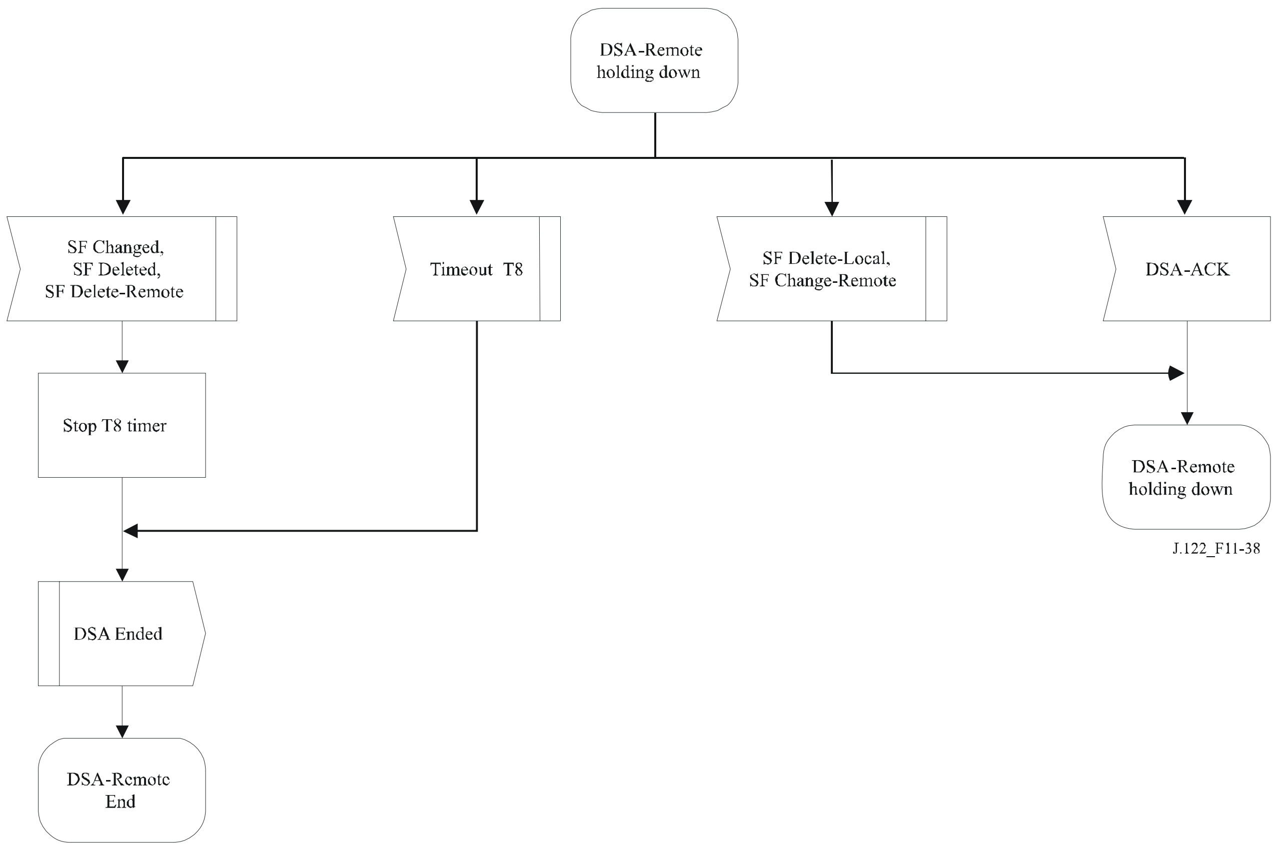

Figure 11-38 – DSA – Remotely-initiated transaction holding down state flow diagram

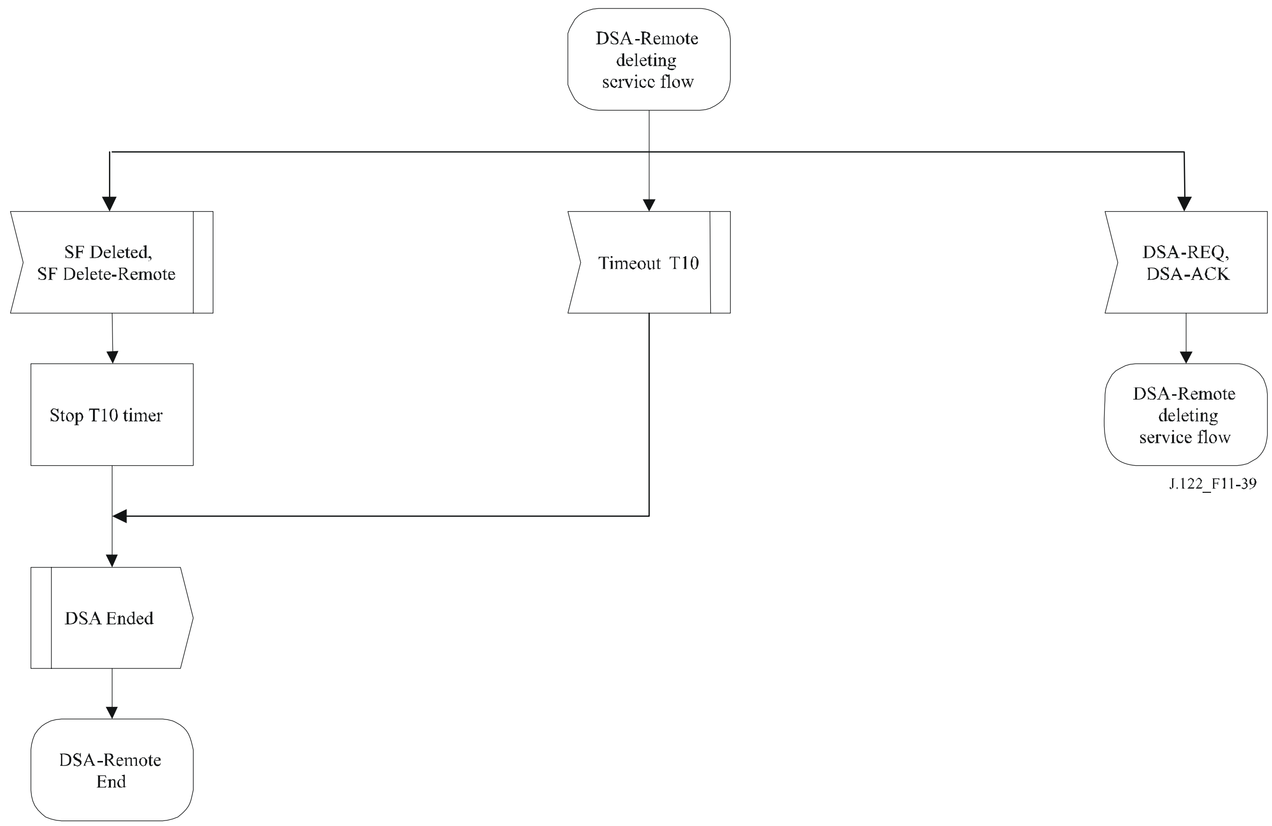

Figure 11-39 – DSA – Remotely-initiated transaction deleting service state flow diagram

11.4.3 Dynamic service change

The dynamic service change (DSC) set of messages is used to modify the flow parameters associated with a service flow. Specifically, DSC can:

- modify the service flow specification;
- add, delete or replace a flow classifier;
- add, delete or set PHS elements.

A single DSC message exchange can modify the parameters of one downstream service flow and/or one upstream service flow.

To prevent packet loss, any required bandwidth change must be sequenced between the application generating the data and the bandwidth parameters of the service flow carrying the data. Because MAC messages can be lost, the timing of service flow parameter changes can vary, and it occurs at different times in the CM and CMTS. Applications should reduce their transmitted data bandwidth before initiating a DSC to reduce the service flow bandwidth, and should not increase their transmitted data bandwidth until after the completion of a DSC increasing the service flow bandwidth.

The CMTS controls both upstream and downstream scheduling. Scheduling is based on data transmission requests and is subject to the limits contained in the current service flow parameters at the CMTS. The timing of service flow parameter changes, and any consequent scheduling changes, is independent of both direction and whether there is an increase or decrease in bandwidth. The CMTS always changes service flow parameters on receipt of a DSC-REQ (CM-initiated transaction) or DSC-RSP (CMTS-initiated transaction).

The CMTS also controls the downstream transmit behaviour. The change in downstream transmit behaviour is always coincident with the change in downstream scheduling (i.e., CMTS controls both and changes both simultaneously).

The CM controls the upstream transmit requests, subject to limits contained in the current service flow parameters at the CM. The timing of service flow parameter changes in the CM, and any consequent CM transmit request behaviour changes, is a function of which device initiated the transaction. For a CM-initiated DSC-REQ, the service flow parameters are changed on receipt of the DSC-RSP from the CMTS. For a CMTS-initiated DSC-REQ, the service flow parameters are changed on receipt of the DSC-REQ from the CMTS.

Any service flow can be deactivated with a dynamic service change command by sending a DSC-REQ message, referencing the service flow identifier, and including a null ActiveQosParamSet. However, if a primary service flow of a CM is deactivated, that CM is de-registered and MUST re-register. Therefore, care should be taken before deactivating such service flows. If a service flow that was provisioned during registration is deactivated, the provisioning information for that service flow MUST be maintained until the service flow is reactivated.

A CM MUST have only one DSC transaction outstanding per service flow. If it detects a second transaction initiated by the CMTS, the CM MUST abort the transaction it initiated and allow the CMTS-initiated transaction to complete.

A CMTS MUST have only one DSC transaction outstanding per service flow. If it detects a second transaction initiated by the CM, the CMTS MUST abort the transaction the CM initiated and allow the CMTS-initiated transaction to complete.

NOTE – Currently anticipated applications would probably control a service flow through either the CM or CMTS, and not both. Therefore, the case of a DSC being initiated simultaneously by the CM and CMTS is considered as an exception condition and treated as one.

11.4.3.1 CM-initiated dynamic service change

A CM that needs to change a service flow definition performs the following operations.

The CM informs the CMTS using a dynamic service change request message (DSC-REQ). The CMTS MUST decide if the referenced service flow can support this modification. The CMTS MUST respond with a dynamic service change response (DSC-RSP) indicating acceptance or rejection. The CM reconfigures the service flow if appropriate, and then MUST respond with a dynamic service change acknowledge (DSC-ACK).

CMTS		CM
		Service flow requires modifying
Receive DSC-REQ	← DSC-REQ —	**Send DSC-REQ**
Validate request		
Modify service flow		
Send DSC-RSP	— DSC-RSP →	**Receive DSC-RSP**
		Modify service flow
Receive DSC-ACK	← DSC-ACK —	**Send DSC-ACK**

Figure 11-40 – CM-initiated DSC

11.4.3.2 CMTS-initiated dynamic service change

A CMTS that needs to change a service flow definition performs the following operations.

The CMTS MUST decide if the referenced service flow can support this modification. If so, the CMTS informs the CM using a dynamic service change request message (DSC-REQ). The CM checks that it can support the service change, and MUST respond using a dynamic service change response (DSC-RSP) indicating acceptance or rejection. The CMTS reconfigures the service flow if appropriate, and then MUST respond with a dynamic service change acknowledgment (DSC-ACK).

CMTS		CM
Service flow requires modifying		
Send DSC-REQ	— DSC-REQ →	**Receive DSC-REQ**
		Modify service flow
Receive DSC-RSP	← DSC-RSP —	**Send DSC-RSP**
Modify service flow		
Send DSC-ACK	— DSC-ACK →	**Receive DSC-ACK**

Figure 11-41 – CMTS-initiated DSC

11.4.3.3 Dynamic service change state transition diagrams

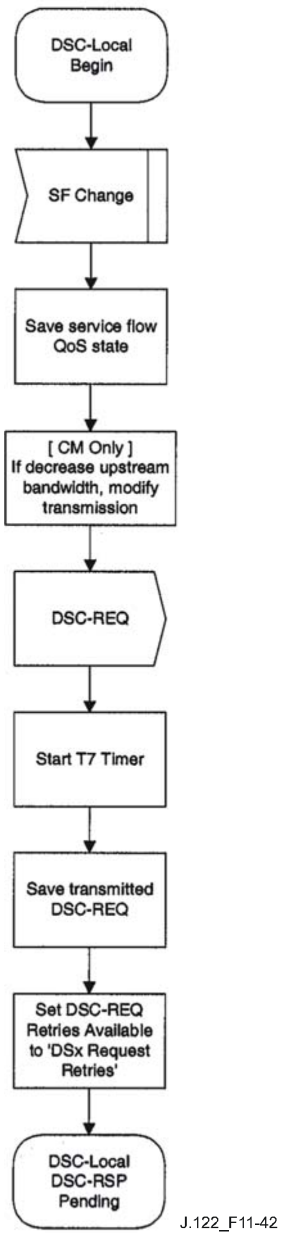

Figure 11-42 – DSC – Locally-initiated transaction begin state flow diagram

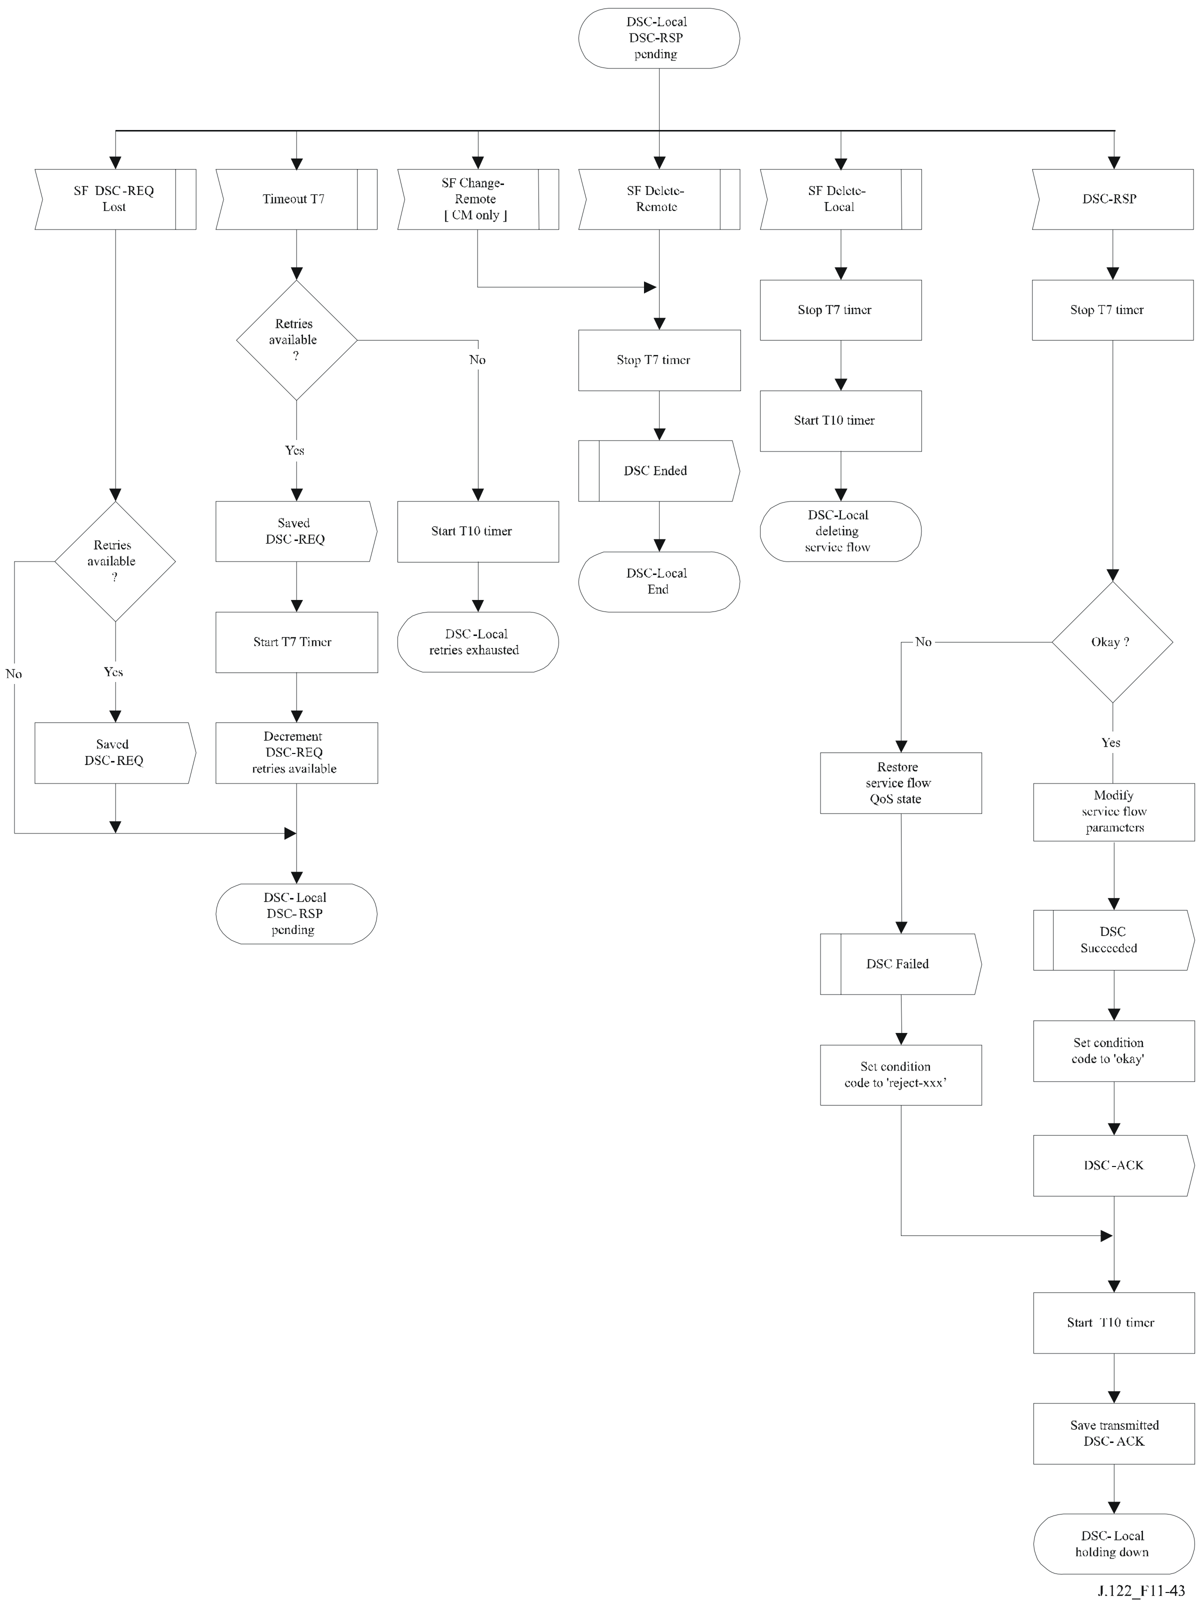

Figure 11-43 – DSC – Locally-initiated transaction DSC-RSP pending state flow diagram

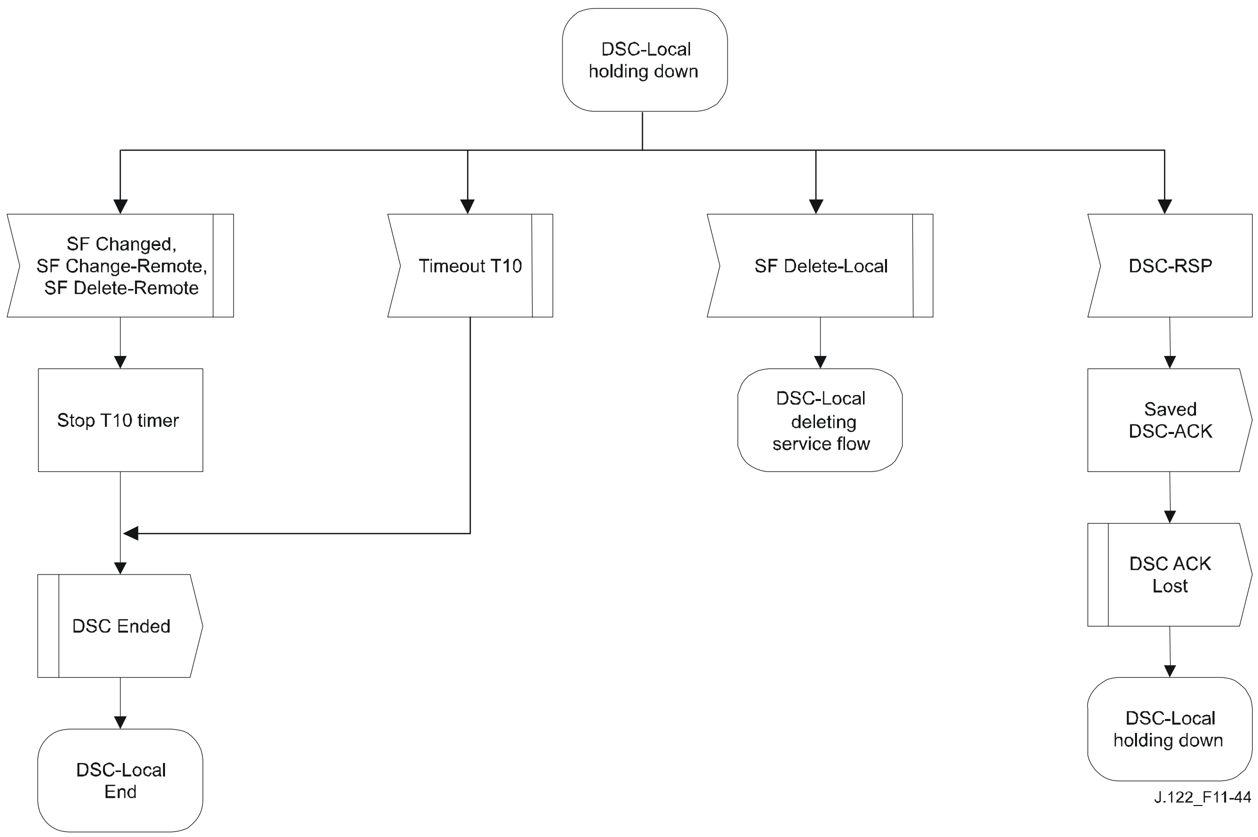

Figure 11-44 – DSC – Locally-initiated transaction holding down state flow diagram

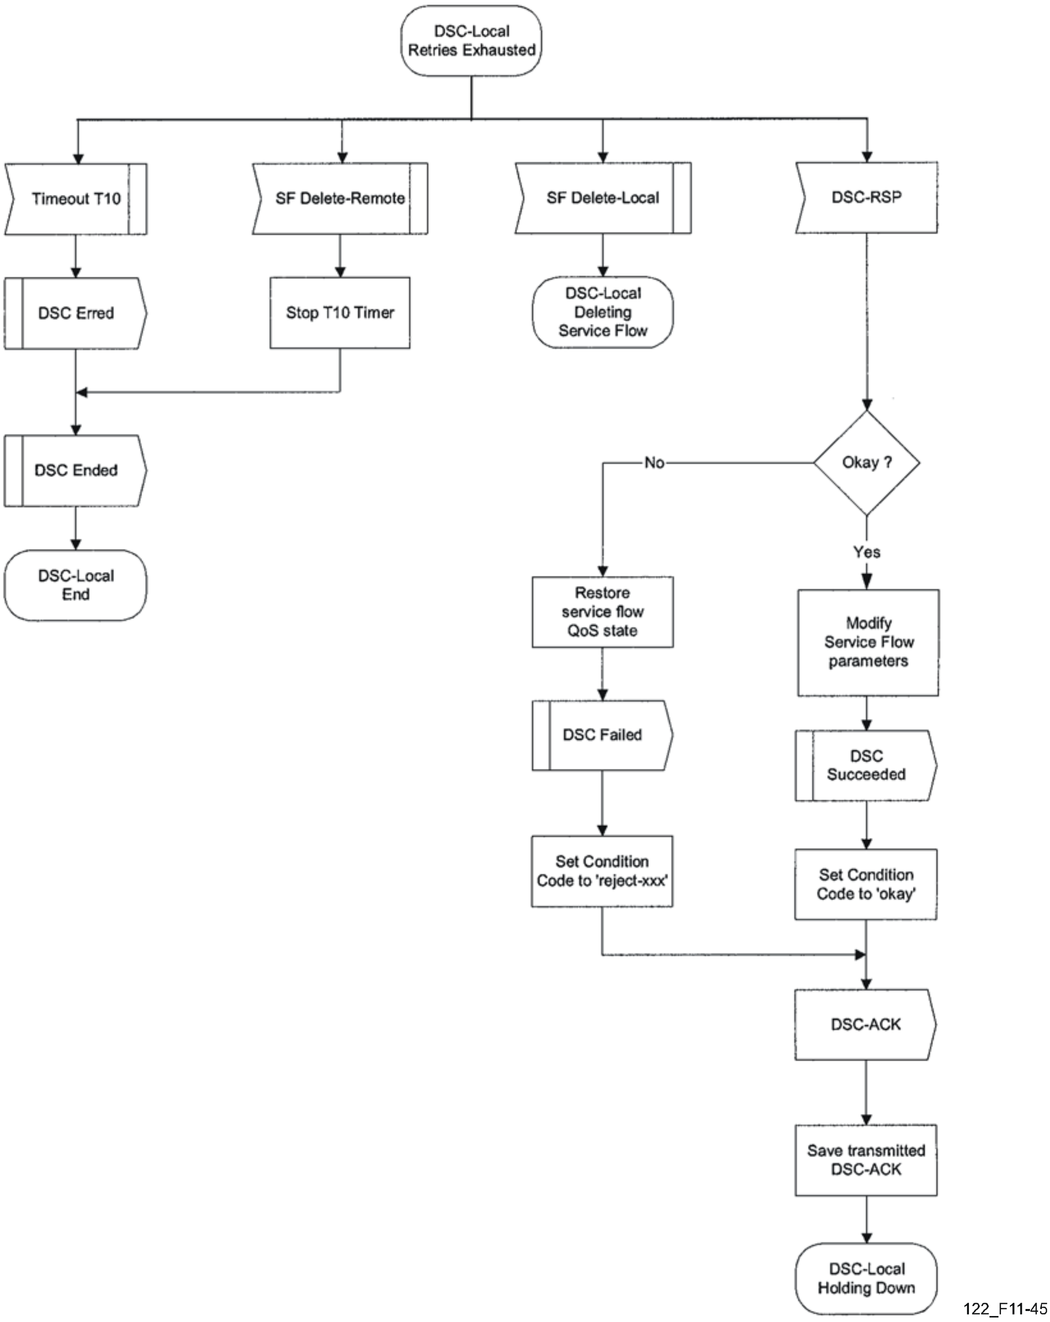

Figure 11-45 – DSC – Locally-initiated transaction retries exhausted state flow diagram

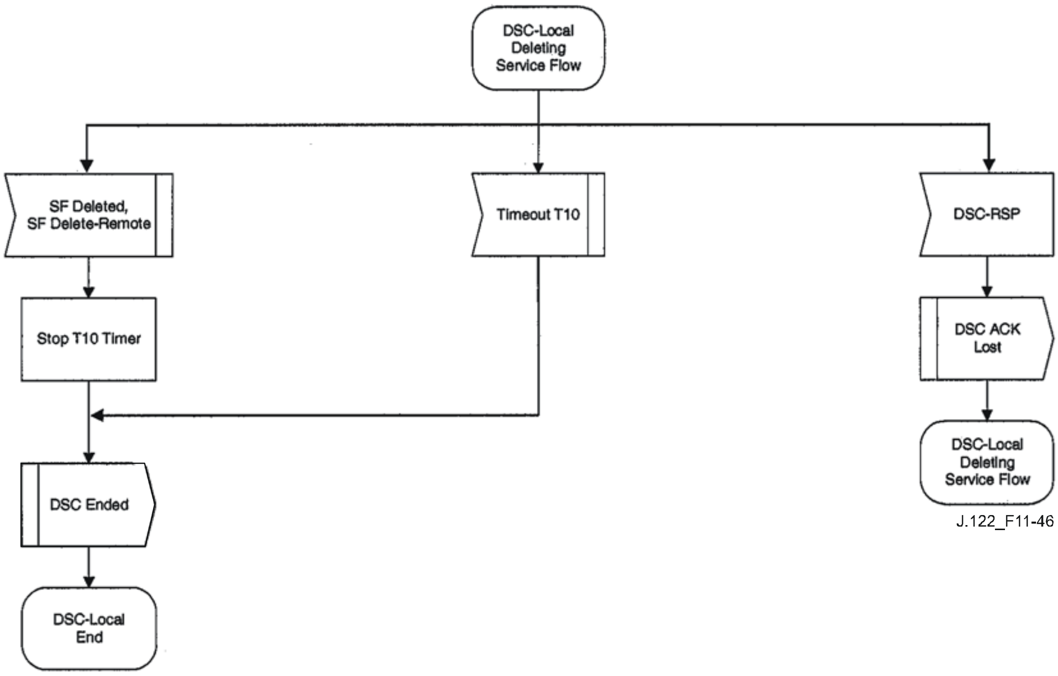

Figure 11-46 – DSC – Locally-initiated transaction deleting service flow state flow diagram

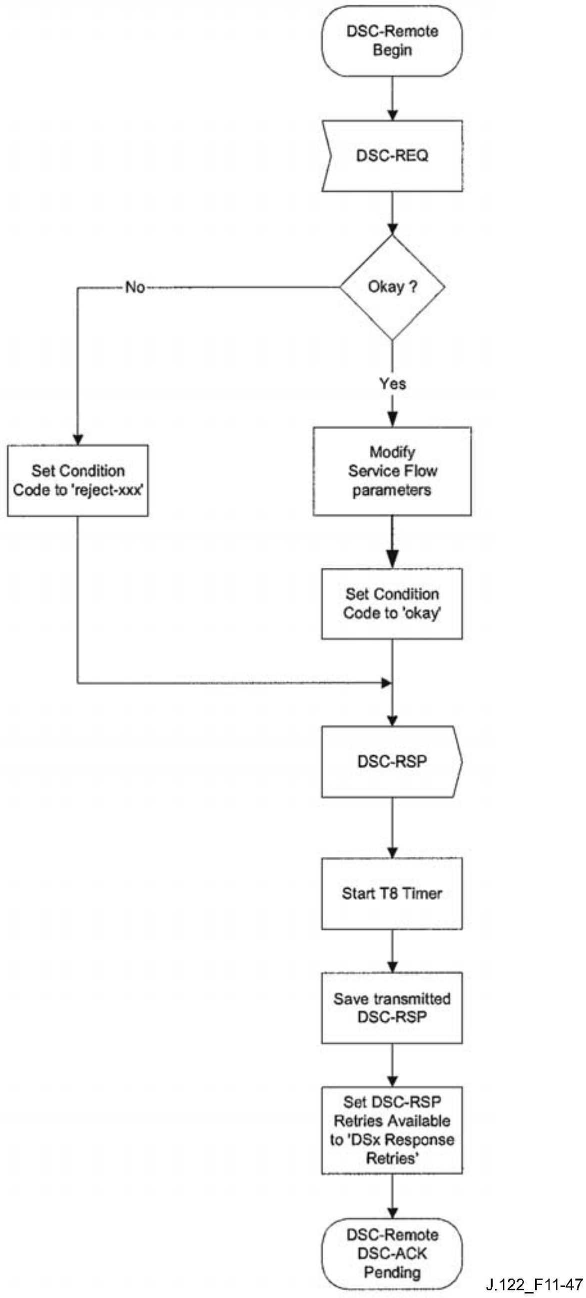

Figure 11-47 – DSC – Remotely-initiated transaction begin state flow diagram

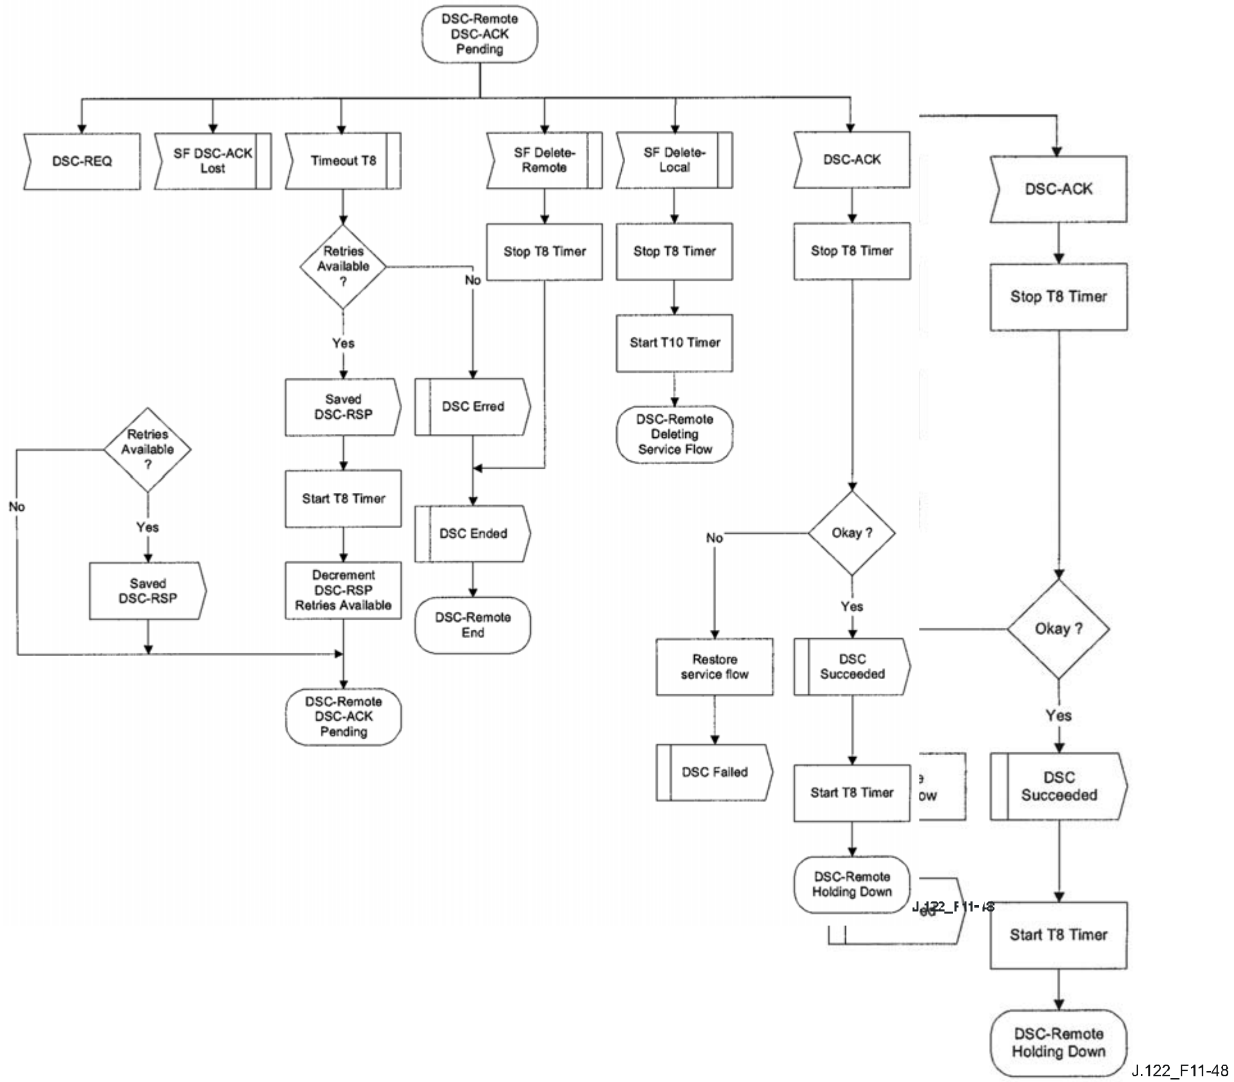

Figure 11-48 – DSC – Remotely-initiated transaction DSC-ACK pending state flow diagram

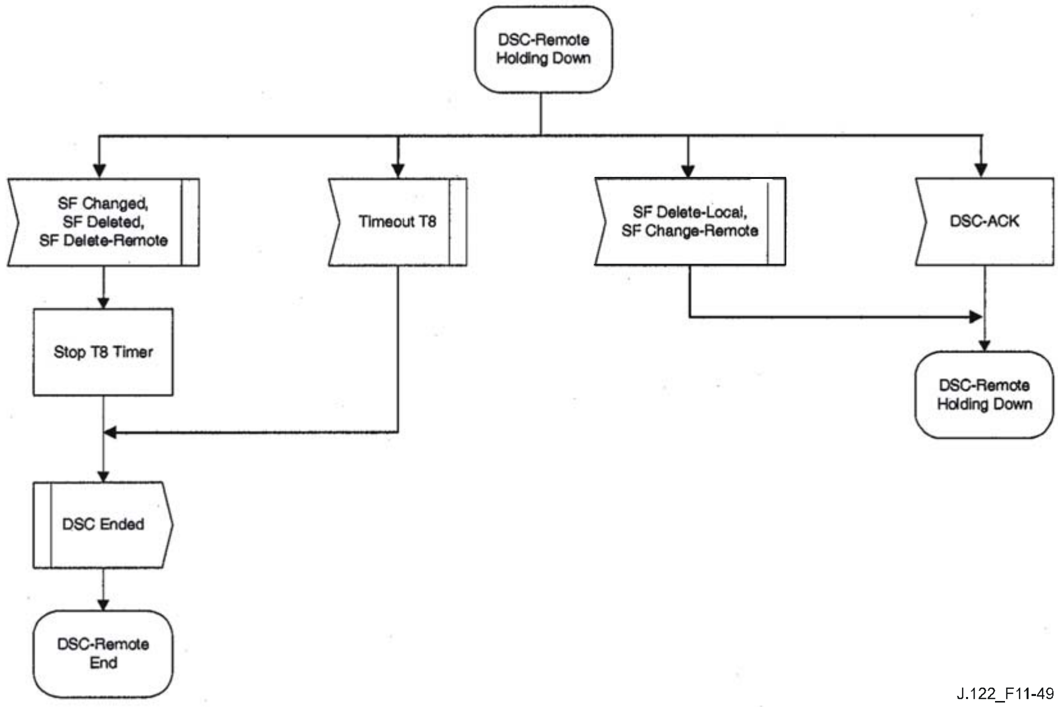

Figure 11-49 – DSC – Remotely-initiated transaction holding down state flow diagram

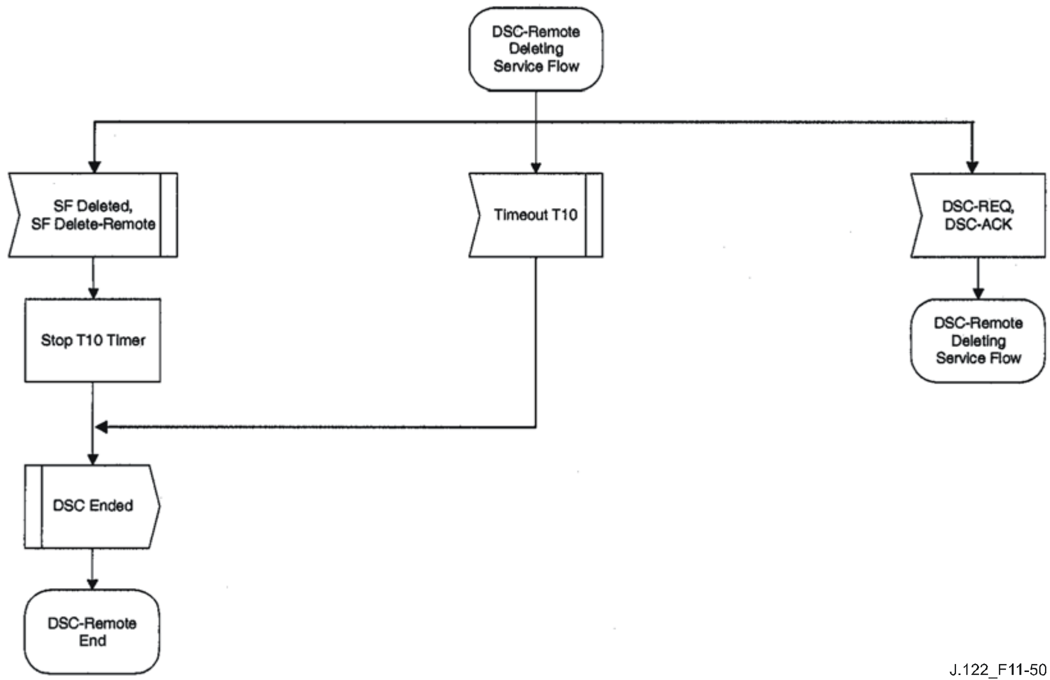

Figure 11-50 – DSC – Remotely-initiated transaction deleting service flow state flow diagram

11.4.4 Dynamic service deletion

Any service flow can be deleted with the dynamic service deletion (DSD) messages. When a service flow is deleted, all resources associated with it are released, including classifiers and PHS. However, if a primary service flow of a CM is deleted, that CM is de-registered and MUST re-register. Also, if a service flow that was provisioned during registration is deleted, the provisioning information for that service flow is lost until the CM re-registers. However, the deletion of a provisioned service flow MUST NOT cause a CM to re-register. Therefore, care should be taken before deleting such service flows.

11.4.4.1 CM-initiated dynamic service deletion

A CM wishing to delete an upstream and/or a downstream service flow generates a delete request to the CMTS using a dynamic service deletion request message (DSD-REQ). The CMTS removes the service flow(s) and generates a response using a dynamic service deletion response message (DSD-RSP). Only one upstream and/or one downstream service flow can be deleted per DSD-Request.

CM		CMTS
Service flow(s) no longer needed		
Delete service flow(s)		
Send DSD-REQ	—— DSD-REQ ——→	**Receive DSD-REQ**
		Verify CM is service flow(s) 'owner'
		Delete service flow(s)
Receive DSD-RSP	←—— DSD-RSP ——	**Send DSD-RSP**

Figure 11-51 – Dynamic service deletion initiated from CM

11.4.4.2 CMTS-initiated dynamic service deletion

A CMTS wishing to delete an upstream and/or a downstream dynamic service flow generates a delete request to the associated CM using a dynamic service deletion request message (DSD-REQ). The CM removes the service flow(s) and generates a response using a dynamic service deletion response message (DSD-RSP). Only one upstream and/or one downstream service flow can be deleted per DSD-Request.

CM		CMTS
		Service flow(s) no longer needed
		Delete service flow(s)
		Determine associated CM for this service flow(s)
Receive DSD-REQ	←—— DSD-REQ ——	**Send DSD-REQ**
Delete service flow(s)		
Send DSD-RSP	—— DSD-RSP ——→	**Receive DSD-RSP**

Figure 11-52 – Dynamic service deletion initiated from CMTS

11.4.4.3 Dynamic service deletion state transition diagrams

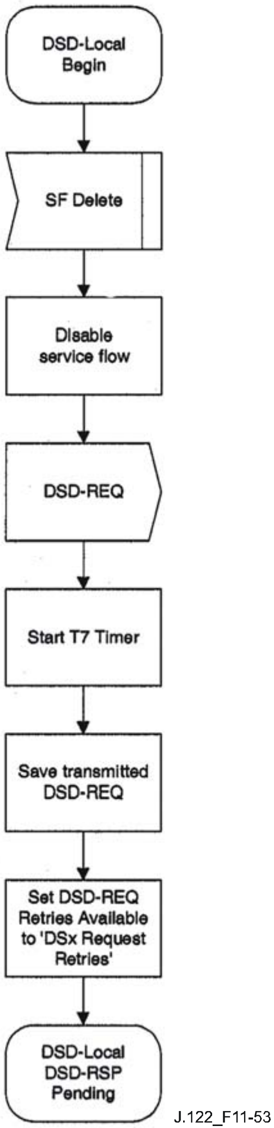

Figure 11-53 – DSD – Locally-initiated transaction begin state flow diagram

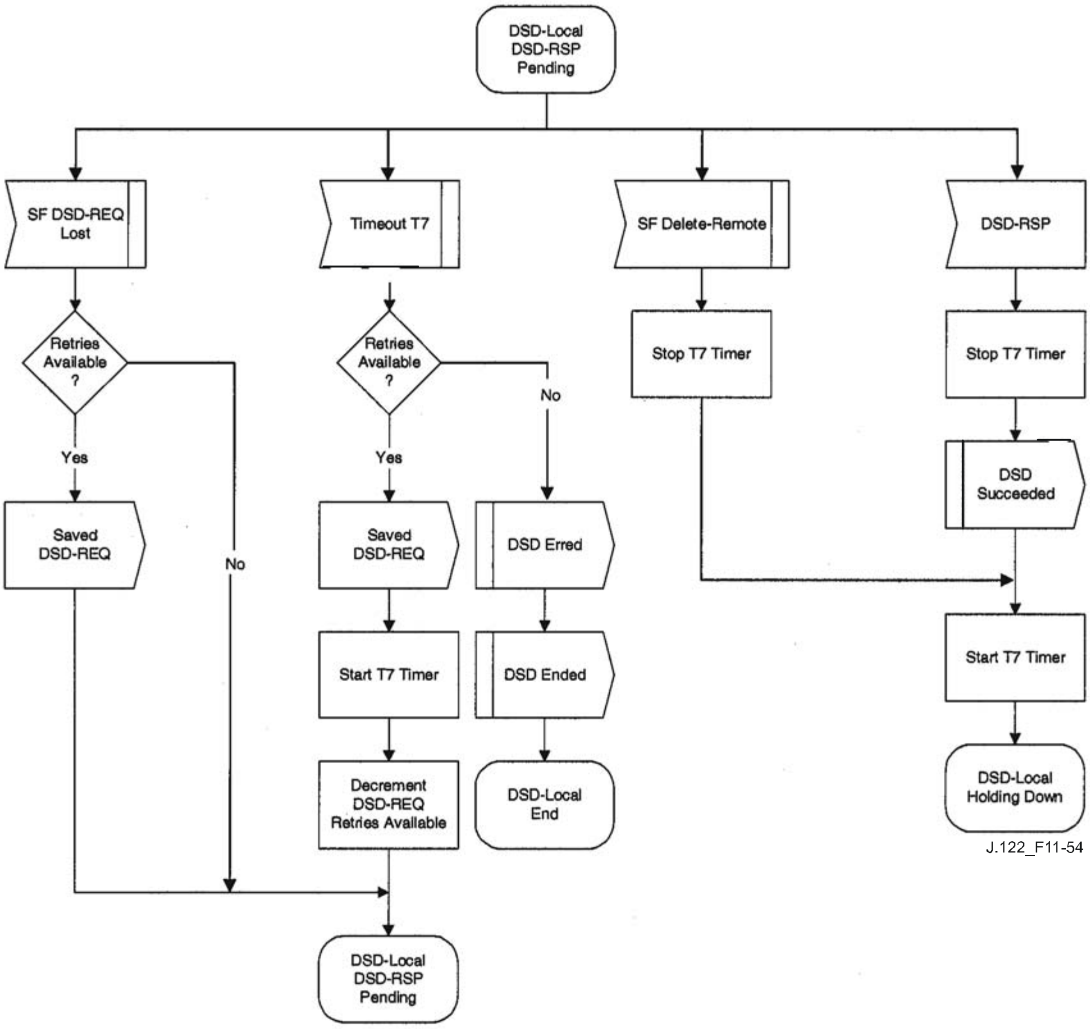

Figure 11-54 – DSD – Locally-initiated transaction DSD-RSP pending state flow diagram

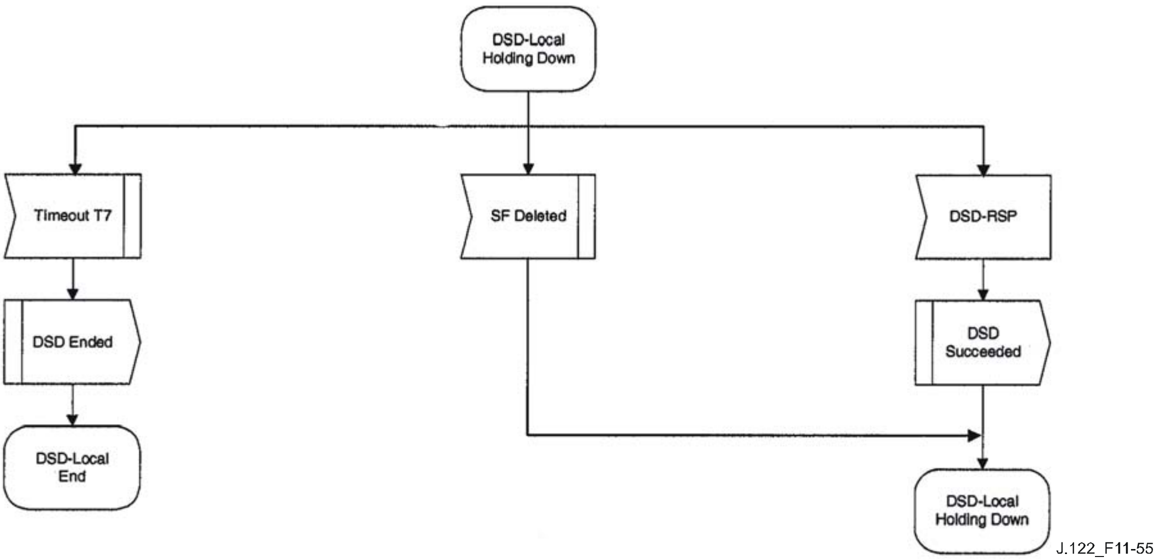

Figure 11-55 – DSD – Locally-initiated transaction holding down state flow diagram

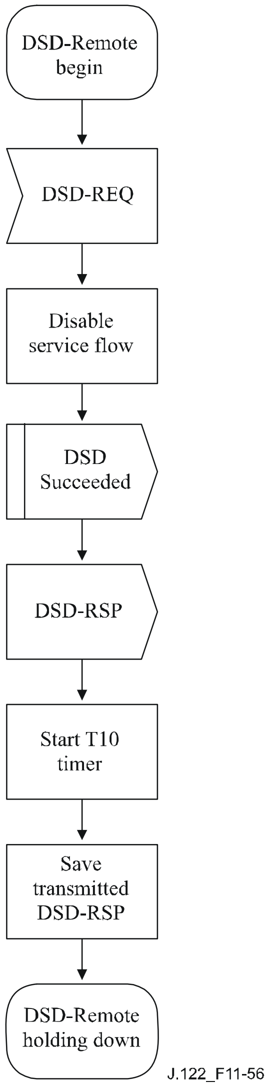

Figure 11-56 – DSD – Remotely-initiated transaction begin state flow diagram

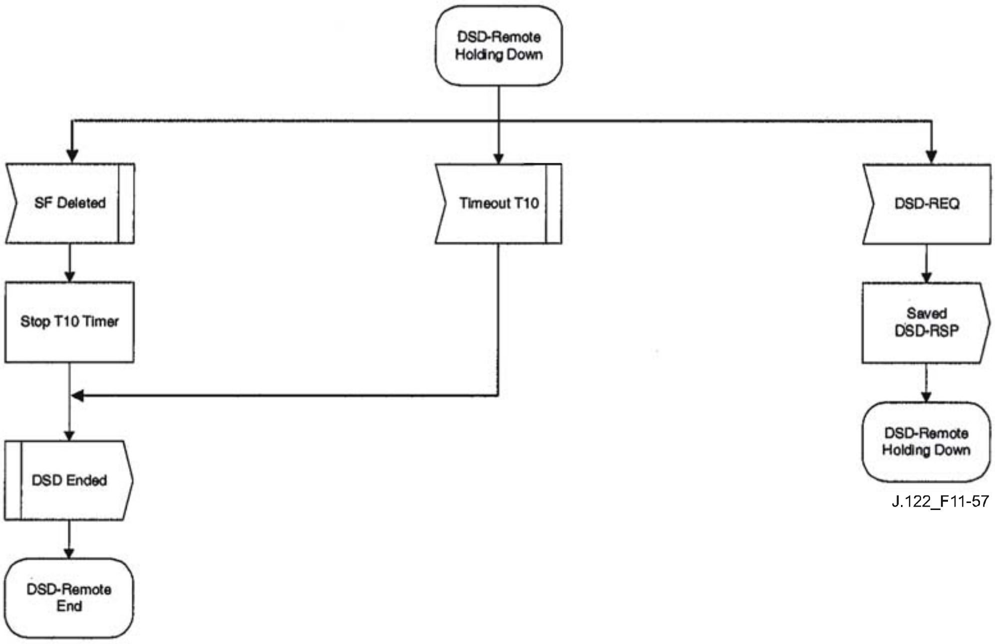

Figure 11-57 – DSD – Remotely-initiated transaction holding down state flow diagram

11.4.5 Dynamically changing downstream and/or upstream channels

11.4.5.1 DCC general operation

At any time after registration, the CMTS MAY use the DCC-REQ to direct the CM to change its downstream and/or upstream channel. The CMTS MUST be capable of performing DCC operations to trigger upstream and/or downstream channel changes within a MAC domain and between MAC domains. Physical layer conditions permitting, the CMTS MUST be capable of executing such channel changes using all initialization techniques (see clause 11.4.5.1.2). This may be done for load balancing (as described in clause 11.4.5.6), noise avoidance or other reasons that are beyond the scope of this Recommendation. In addition, the CMTS MUST support DCC operations triggered via SNMP (see [DOCS5]). Figures 11-58 through 11-61 show the procedure that MUST be followed by the CMTS when performing a dynamic channel change. Figures 11-62 through 11-65 show the corresponding procedure that MUST be followed by a CM when performing a dynamic channel change.

The DCC command can be used to change only the upstream frequency, only the downstream frequency, or both the upstream and downstream frequencies. When only the upstream or only the downstream frequency is changed, the change is typically within a MAC domain. When both the upstream and downstream frequencies are changed, the change may be within a MAC domain, or between MAC domains.

The upstream channel ID SHOULD be unique between the old and new channels. In this context, the old channel refers to the channel that the CM was on before the jump, and the new channel refers to the channel that the CM is on after the jump.

Upon synchronizing with the new upstream and/or downstream channel, the CM MUST use the technique specified in the DCC-REQ initialization technique TLV, if present, to determine if it should perform re-initialization, only ranging, or neither. If this TLV is not present in DCC-REQ, the CM MUST re-initialize its MAC on the new channel assignment (refer to clause 11.2). If the CM has been instructed to re-initialize, then the CMTS MUST NOT wait for a DCC-RSP to occur on the new channel.

If the CM is being moved within a MAC domain, a re-initialization may not be required. If the CM is being moved between MAC domains, a re-initialization may be required. Re-initializing, if requested, is done with the new upstream and downstream channel assignments. It includes obtaining upstream parameters, establishing IP connectivity, establishing time of day, transferring operational parameters, registering, and initializing baseline privacy. If re-initialization is performed, the CM MUST NOT send a DCC-RSP on the new channel.

The decision to re-range is based upon the CMTS's knowledge of any path diversity that may exist between the old and new channels, or if any of the fundamental parameters of the upstream or downstream channel, such as modulation rate, modulation type, or mini-slot size, have changed.

When DCC-REQ does not involve re-initialization or re-ranging, the design goal of the CM will typically be to minimize the disruption of traffic to the end user. To achieve this goal, a CM MAY choose to continue to use QoS resources (such as bandwidth grants) on its current channel after receiving a DCC-REQ and before actually executing the channel change. The CM might also need this time to flush internal queues or reset state machines prior to changing channels.

The CM MAY continue to use QoS resources on the old channel, including the transmission and reception of packets, after sending a DCC-RSP (depart) message and prior to the actual jump. The CM MAY use QoS resources on the new channel, including the transmission and reception of packets, after the jump and prior to sending a DCC-RSP (arrive) message. The CMTS MUST NOT use the DCC-RSP (depart) message to remove QoS resources on the old channel. The CMTS MUST NOT wait for a DCC-RSP (arrive) message on the new channel before allowing QoS resources to be used. This provision is to allow the unsolicited grant service to be used on the old and new channel with a minimum amount of disruption when changing channels.

The CMTS MUST hold the QoS resources on the old channel until a time of T13 has passed after the last DCC-REQ that was sent, or until it can internally confirm the presence of the CM on the new channel assignment. The CM MUST execute the departure from the old channel before the expiry of T13. The CM MAY continue to use QoS resources on the old channel after responding with DCC-RSP and before the expiry of T13.

If the CM is commanded to perform initial or station maintenance or to use the channel directly, the destination CMTS MUST hold the QoS resources on the new channel until a time of T15 has passed after the last DCC-REQ was sent if the CM has not yet been detected on the new channel. If the CM is commanded to re-initialize the MAC, then QoS resources are not reserved on the destination CMTS, and T15 does not apply. If, in the process of a dynamic channel change with a non-zero initialization technique, the CMTS detects that the CM has re-initialized the MAC before completing the channel change, the CMTS MAY de-allocate the resources that were previously allocated to the modem on the new channel before the expiration of T15.

The T15 timer represents the maximum time period for the CM to complete the move to the destination CMTS, and is based on the TLV encodings (i.e., initialization technique TLV) included in the DCC-REQ message and the local configuration of the destination CMTS.

The destination CMTS SHOULD calculate and limit T15 based on internal policy according to the guidelines in clause 11.4.5.1.1.

If the initialization technique of initial ranging is utilized and if the CM arrives after T15 has passed, attempting to use resources on the new channel that have been removed (ranging or requesting bandwidth on a SID that has been deleted), the CMTS MUST send a ranging abort to the CM in order to cause the DCC transaction to fail, ensuring that the CM and CMTS states remain in sync.

When a CM is moved between downstream channels on different IP subnets using initialization techniques other than re-initialize the MAC, a network connectivity issue may occur because no DHCP process is indicated as part of the DCC operation. The CM MAY implement a vendor-specific feature to deal with this situation. The CMTS SHOULD take this issue into account when sending a DCC-REQ and SHOULD direct the CM to use the appropriate initialization technique TLV to ensure no IP connectivity loss as a result of DCC.

Once the CM changes channels, all previous outstanding bandwidth requests made via the request IE or request/data IE are invalidated, and the CM MUST re-request bandwidth on the new channel. In the case of unsolicited grant service in the upstream, the grants are implicit with the QoS reservations, and do not need to be re-requested.

11.4.5.1.1 Derivation of T15 timer

The maximum value noted for the T15 timer denotes the maximum amount of time that the CMTS should reserve resources on the new channel. This value is not to be used to represent acceptable performance.

The equation below describes the method for calculating the value of T15.

$$T15 = \mathit{CmJumpTime} + \mathit{CmtsRxRngReq}$$

Each of the variables in the equation calculating the T15 timer is explained in the table below.

Variable	Explanation and value
CmJumpTime	This is the CM's indication to the CMTS of when it intends to start the jump and how long it will take to jump. For a downstream change, it includes the time for the CM to synchronize to the downstream parameters on the destination channel, such as QAM symbol timing, FEC framing and MPEG framing. It incorporates CM housecleaning on the old channel. It also incorporates one T11 period for the CM to process and receive the DCC-REQ message. This optional value is computed by the CM and returned in DCC-RSP (depart). If the CM does not specify the jump time TLVs, then the destination CMTS assumes that the value is 1.3 seconds. This recognizes the fact that the CM may continue to use the old channel until the expiry of the T13 timer. If the CM specifies the jump time TLVs, then the destination CMTS uses the specified value.
CmtsRxRngReq	This variable represents the time for the CM to receive and use a ranging opportunity, and for the CMTS to receive and process the RNG-REQ. For the initialization technique of use directly, this value is two times the CMTS time period between unicast ranging opportunities plus 20-40 milliseconds for MAP and RNG-REQ transmission time and CMTS RNG-REQ processing time. For the initialization technique of station maintenance, this value is two times the CMTS time period between unicast ranging opportunities plus 20-40 milliseconds for MAP and RNG-REQ transmission time and CMTS RNG-REQ processing time. For the initialization technique of initial maintenance, this value is 30 seconds. Because the variables involved in initial maintenance are not strictly under the control of the CMTS, the computation of this factor is uncertain.

The minimum value of the T15 timer is four seconds; this was derived by quadrupling the value of the T13 timer. The maximum value of the T15 timer is 35 seconds.

11.4.5.1.2 Initialization technique

There are many factors that drive the selection of an initialization technique when commanding a dynamic channel change. While it is desirable to provide the minimum of interruption to existing

QoS services such as voice over IP or streaming video sessions, a CM will not be able to successfully execute a channel change given an initialization technique that is unsuitable for the cable plant conditions. In some cases, a CM will impair the new channel given an unsuitable initialization technique. For instance, consider the use of initialization technique 4 (use the new channel(s) directly without re-initializing or performing initial or station maintenance) when there is a significant difference in propagation delay between the old and new channel. Not only will the CM be unable to communicate with the CMTS after the channel change, but its transmissions may also interfere with the transmissions of other CMs using the channel.

Careful consideration must be given to the selection of an initialization technique. Some restrictions are listed below. This list is not exhaustive, but is intended to prevent the use of a particular initialization technique under conditions where its use could prevent the CM from successfully executing the channel change. Packets may be dropped under some conditions during channel changes; applications that are running over the DOCSIS path should be able to cope with the loss of packets that may occur during the time that the CM changes channels.

If a CM performs a channel change without performing a re-initialization (as defined in clause 8.3.20.1.3), all the configuration variables of the CM MUST remain constant (BPI keys, IP address, classifiers, PHS Rules, etc.), with the exception of the configuration variables which are explicitly changed by the DCC-REQ message encodings (as defined in clauses 8.3.20.1.1 through 8.3.20.1.7). The CM will not be aware of any configuration changes other than the ones that have been supplied in the DCC command, so consistency in provisioning between the old and new channels is important. Note that regardless of the initialization technique, the CPE will not be aware of any configuration changes and will continue to use its existing IP address.

11.4.5.1.2.1 Initialization technique zero

The use of initialization technique 0 (re-initialize the MAC) results in the longest interruption of service. The CMTS MUST signal the use of this technique when QoS resources will not be reserved on the new channel(s).

11.4.5.1.2.2 Initialization technique one

The use of initialization technique 1 (broadcast initial ranging) may also result in a lengthy interruption of service. However, this interruption of service is mitigated by the reservation of QoS resources on the new channel(s). The service interruption can be further reduced if the CMTS supplies downstream parameter sub-TLVs and the UCD substitution TLV in the DCC-REQ in addition to providing more frequent initial ranging opportunities on the new channel.

11.4.5.1.2.3 Initialization technique two

The use of initialization technique 2 (unicast initial ranging) offers the possibility of only a slight interruption of service.

In order to use this technique, the CMTS MUST:

- Synchronize timestamps (and downstream symbol clocks for S-CDMA support) across the downstream channels involved, and specify SYNC substitution sub-TLV with a value of 1 if the downstream channel is changing.
- Include the UCD substitution in the DCC message if the upstream channel is changing.

However, the CMTS MUST NOT use this technique if:

- The DCC-REQ message requires the CM to switch between S-CDMA and TDMA.
- Propagation delay differences between the old and new channels will cause the CM burst timing to exceed the ranging accuracy requirements of clause 6.2.19.1.
- Attenuation or frequency response differences between the old and new upstream channels will cause the received power at the CMTS to be outside the limits of reliable reception.

11.4.5.1.2.4 Initialization technique three

The use of initialization technique 3 (initial ranging or periodic ranging) offers the possibility of only a slight interruption of service. This value might be used when there is uncertainty when the CM may execute the DCC command and thus a chance that it might miss station maintenance slots. However, the CMTS MUST NOT use this technique if the conditions for using techniques 1 and 2 are not completely satisfied.

11.4.5.1.2.5 Initialization technique four

The use of initialization technique 4 (use the new channel without re-initialization or ranging) results in the least interruption of service.

In order to use this technique, the CMTS MUST:

- Synchronize timestamps (and downstream symbol clocks for S-CDMA support) across the downstream channels involved and specify SYNC substitution sub-TLV with a value of 1 if the downstream channel is changing.
- Include the UCD substitution in the DCC message if the upstream channel is changing.

However, the CMTS MUST NOT use this technique if:

- The CM is operating in S-CDMA mode and any of the following parameters are changing:
  – Modulation rate;
  – S-CDMA mode enable;
  – S-CDMA US ratio numerator 'M';
  – S-CDMA US ratio numerator 'N';
  – Downstream channel.
- The DCC-REQ message requires the CM to switch between S-CDMA and TDMA.
- Attenuation or frequency response differences between the old and new upstream channels will cause the received power at the CMTS to change by more than 6 dB.
- Propagation delay differences between the old and new channels will cause the CM burst timing to exceed the ranging accuracy requirements of clause 6.2.19.1.
- Micro-reflections on the new upstream channel will result in an unacceptable PER (greater than 1%) with the pre-equalizer coefficients set to the initial setting (refer to clause 6.2.15).

The CMTS SHOULD NOT use initialization technique 4 when the new or old channel is S-CDMA and the modulation rate is changing.

11.4.5.2 DCC exception conditions

If a CM issues a DSA-REQ or DSC-REQ for more resources, and the CMTS needs to do a DCC to obtain those resources, the CMTS will reject the DSA or DSC command without allocating any resources to the CM. The CMTS includes a confirmation code of "reject-temporary-DCC" (refer to clause C.1.3.1) in the DSC-RSP message to indicate that the new resources will not be available until a DCC is received. The CMTS will then follow the DSA or DSC transaction with a DCC transaction.

After the CM jumps to a new channel and completes the DCC transaction, the CM retries the DSA or DSC command. If the CM has not changed channels after the expiry of T14, as measured from the time that the CM received DSA-RSP or DSC-RSP from the CMTS, then the CM MAY retry the resource request.

If the CMTS needs to change channels in order to satisfy a resource request other than a CM-initiated DSA or DSC command, then the CMTS should execute the DCC command first, and then issue a DSA or DSC command.

If the provisioning system default is to specify the upstream channel ID, the downstream frequency, and/or a downstream channel list in the configuration file, care should be taken when using DCC, particularly when using initialization technique 0 (re-initialize MAC). If a CMTS does a DCC with re-initialize, the configuration file could cause the CM to come back to the original channel. This would cause an infinite loop.

The CMTS MUST NOT issue a DCC command if the CMTS has previously issued a DSA, or DSC command, and that command is still outstanding. The CMTS MUST NOT issue a DCC command if the CMTS is still waiting for a DSA-ACK or DSC-ACK from a previous CM initiated DSA-REQ or DSC-REQ command.

The CMTS MUST NOT issue a DSA or DSC command if the CMTS has previously issued a DCC command, and that command is still outstanding.

If the CMTS issues a DCC-REQ command and the CM simultaneously issues a DSA-REQ or DSC-REQ, then the CMTS command takes priority. The CMTS responds with a confirmation code of "reject-temporary" (refer to clause C.1.3.1). The CM proceeds with executing the DCC command.

If the CM is unable to achieve communications with a CMTS on the new channel(s), it MUST return to the previous channel(s) and re-initialize its MAC. The previous channel assignment represents a known good operating point which should speed up the re-initialization process. Also, returning to the previous channel provides a more robust operational environment for the CMTS to find a CM that fails to connect on the new channel(s).

If the CMTS sends a DCC-REQ and does not receive a DCC-RSP within time T11, it MUST retransmit the DCC-REQ up to a maximum of "DCC-REQ Retries" (Annex B) before declaring the transaction a failure. Note that if the DCC-RSP was lost in transit and the CMTS retries the DCC-REQ, the CM may have already changed channels.

If the CM sends a DCC-RSP on the new channel and does not receive a DCC-ACK from the CMTS within time T12, it MUST retry the DCC-RSP up to a maximum of "DCC-RSP retries" (Annex B).

If the CM receives a DCC-REQ with the upstream channel ID TLV, if present, equal to the current upstream channel ID, and the downstream frequency TLV, if present, is equal to the current downstream frequency, then the CM MUST consider the DCC-REQ as a redundant command. The remaining DCC-REQ TLV parameters MUST NOT be executed, and the CM MUST return a DCC-RSP, with a confirmation code of "reject-already-there", to the CMTS (refer to clause C.4.1).

If a DOCS 2.0-enabled CM is moved from a type 1 channel to a type 2 channel via a DCC using an initialization technique other than re-initialize the MAC, the CM MUST operate in 2.0 mode on the destination channel, basing its requests on IUCs 9 and 10 instead of IUCs 5 and 6. If a CM in which DOCS 2.0 is disabled in registration is moved from a type 1 channel to a type 2 channel via a DCC with initialization technique other than re-initialize the MAC, the CM MUST continue to base its requests on the destination channel on IUCs 5 and 6. A CM in which DOCSIS 2.0 is disabled in registration MUST consider a type 3 channel to be invalid for any DCC using an initialization technique other than re-initialize the MAC.

If the CM is moved via a DCC using the initialization technique of re-initialize the MAC, the previous 2.0 enable setting is not applicable. If DOCS 2.0 was previously disabled and the target upstream channel is type 2 or type 3, the 2.0 disable setting is discarded and the CM MUST use the target upstream channel; however, after re-initialization, the CM MUST operate in the 2.0 mode defined in the config file acquired in the re-initialization process.

11.4.5.3 DCC state transition diagrams

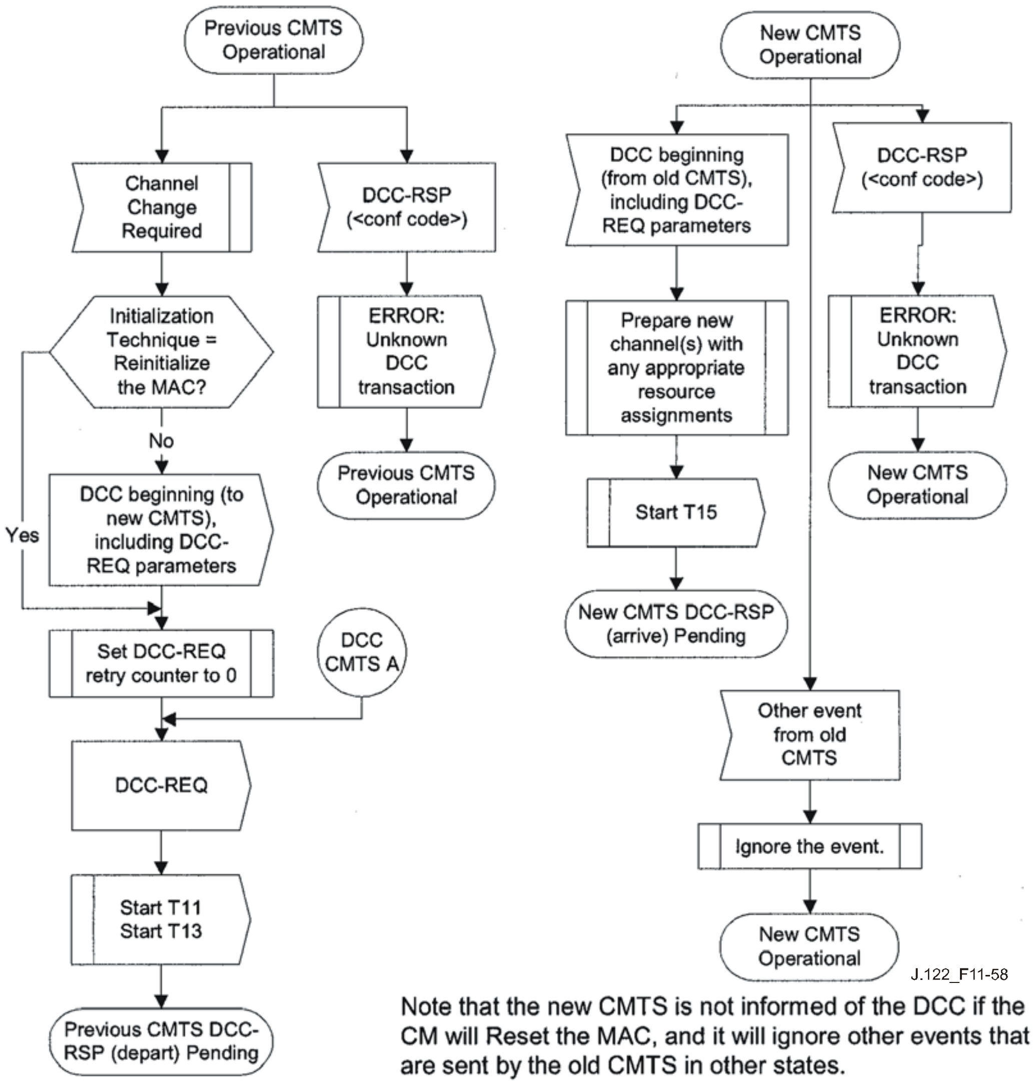

Figure 11-58 – Dynamically changing channels: CMTS view, part 1

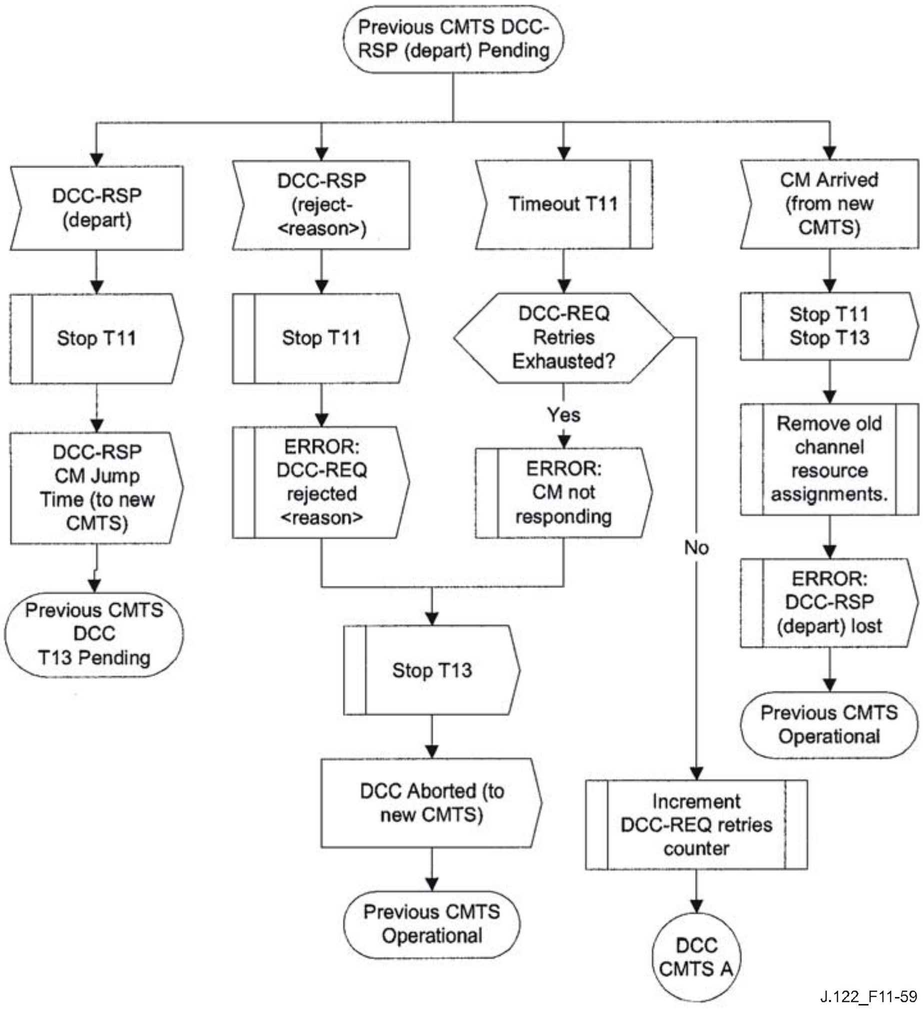

Figure 11-59 – Dynamically changing channels: CMTS view, part 2

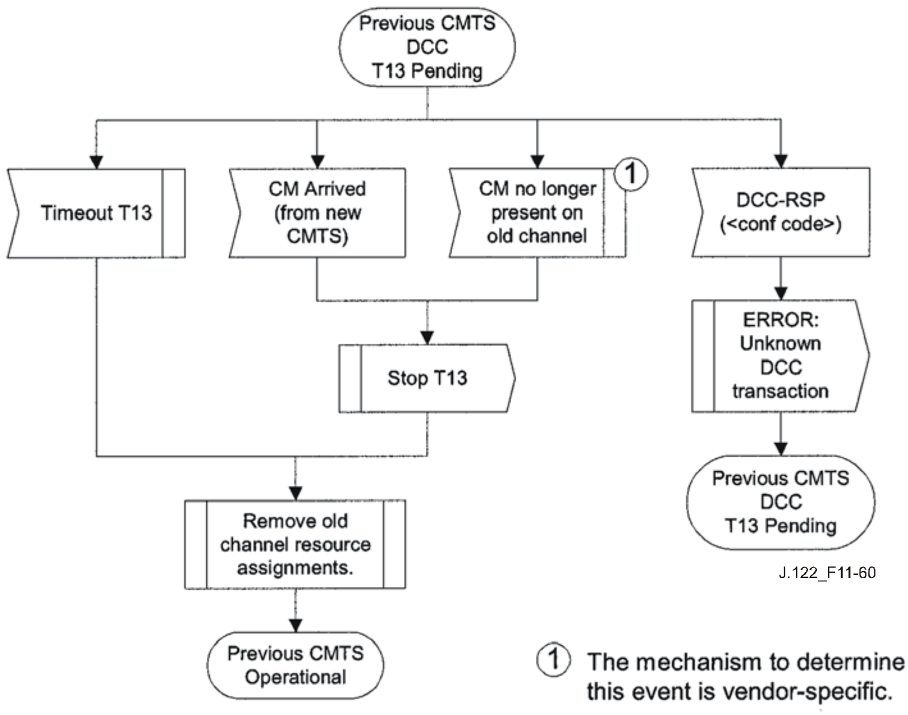

Figure 11-60 – Dynamically changing channels: CMTS view, part 3

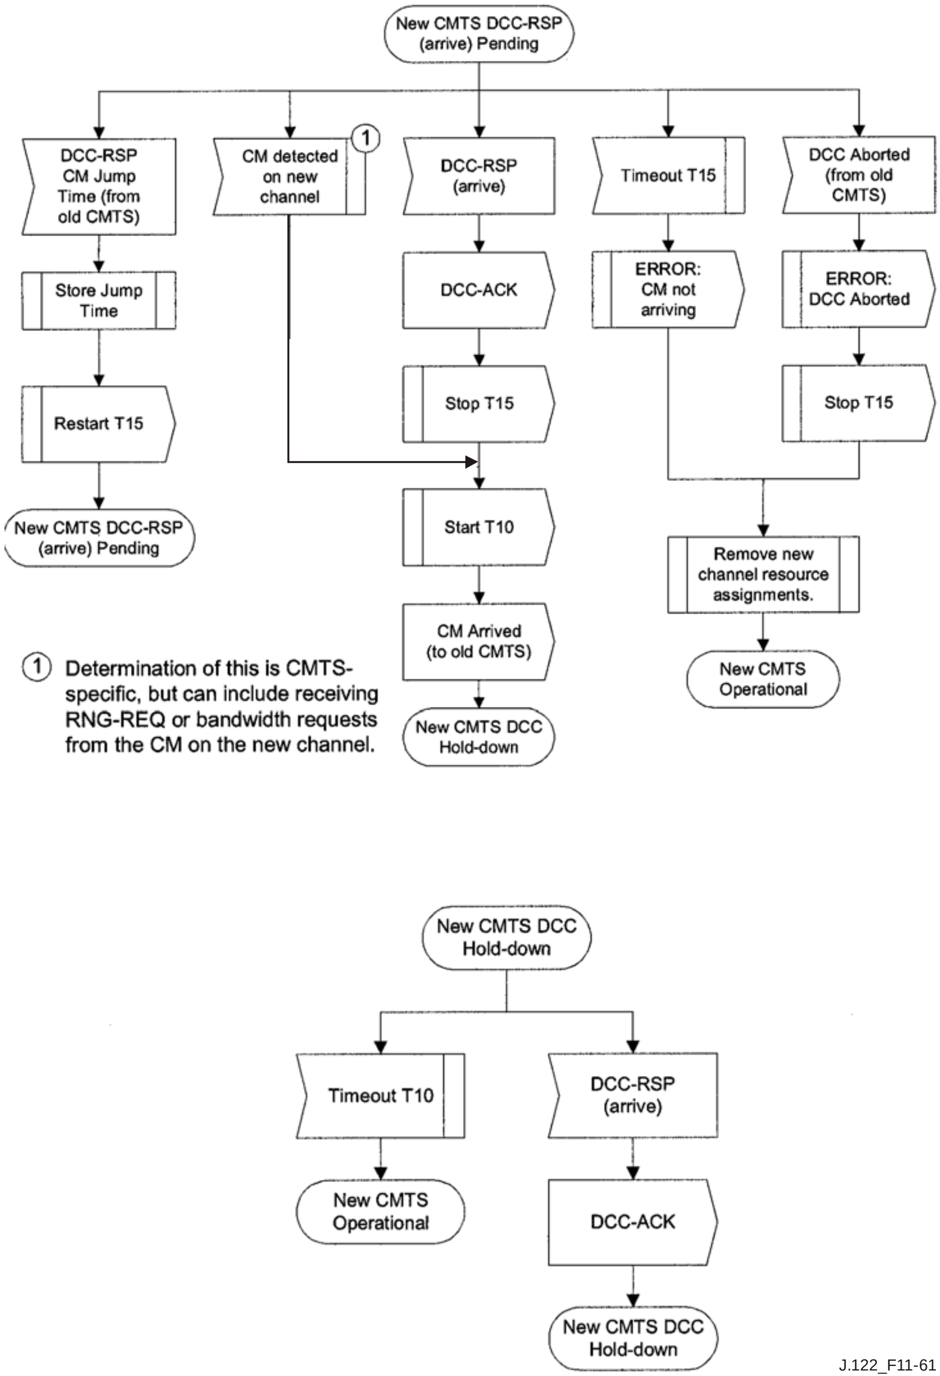

Figure 11-61 – Dynamically changing channels: CMTS view, part 4

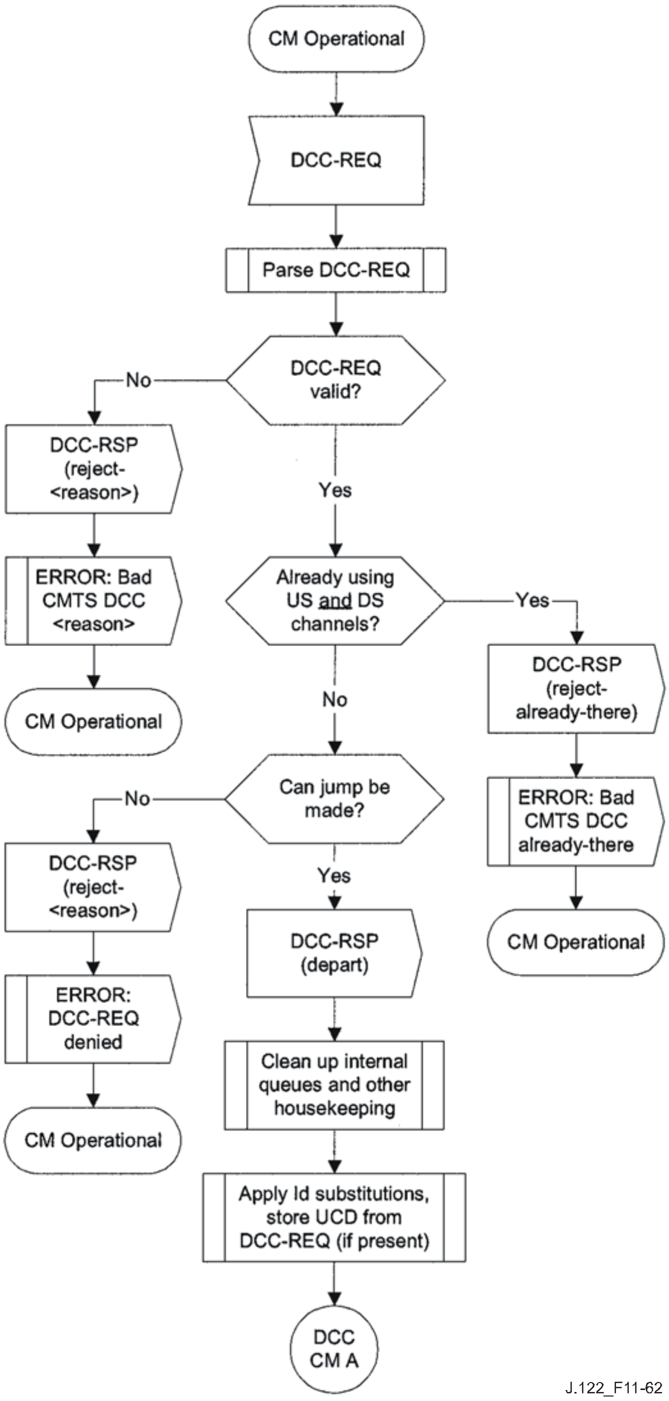

Figure 11-62 – Dynamically changing channels: CM view, part 1

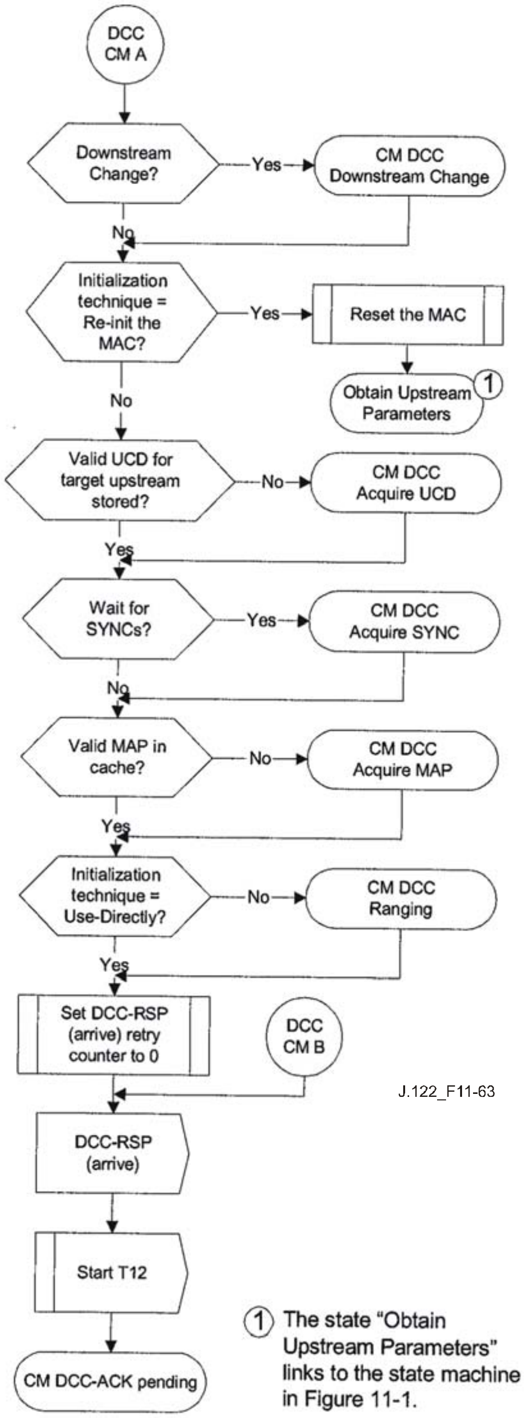

Figure 11-63 – Dynamically changing channels: CM view, part 2

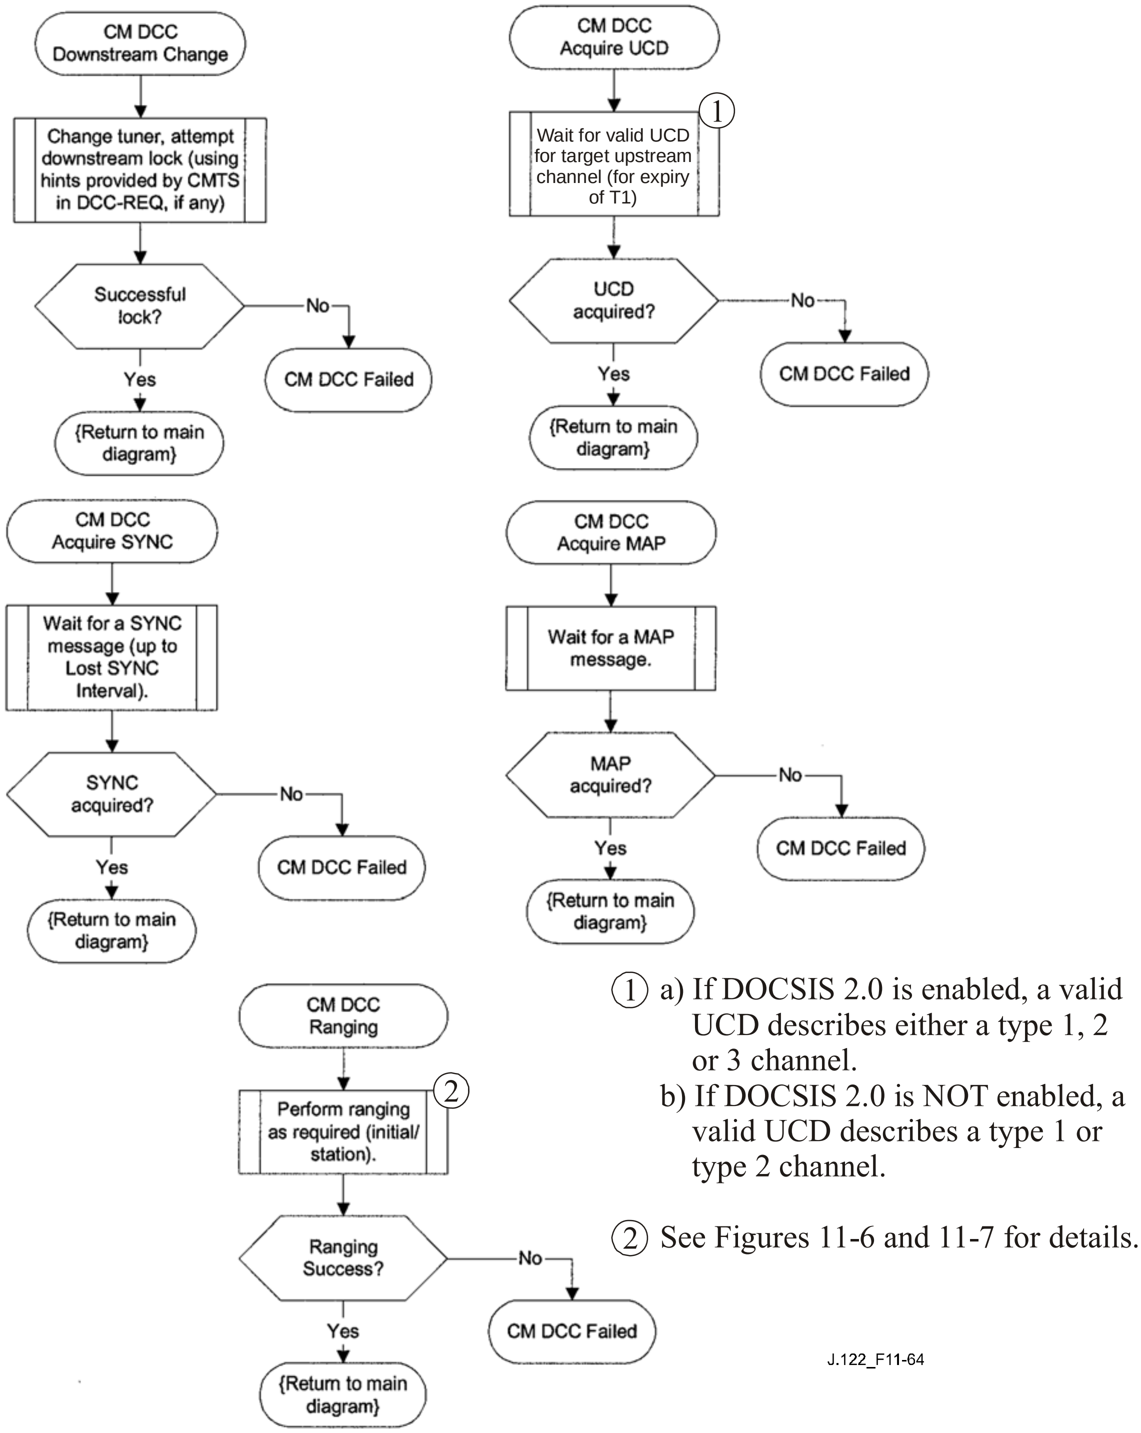

Figure 11-64 – Dynamically changing channels – CM view, part 3

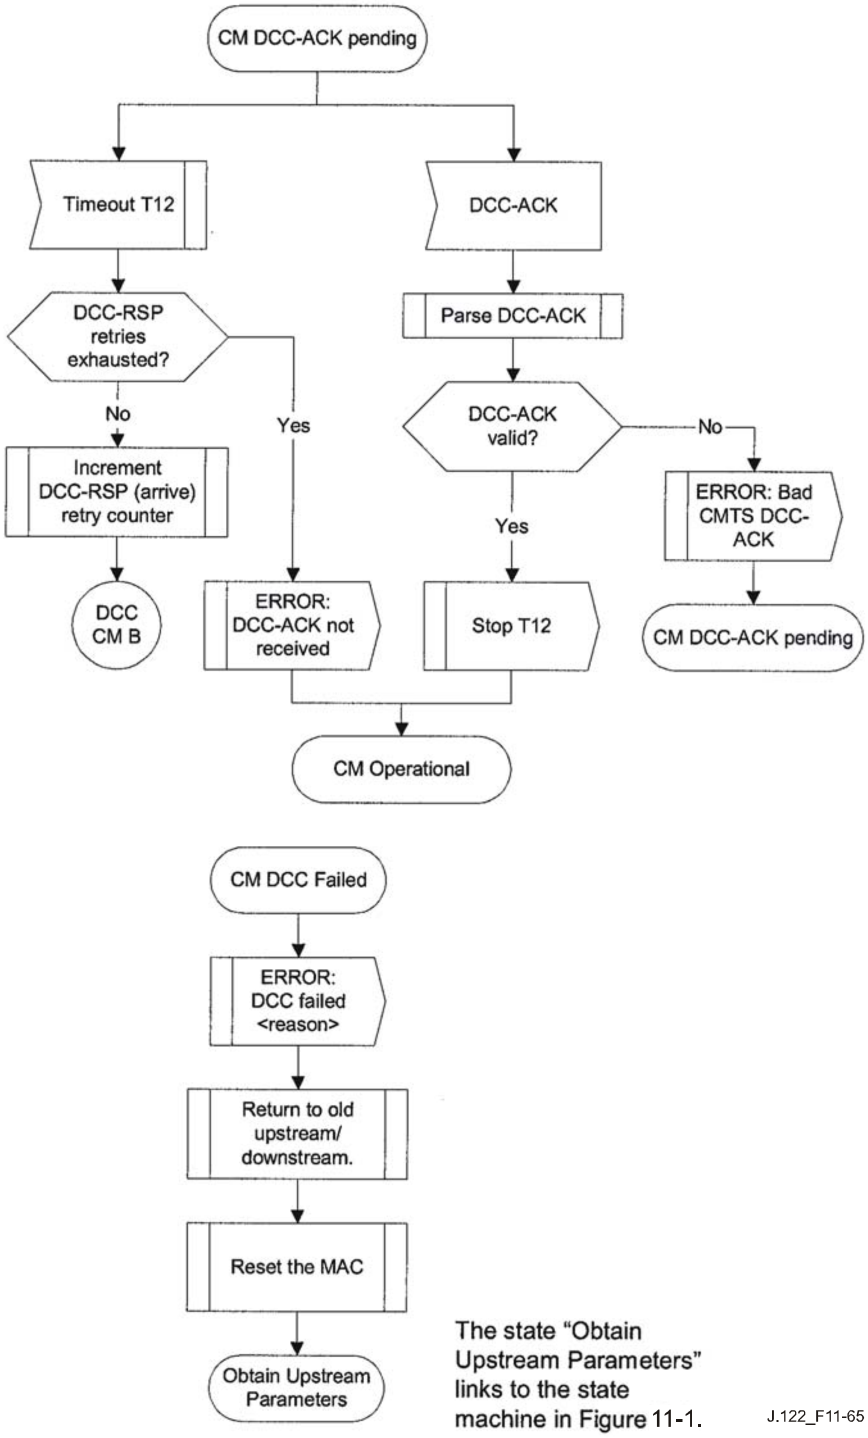

Figure 11-65 – Dynamically changing channels – CM view, part 4

11.4.5.4 DCC performance

The purpose of a DCC is to move the CM to a new upstream and/or downstream channel with little interruption of service. There are many factors that affect the performance of a DCC transaction including CM housecleaning, initialization technique and the number of TLV hints given by the current CMTS in the DCC-REQ message. Each of these factors is individually discussed in Table 11-1.

Table 11-1 – Factors affecting DCC performance

TLV type	Value	Explanation
Initialization technique	Absent or 0 re-initialize MAC	There are no performance requirements in this case. The CM will arrive on the destination CMTS after initialization occurs.
	1 – Broadcast initial ranging	There are low performance expectations in this case because many factors affect the performance, such as collisions and ranging backoff. The CM should arrive on the destination CMTS as quickly as possible.
	2 – Unicast initial ranging	The DCC transaction SHOULD complete within 1.5 seconds after the start of jump.
	3 – Broadcast or unicast initial ranging	The CMTS does not know which ranging technique the CM will utilize. The CM should arrive on the destination CMTS as quickly as possible.
	4 – Use channel directly	The DCC transaction SHOULD complete within one second after the start of jump.
DS parameter		The CMTS SHOULD include the downstream parameter TLVs for station maintenance and use directly initialization techniques that are expected to occur quickly.
UCD substitution		The CMTS MUST set the UCD substitution TLV according to clause 11.4.5.1.2.
SYNC substitution		The CMTS MUST set the SYNC substitution TLV according to clause 11.4.5.1.2.
CM jump time		The length of jump TLV SHOULD be less than one second for downstream channel changes that include the downstream parameter TLVs or for upstream-only channel changes.

The DCC transaction is defined from the perspective of both the CM and the CMTS for the discussion on performance in Table 11-1. From the perspective of the CM, the DCC transaction begins when the CM receives the DCC-REQ message from the CMTS and completes when the CM receives the DCC-ACK message from the CMTS. From the perspective of the CMTS, the DCC transaction begins when the CMTS sends the DCC-REQ message to the CM and completes when the CMTS receives the DCC-RSP (arrive) message from the CM.

When the DCC-REQ does not contain UCD substitution TLVs and/or specifies an initialization technique of initial maintenance, station maintenance or use directly, the destination CMTS SHOULD increase the probability that the CM will arrive quickly by using the CM jump time TLVs specified in the DCC-RSP (depart) to adjust the transmission of UCDs and ranging opportunities such that they coincide with the time when CM has estimated that it will arrive, and SHOULD increase the frequency of UCDs and/or ranging opportunities during this period.

11.4.5.5 Example operation

11.4.5.5.1 Example signalling

Figure 11-66 shows an example of the use of DCC and its relation to the other DOCS MAC messages. In particular, this example describes a scenario where the CM attempts to allocate new resources with a DSA message. The CMTS temporarily rejects the request, tells the CM to change channels, and then the CM re-requests the resources. This example (not including all exception conditions) is described below. Refer to clause 11.2 for more detail.

a) An event occurs, such as the CM issuing a DSA-REQ message.

b) The CMTS decides that it needs the CM to change channels in order to service this resource request. The CMTS responds with a DSA-RSP message which includes a confirmation code of "reject-temporary-DCC" (refer to clause C.1.3.1) in the DSC-RSP message to indicate that the new resources are not available until a DCC is received. The CMTS now rejects any further DSA or DSC messages until the DCC command is executed.

c) The CMTS initiates QoS reservations on the new upstream and/or downstream channels. The QoS reservations include the new resource assignment along with all the current resource assignments assigned to the CM. In this example, both the upstream and downstream channels are changed.

d) To facilitate a near-seamless channel change, since the CMTS is not sure exactly when the CM will switch channels, the CMTS duplicates the downstream packet flow on the old and new downstream channels.

e) The CMTS issues a DCC-REQ command to the CM.

f) The CM cleans up its queues and state machines as appropriate, sends a DCC-RSP (depart), and changes channels.

g) If there was a downstream channel change, the CM synchronizes to the QAM symbol timing, synchronizes the FEC framing, and synchronizes with the MPEG framing.

h) If the CM has been instructed to re-initialize, it does so with the new upstream and/or downstream channel assignment. The CM exits from the flow of events described here, and enters the flow of events described in clause 11.2, starting with the recognition of a downstream SYNC message.

i) The CM searches for a UCD message unless it has been supplied with a copy.

j) The CM waits for a downstream SYNC message unless it has been instructed not to wait for one.

k) The CM collects MAP messages unless it already has them available in its cache.

l) The CM performs ranging if required by the initialization technique TLV.

m) The CM resumes normal data transmission with its new resource assignment.

n) The CM sends a DCC-RSP (arrive) message to the CMTS.

o) The CMTS responds with a DCC-ACK.

p) The CMTS removes the QoS reservations from the old channels. If the downstream packet flow was duplicated, the packet duplication would also be removed on the old downstream channel.

q) The CM re-issues its DSA-REQ command.

r) The CMTS reserves the requested resources and responds with a DSA-RSP.

s) The CM finishes with a DSA-ACK.

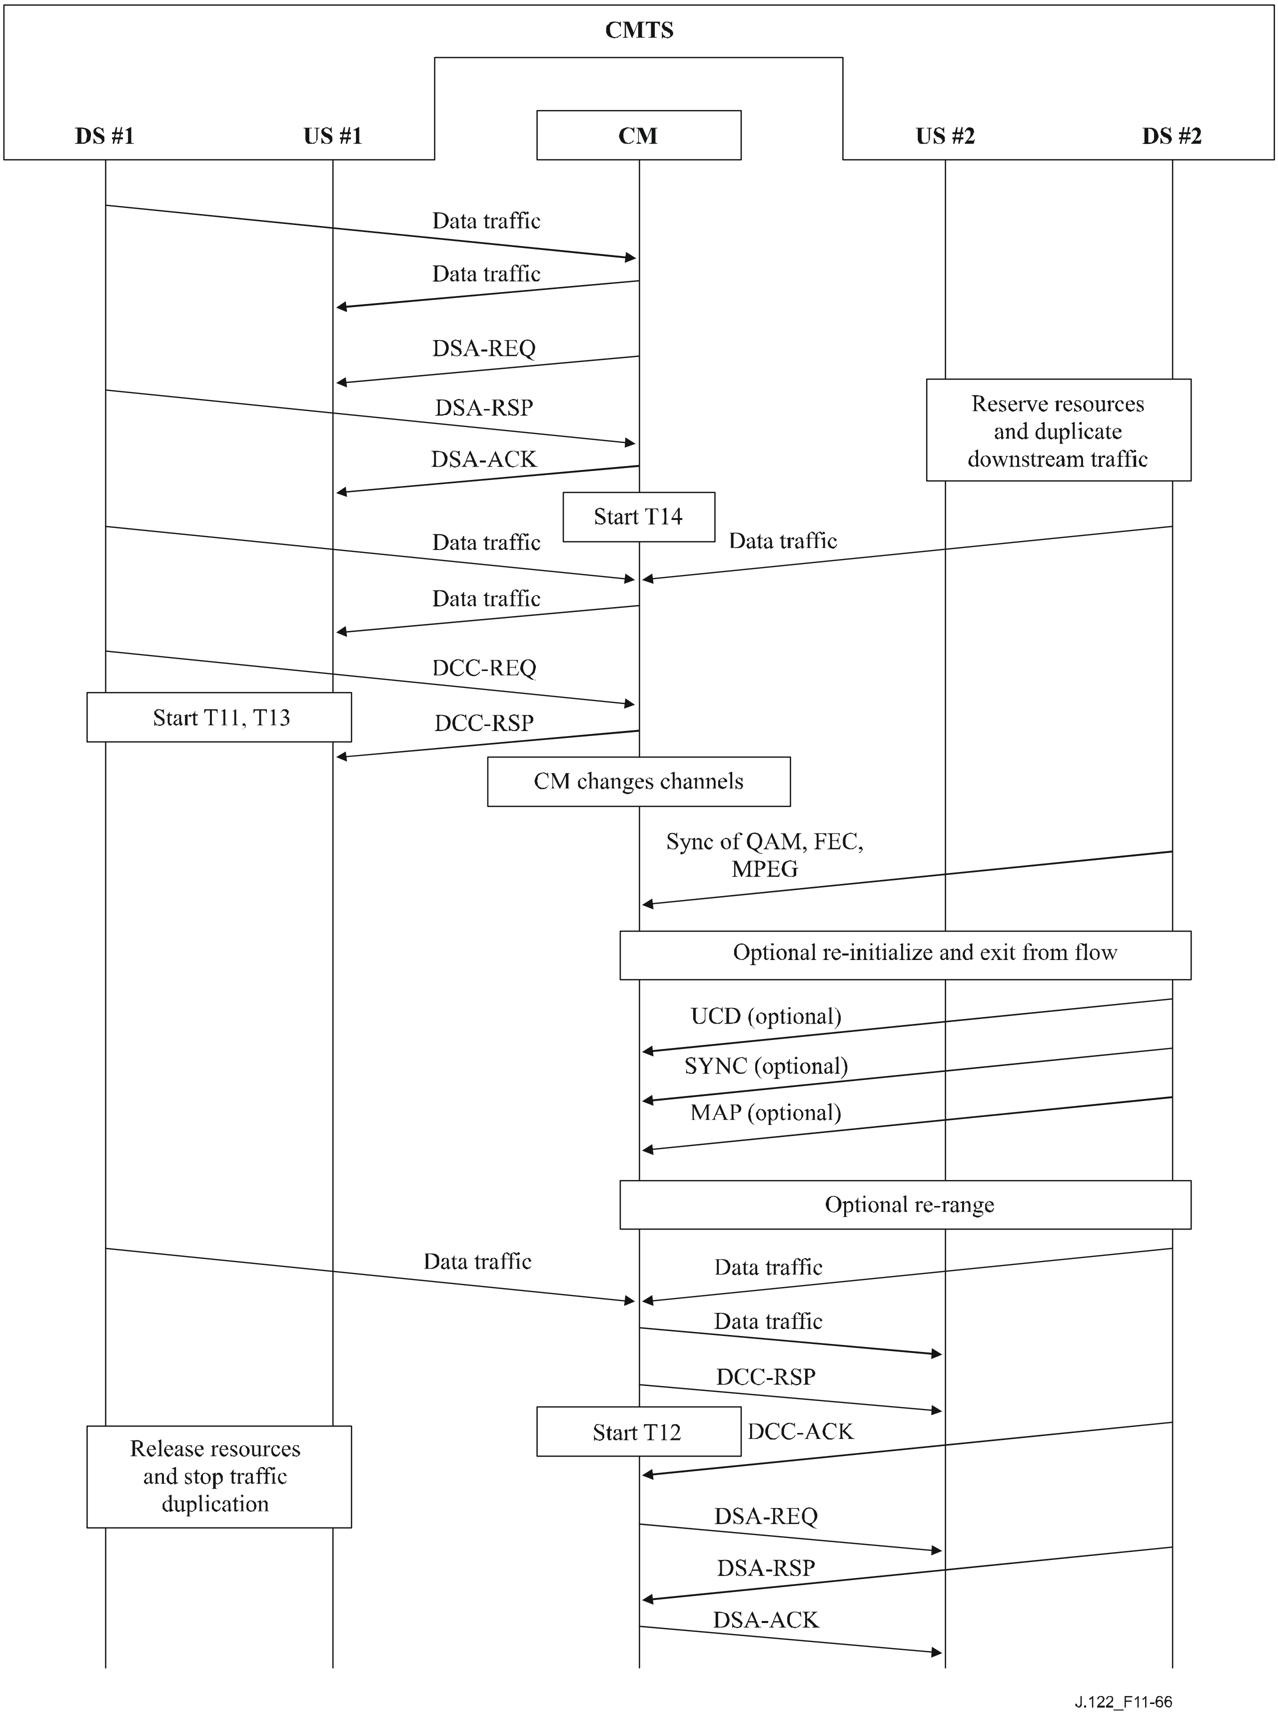

Figure 11-66 – DCC example operational flow

The states for the old and new CMTSs are shown as separate flow diagrams, since the old and new CMTS may be different. If the CMTSs are the same (e.g., the same MAC domain), the CMTS will need to run both sets of state machines concurrently.

The flow diagrams show points where explicit signalling between the old and new CMTS is desirable, especially for near-seamless operation. The mechanism for this signalling is beyond the scope of this Recommendation.

Note that the flow diagrams for both old and new CMTSs have been carefully crafted to handle many error conditions, such as:

- If the CM does not respond to the DCC-REQ (or responds with a reject conf code) and does not move, then it will be allowed to remain on the old channel. Resources on the new channel will be released (old CMTS signals DCC aborted to the new CMTS).
- If the CM DCC-RSP (depart) is lost, but the CM moves and arrives on the new CMTS, the new CMTS will signal that the CM has arrived to the old CMTS, allowing it to remove resources.
- If the CM DCC-RSP (depart) is received and the CM DCC-RSP (arrive) is lost, but the new CMTS otherwise detects the presence of the CM, the DCC transaction is considered successful, and the CM is allowed to remain on the new channel.
- If the CM DCC-RSP (depart) and (arrive) are lost, but the new CMTS otherwise detects the presence of the CM, the new CMTS will signal that the CM has arrived to the old CMTS, allowing it to remove resources, and the CM is allowed to remain on the new channel.
- If the CM DCC-RSP (depart) is received, but the CM never arrives, the new CMTS will detect this and remove resources after T15 expires.
- If the CM DCC-RSP (depart) is lost and the CM never arrives, the old CMTS will signal DCC aborted to the new CMTS, allowing it to remove resources. The old CMTS will use a different mechanism outside the scope of the DCC flow diagrams (such as ranging time-out) to remove resources on the old channels.
- If the CMTS DCC-ACK is lost and the DCC-RSP retry counter is expired, the CM will log an error and continue to the operational state.

There is a race condition that is not addressed in the flow diagrams; if the CM DCC-RSP (depart) is lost, the old CMTS will signal DCC aborted to the new CMTS. If the CM is in the process of moving, but has not yet arrived, the new CMTS will remove resources. This will prevent the CM from arriving successfully, unless it is able to complete the jump and arrive in approximately 1.2 seconds (3 retries of the DCC-REQ).

11.4.5.5.2 Example timing

11.4.5.5.2.1 Upstream and downstream change (use channel directly: CMTS supplies all TLV hints)

In this example, the current CMTS sends a DCC-REQ message requesting that the CM switch both upstream and downstream channels. The DCC-REQ message includes the UCD substitution TLV, the SYNC substitution TLV, the downstream parameter TLVs, and the initialization technique TLV of 4 (use channel directly). The CM does not include the CM jump time TLV in the DCC-RSP.

The destination CMTS has the following local parameters:

- UCD interval: 1 s;
- SYNC interval: 10 ms;
- unicast ranging interval: 1 s.

The destination CMTS calculates the T15 timer value. The definition of the formula used to determine T15 is shown below. The variables used in calculating T15 are explained in the table below.

T15 = CmJumpTime + CmtsRxRngReq

T15 = 1.3 s + (2.04 s) = 3.34 s

Since 3.34 s is less than the minimum value of the T15 timer, the CMTS sets the T15 timer to the minimum value for 4 seconds.

Variable	Value	Explanation
CmJumpTime	1.3 s	Since the CM did not include the optional jump time TLV, the CMTS will use the default value of 1.3 seconds.
CmtsRxRngReq	2.02 s 2 × (1 sec) + 40 ms	Two times the CMTS time period between unicast ranging opportunities plus 20-40 milliseconds for MAP and RNG-REQ transmission time and CMTS RNG-REQ processing time.

The CM synchronizes to the downstream parameters on the new channel, applies the UCD supplied in the DCC-REQ, collects MAP messages on the new channel, and resumes normal data transmission on the destination channels. This occurs within the recommended performance of 1 second.

11.4.5.5.2.2 Upstream and downstream change (station maintenance: CMTS supplies no TLV hints)

In this example, the current CMTS sends a DCC-REQ message requesting that the CM switch both upstream and downstream channels. The DCC-REQ message includes the initialization technique TLV of 2 (perform station maintenance). It also includes the required UCD substitution TLV and SYNC substitution sub-TLV. The CM does not include the CM jump time TLV in the DCC-RSP.

The destination CMTS has the following local parameters:

- UCD interval: 1 s;
- SYNC interval: 10 ms;
- unicast ranging interval: 5 s.

The destination CMTS starts scheduling the CM immediately after it sends the DCC-REQ. The destination CMTS calculates the T15 timer value. The definition of the formula used to determine T15 is shown below. The variables used in calculating T15 are explained in the table below.

T15 = CmJumpTime + CmtsRxRngReq

T15 = 1.3 s + (2.04 s) = 3.34 s

Variable	Value	Explanation
CmJumpTime	1.3 s	Since the CM did not include the optional jump time TLV, the CMTS will use the default value of 1.3 seconds.
CmtsRxRngReq	10.04 s 2 × (5 s) + 40 ms	Two times the CMTS time period between unicast ranging opportunities plus 20-40 milliseconds for MAP and RNG-REQ transmission time and CMTS RNG-REQ processing time.

The CM should synchronize to the downstream parameters on the new channel, apply the UCD message provided, collect MAP messages on the destination channel without waiting for a downstream SYNC on the destination channel, perform station maintenance on the destination channel, and resume normal data transmission on the destination channels.

These events occur in less than two seconds; this is within the acceptable performance criteria. The DCC transaction occurs within the recommended four second sum of CM jump time and two ranging intervals (0 + 2 s = 2 s).

11.4.5.6 Autonomous load balancing

The CMTS MUST be capable of autonomous load balancing of CMs across all of its upstream and downstream channels using the DCC-REQ message, plant configuration permitting.

Load balancing is manageable on a per-CM basis. The CMTS assigns each CM:

- to a set of channels (a load balancing group) among which it can be moved by the CMTS;
- a policy that governs if and when the CM can be moved;
- a priority value that can be used by the CMTS in order to select CMs to move.

11.4.5.6.1 Load balancing groups

A load balancing group is a cluster of downstream and associated upstream physical channels among which modems can be autonomously load balanced. The operator must define load balancing groups to be consistent with the physical plant topology. Load balancing groups can be defined in order to cater for a specific class of service (e.g., residential or business).

There are two types of load balancing groups, general load balancing groups and restricted load balancing groups. A restricted load balancing group is associated with a specific, provisioned set of cable modems, while general load balancing groups are open for CMs that are not provisioned into a restricted load balancing group. General load balancing groups can only be used when the plant topology allows the location of CMs to be unambiguously determined by the upstream channel on which they have registered. Restricted load balancing groups are used to accommodate a topology-specific or provisioning-specific restriction (such as a set of channels reserved for business customers). A restricted load balancing group could be a subset of one or more general load balancing groups.

The CMTS MUST NOT associate an upstream channel with more than one general load balancing group. The CMTS can associate an upstream channel with any number of restricted load balancing groups (and potentially a single general load balancing group as well). The CMTS can associate a downstream channel with any number of load balancing groups, general and/or restricted.

The CMTS assigns modems to general load balancing groups based on the upstream channel on which they are operating. The CMTS will assign a modem to a restricted load balancing group only if it is explicitly provisioned (via SNMP or configuration file TLV) to be a member of that group. The CMTS MUST NOT assign a CM to more than one load balancing group.

When the CMTS receives a REG-REQ message, the CMTS MUST identify whether this CM has been assigned to a restricted load balancing group via SNMP (see [DOCS5]). If the CM is not assigned to a restricted load balancing group via SNMP, the CMTS MUST check for the presence of the CM load balancing group TLV in the REG-REQ message to identify assignment to a restricted load balancing group. If the CM is assigned to a restricted load balancing group via SNMP, the CMTS MUST ignore the CM load balancing group TLV in the REG-REQ message. If the REG-REQ contains a general load balancing group ID or a group ID that is not defined on the CMTS, the CMTS MUST ignore the group ID.

If the CM has been assigned to a restricted load balancing group (either via SNMP or via the config file), and the CMTS detects that the CM is registering on a channel pair that is not associated with the assigned load balancing group, when registration completes, the CMTS MUST initiate a DCC-REQ to move the CM to the appropriate channel(s) in the assigned group. If the CM has not been assigned to a restricted load balancing group via SNMP or the config file, the CMTS MUST assign it to the general load balancing group that corresponds to the upstream channel on which it registered. If the upstream channel is not associated with a general load balancing group, the CMTS MUST NOT assign the CM to any load balancing group.

As part of autonomous load balancing operations, the CMTS MUST adhere to the following restrictions:

- The CMTS MUST NOT direct a CM to move to a channel outside the load balancing group to which it is assigned.
- If a CM is not assigned to a load balancing group, the CMTS MUST NOT move the CM.
- The CMTS MUST NOT move a DOCSIS 1.0 – or DOCSIS 1.1 – compliant CM, or a DOCSIS 2.0 – compliant CM that has 2.0 mode disabled, to a type 3 upstream channel.

If, during normal operation, the CMTS is directed (via SNMP or otherwise) to move a CM to a channel that is outside the load balancing group to which it is assigned, the CMTS MUST assign the CM to the appropriate general load balancing group based on the new upstream channel, or to no load balancing group if the new upstream channel is not associated with a general load balancing group. Note that a CM provisioned into a restricted load balancing group will (unless the provisioning is changed) be returned to its provisioned group by the CMTS upon a re-initialize MAC operation. Thus, a manual DCC operation that moves a CM outside of its restricted load balancing group will generally be only a temporary move, until the CM re-initializes. Further, specifying initialization technique zero in the DCC request would result in the CM being returned immediately to its provisioned group.

11.4.5.6.2 Initialization techniques

The description of a load balancing group includes the initialization technique(s) that could be used when autonomously load balancing a cable modem within the group. The initialization technique definition for each load balancing group is represented in the form of a bit map, with each bit representing a specific technique (bits 0-4).

The CMTS MUST support, via SNMP, the override of this initialization technique for channel changes between specific pairs of logical upstream channels. If an override is present for a particular pair of channels, the CMTS MUST use the initialization technique for the channel pair. Otherwise, the CMTS MUST use the default initialization technique definition for the load balancing group.

11.4.5.6.3 Load balancing policies

Load balancing policies allow control over the behaviour of the autonomous load balancing process on a per-CM basis. A load balancing policy is described by a set of conditions (rules) that govern the autonomous load balancing process for the CM. When a load balancing policy is defined by multiple rules, all of the rules apply in combination.

Load balancing rules and the load balancing policy definition mechanism have been created to allow for specific vendor-defined load balancing actions. However, there are two basic rules that the CMTS is required to support.

The CMTS MUST support the following basic rules:

- Prohibit load balancing using a particular CM.
- Prohibit load balancing using a particular CM during certain times of the day.

The policy ID value of zero is reserved to indicate the CMTS's basic load balancing mechanism, which does not need to be defined by a set of rules.

Each load balancing group has a default load balancing policy. During the registration process, the CMTS MUST assign the CM a load balancing policy ID. The policy ID may be assigned to a cable modem via the cable modem config file. The CMTS MUST assign the CM the load balancing policy ID provisioned in the config file and sent in the REG-REQ, if it exists. Otherwise, the CMTS MUST assign the CM the default policy ID defined for the load balancing group.

The per-CM load balancing policy ID assignment can be modified at any time while the CM is in the operational state via SNMP or internal CMTS processes, however the policy ID is always overwritten upon receipt of a REG-REQ message.

11.4.5.6.4 Load balancing priorities

A load balancing priority is an index that defines a rank or level of importance, which is used to apply differential treatment between CMs in the CMTS's load balancing decision process.

In general, a lower load balancing priority indicates a higher likelihood that a CM will be moved due to load balancing operations. The CMTS MAY take many factors into account when selecting a CM to move, of which priority is only one. When other factors are equal, the CMTS SHOULD preferentially move a CM with lower load balancing priority over one with higher load balancing priority.

The CMTS MUST associate each cable modem with a load balancing priority. Priority may be assigned to a cable modem via the cable modem config file. The CMTS MUST assign the CM the load balancing priority provisioned in the config file and sent in the REG-REQ, if it exists. If a cable modem has not been assigned a priority, the cable modem is associated with the default (lowest) load balancing priority value of zero.

The per-CM load balancing priority assignment can be modified at any time while the CM is in the operational state via SNMP or internal CMTS processes as dictated by a specific load balancing policy; however, the priority assignment is always overwritten upon receipt of a REG-REQ message.

11.4.5.6.5 Load balancing and multicast

The IGMP protocol requires hosts to suppress IGMP messages that are not necessary for the router to maintain multicast group membership state. Clause 5.3.1 extends these IGMP requirements to the DOCSIS access network by requiring CMs to suppress messages that are deemed to be superfluous for the CMTS. As a result, the CMTS is not guaranteed to be aware of multicast group membership on a per-CM basis. For an active multicast group, there could be any number of CMs that have group members and that are actively forwarding multicast traffic, but that have not sent a membership report to the CMTS. This lack of CMTS awareness can create a situation in which load balancing and multicast conflict.

In order to efficiently manage multicast traffic and balance load across a load balancing group, it is reasonable to expect that the CMTS might attempt to reduce the amount of duplicated multicast traffic by consolidating all members for a specific multicast group to a single downstream channel in the load balancing group. More generally, a load balancing algorithm will perform more effectively if it takes into account both the unicast and multicast traffic load for each CM when making decisions on where and when to move CMs. This is made difficult when the CMTS is unaware of which CMs have multicast group members and which do not.

If a CM with active multicast sessions is moved from its current downstream to a new downstream that is not carrying the multicast traffic, the session will be interrupted until the CM or CPE sends a membership report. In order to reduce the interruption of multicast service, CMs that implement active IGMP mode (see clause 5.3.1) SHOULD send a membership report for all active multicast groups upon completion of a DCC operation that involves a downstream channel change.

The multicast issues are alleviated to some degree when BPI+ is enabled, and are alleviated further when multicast traffic is encrypted using dynamic security associations (see [DOCS8]).

When BPI+ is enabled, a CM will, upon receiving an IGMP "join" message on its CPE interface, send a SA map request message to the CMTS. Since this message is only sent at the moment multicast group membership begins, it does not provide any indication of ongoing membership. Because multicast group membership can be transient, the past receipt of a map request for a particular multicast group, although necessary, is not a sufficient condition to alert the CMTS that the CM currently has members for that multicast group. The absence of a map request is sufficient evidence that the CM does not have members for the multicast group.

If the multicast traffic for a particular multicast group is encrypted using a dynamic security association, the CMTS can monitor the reception of TEK key requests and gain knowledge of multicast group membership. Since it is optional functionality for a CM to stop the TEK state machine (and discontinue sending key requests) when there are no longer members for multicast groups mapped to a particular security association, the continued receipt of key requests by the CMTS does not necessarily indicate continued multicast group membership. The lack of continuing key requests, however, does indicate lack of members.

11.4.5.6.6 Examples demonstrating autonomous load balancing

Figure 11-67 shows an example combining network which illustrates the definition of general load balancing groups and the use of restricted load balancing groups to resolve topological ambiguities.

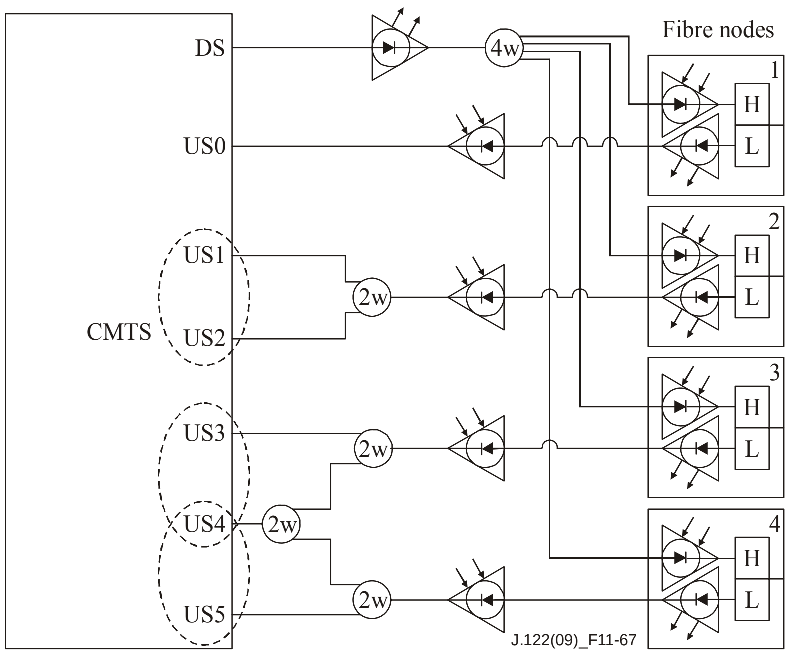

Figure 11-67 – Example combining network 1

In this example, there are six upstream channels (US0-US5) that are members of a single MAC domain. All six upstream channels are associated with a single downstream channel (DS). The downstream is split over all four fibre nodes, while the six upstreams return from the four nodes via the combining network shown, such that each upstream channel is not physically connected to each fibre node. In particular, fibre node 1 connects to US0 only, fibre node 2 connects to both US1 and US2, fibre node 3 connects to both US3 and US4, and fibre node 4 connects to both US4 and US5.

In this situation, the load balancing groups could be defined as follows:

Load balancing group 1:

Group ID:	1
Type:	General
Downstream channels:	DS

Upstream channels: US1, US2

Load balancing group 2:

Group ID: 2

Type: Restricted

Downstream channels: DS

Upstream channels: US3, US4

Load balancing group 3:

Group ID: 3

Type: Restricted

Downstream channels: DS

Upstream channels: US4, US5

Note that a REG-REQ on either upstream channel US1 or US2 uniquely identifies the load balancing group to which a CM can be assigned, hence those two channels form the general load balancing group 1. Upstream channels US3-US5 have a more complex topology, since US4 is shared across two fibre nodes. To resolve the topological ambiguities that would arise by a REG-REQ received on US4, two restricted load balancing groups have been defined (Group IDs 2 and 3). In order to be load balanced, each CM that is attached to fibre node 3 would need to be provisioned to be a member of restricted load balancing group 2, while each CM attached to fibre node 4 would need to be provisioned into restricted load balancing group 3. If a CM were to register on one of these channels without having been provisioned into the appropriate restricted load balancing group, the CMTS would not associate the CM with any load balancing group (which results in the CM not being load balanced).

Also, note that US0 is not a member of any load balancing group. CMs that register on that upstream channel will not be load balanced to another channel.

Figure 11-68 shows a second example, in which two MAC domains are shared across two fibre nodes in a complex combining network. In this example, a pair of upstream channels (one from each MAC domain) are set aside for a particular customer group (e.g., business customers), a restricted load balancing group is formed to allow load balancing for those customers.

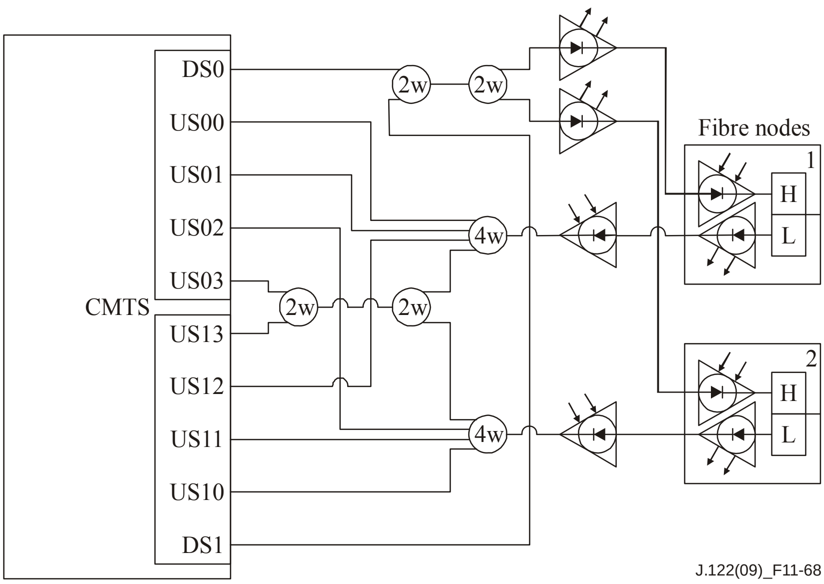

Figure 11-68 – Example combining network 2

Load balancing group 1:

Group ID:	1
Type:	General
Downstream channels:	DS0, DS1
Upstream channels:	US00, US01, US12
Subgroup:	DS0, US00, US01

Load balancing group 2:

Group ID:	2
Type:	General
Downstream channels:	DS0, DS1
Upstream channels:	US10, US11, US02
Subgroup:	DS1, US10, US11

Load balancing group 3:

Group ID:	3
Type:	Restricted
Downstream channels:	DS0, DS1
Upstream channels:	US03, US13

11.5 Fault detection and recovery

Fault detection and recovery occurs at multiple levels.

- At the physical level, FEC is used to correct errors where possible – refer to clause 6 for details.
- The MAC protocol protects against errors through the use of checksum fields across both the MAC header and the data portions of the packet – refer to clause 8 for details.
- All MAC management messages are protected with a CRC covering the entire message, as defined in clause 8. Any message with a bad CRC MUST be discarded by the receiver.

Table 11-2 shows the recovery process that MUST be taken following the loss of a specific type of MAC message.

Table 11-2 – Recovery process on loss of specific MAC messages

Message name	Action following message loss
SYNC	The CM can lose SYNC messages for a period of the lost SYNC interval (see Annex B) before it has lost synchronization with the network. A CM that has lost synchronization MUST NOT use the upstream and MUST try to re-establish synchronization.
UCD	During CM initialization, the CM MUST receive a usable UCD[a)] before transmitting on the upstream. When in the "obtain upstream parameters" state of CM initialization process, if the CM does not receive a usable UCD within the T1 time-out period, the CM MUST NOT transmit on the upstream and MUST scan for another downstream channel. After receiving a usable UCD, whenever the CM receives an unusable UCD or a MAP with a UCD count that does not match the configuration change count of the last UCD received, the CM MUST NOT transmit on the upstream and MUST start the T1 timer. If the T1 timer expires under these circumstances, the CM MUST reset and re-initialize its MAC connection.
MAP	A CM MUST NOT transmit without a valid upstream bandwidth allocation. If a MAP

Table 11-2 – Recovery process on loss of specific MAC messages

Message name	Action following message loss
	is missed due to error, the CM MUST NOT transmit for the period covered by the MAP.
RNG-REQ RNG-RSP	If a CM fails to receive a valid ranging response within a defined time-out period after transmitting a request, the request MUST be retried a number of times (as defined in Annex B). Failure to receive a valid ranging response after the requisite number of attempts MUST cause the modem to reset and reinitialize its MAC connection.
REG-REQ REG-RSP	If a CM fails to receive a valid registration response within a defined time-out period after transmitting a request, the request will be retried a number of times (as defined in Annex B). Failure to receive a valid registration response after the requisite number of attempts will cause the modem to reset and reinitialize its MAC connection.
UCC-REQ UCC-RSP	If a CMTS fails to receive a valid upstream channel change response within a defined time-out period after transmitting a request, the request MUST be retried a number of times (as defined in Annex B). Failure to receive a valid response after the requisite number of attempts MUST cause the CMTS to consider the CM as unreachable.
[a] A usable UCD is one that contains legal profiles that the modem can understand. The CM MAY also require that the UCD count of the MAPs received match the configuration change count field of the last received UCD before it considers the UCD as usable.	

Annex D of [DOCS5] contains a list of error codes with more useful information as to the failure of the PHY and MAC layers. Refer to clause 8.2.8 for additional information.

Messages at the network layer and above are considered to be data packets by the MAC sublayer. These are protected by the CRC field of the data packet and any packets with bad CRCs are discarded. Recovery from these lost packets is in accordance with the upper layer protocol.

11.5.1 Prevention of unauthorized transmissions

A CM SHOULD include a means for terminating RF transmission if it detects that its own carrier has been on continuously for longer than the longest possible valid transmission.

12 Supporting future new cable modem capabilities

12.1 Downloading cable modem operating software

A CMTS SHOULD be capable of being remotely reprogrammed in the field via a software download over the network.

The cable modem MUST be capable of being remotely reprogrammed in the field via a software download over the network. This software download capability MUST allow the functionality of the cable modem to be changed without requiring that cable system personnel physically revisit and reconfigure each unit. It is expected that this field programmability will be used to upgrade cable modem software to improve performance, accommodate new functions and features (such as enhanced class-of-service support), correct any design deficiencies discovered in the software, and to allow a migration path as the data-over-cable interface Recommendation evolves.

The mechanism used for download MUST be TFTP file transfer. The mechanism by which transfers are secured and authenticated is in [DOCS8]. The transfer MUST be initiated in one of two ways:

- An SNMP manager requests the CM to upgrade.

- If the software upgrade file name in the CM's configuration file does not match the current software image of the CM, the CM MUST request the specified file via TFTP from the software server.

NOTE – The software server IP address is a separate parameter. If present, the CM MUST attempt to download the specified file from this server. If not present, the CM MUST attempt to download the specified file from the configuration file server.

The CM MUST verify that the downloaded image is appropriate for itself. If the image is appropriate, the CM MUST write the new software image to non-volatile storage. Once the file transfer is completed successfully, the CM MUST restart itself with the new code image.

If the CM is unable to complete the file transfer for any reason, it MUST remain capable of accepting new software downloads (without operator or user interaction), even if power or connectivity is interrupted between attempts. The CM MUST log the failure and MAY report it asynchronously to the network manager.

Following upgrade of the operational software, the CM MAY need to follow one of the procedures described above in order to change channels to use the enhanced functionality.

If the CM is to continue to operate in the same upstream and downstream channels as before the upgrade, then it MUST be capable of interworking with other CMs that may be running previous software releases.

Where software has been upgraded to meet a new version of the Recommendation, it is critical that it MUST interwork with the previous version in order to allow a gradual transition of units on the network.

Annex A

Well-known addresses

(This annex forms an integral part of this Recommendation)

A.1 MAC addresses

MAC addresses described here are defined using the Ethernet/ISO/IEC 8802-3 convention as bit-little-endian.

The following multicast address MUST be used to address the set of all CM MAC sublayers; for example, when transmitting allocation map PDUs:

01-E0-2F-00-00-01

The addresses in the range

01-E0-2F-00-00-02 through 01-E0-2F-00-00-0F

are reserved for future definition. Frames addressed to any of these addresses SHOULD NOT be forwarded out of the MAC-sublayer domain.

A.2 MAC service IDs

The following MAC service IDs have assigned meanings. Those not included in this table are available for assignment, either by the CMTS or administratively.

A.2.1 All CMs and no CM service IDs

The following service IDs are used in MAPs for special purposes or to indicate that any CM can respond in the corresponding interval.

0x0000 is addressed to no CM. This address is typically used when changing upstream burst parameters so that CMs have time to adjust their modulators before the new upstream settings take effect. This is also the "initialization SID" used by the CM during initial ranging.

0x3FFF is addressed to all CMs. It is typically used for broadcast request intervals or initial maintenance intervals.

A.2.2 Well-known multicast service IDs

The following service IDs are only used for request/data IEs. They indicate that any CM can respond in a given interval, but that the CM must limit the size of its transmission to a particular number of mini-slots (as indicated by the particular multicast SID assigned to the interval).

0x3FF1-0x3FFE is addressed to all CMs. IDs in this range are available for small data PDUs as well as requests (used only with request/data IEs). The last digit indicates the frame length and transmission opportunities as follows:

- 0x3FF1: Within the interval specified, a transmission may start at any mini-slot, and must fit within one mini-slot.
- 0x3FF2: Within the interval specified, a transmission may start at every other mini-slot, and must fit within two mini-slots (e.g., a station may start transmission on the first mini-slot within the interval, the third mini-slot, the fifth, etc.).
- 0x3FF3: Within the interval specified, a transmission may start at any third mini-slot, and must fit within three mini-slots (e.g., starts at first, fourth, seventh, etc.).
- 0x3FF4: Starts at first, fifth, ninth, etc.
- 0x3FFD: Starts at first, fourteenth (14th), twenty-seventh (27th), etc.

- 0x3FFE: Within the interval specified, a transmission may start at any 14th mini-slot, and must fit within 14 mini-slots.

A.2.3 Priority request service IDs

The following service IDs (0x3Exx) are reserved for request IEs (refer to clause C.2.2.5.1, "traffic priority").

- If 0x01 bit is set, priority zero can request.
- If 0x02 bit is set, priority one can request.
- If 0x04 bit is set, priority two can request.
- If 0x08 bit is set, priority three can request.
- If 0x10 bit is set, priority four can request.
- If 0x20 bit is set, priority five can request.
- If 0x40 bit is set, priority six can request.
- If 0x80 bit is set, priority seven can request.

Bits can be combined as desired by the CMTS upstream scheduler for any request IUCs.

A.3 MPEG PID

All DOCS data MUST be carried in MPEG-2 packets with the header PID field set to 0x1FFE.

Annex B

Parameters and constants

(This annex forms an integral part of this Recommendation)

System	Name	Time reference	Minimum value	Default value	Maximum value
CMTS	Sync interval	Nominal time between transmission of SYNC messages (see clause 8.3.2)			200 ms
CMTS	UCD interval	Time between transmission of UCD messages (see clause 8.3.3, "upstream channel descriptor (UCD)")			2 s
CMTS	Max MAP pending	The number of mini-slots that a CMTS is allowed to map into the future (see clause 8.3.4)			4096 mini-slot times
CMTS	Ranging interval	Time between transmission of broadcast ranging requests (see clause 9.3.3)			2 s
CM	Lost sync interval	Time since last received sync message before synchronization is considered lost			600 ms
CM	Contention ranging retries	Number of retries on contention ranging requests (see clause 11.2.4)	16		
CM CMTS	Invited ranging retries	Number of retries on inviting ranging requests (see clause 11.2.4)	16		
CM	Request retries	Number of retries on bandwidth allocation requests	16		
CM CMTS	Registration request/ response retries	Number of retries on registration requests/responses	3		
CM	Data retries	Number of retries on immediate data transmission	16		
CMTS	CM MAP processing time	Time provided between arrival of the last bit of a MAP at a CM and effectiveness of that MAP (see clause 9.1.5)	(200 + M/5.12) µs (see clause 6.2.17)		
CMTS	CM ranging response processing time	Minimum time allowed for a CM following receipt of a ranging response before it is expected to reply to an invited ranging request	1 ms		
CMTS	CM configuration	The maximum time allowed for a CM, following receipt of a configuration file, to send a registration request to a CMTS	30 s		

System	Name	Time reference	Minimum value	Default value	Maximum value
CM	T1	Wait for UCD time-out			5 × UCD interval maximum value
CM	T2	Wait for broadcast ranging time-out			5 × ranging interval
CM	T3	Wait for ranging response	50 ms	200 ms	200 ms
CM	T4	Wait for unicast ranging opportunity. If the pending-till-complete field was used earlier by this modem, then the value of that field must be added to this interval	30 s		35 s
CMTS	T5	Wait for upstream channel change response			2 s
CM CMTS	T6	Wait for REG-RSP and REG-ACK			3 s
CM CMTS	Mini-slot size for 1.x channels	Size of mini-slot for upstream transmission. For channels that support DOCS 1.x CMs	32 modulation intervals		
CM CMTS	Mini-slot size for DOCS 2.0 Only Channels	Size of mini-slot for upstream transmission. For channels that do not support DOCS 1.x CMs	16 symbols		
CM CMTS	Timebase tick	System timing unit	6.25 µs		
CM CMTS	DSx request retries	Number of time-out retries on DSA/DSC/DSD requests	3		
CM CMTS	DSx response retries	Number of timeout retries on DSA/DSC/DSD responses	3		
CM CMTS	T7	Wait for DSA/DSC/DSD response timeout			1 s
CM CMTS	T8	Wait for DSA/DSC acknowledge timeout			300 ms
CM	TFTP backoff start	Initial value for TFTP backoff	1 s		
CM	TFTP backoff end	Last value for TFTP backoff	16 s		
CM	TFTP request retries	Number of retries on TFTP request	16		
CM	TFTP download retries	Number of retries on entire TFTP downloads	3		
CM	TFTP wait	The wait between TFTP retry sequences	10 min		
CM	ToD retries	Number of retries per ToD retry period	3		
CM	ToD retry period	Time period for ToD retries	5 min		

System	Name	Time reference	Minimum value	Default value	Maximum value
CMTS	T9	Registration timeout, the time allowed between the CMTS sending a RNG-RSP (success) to a CM, and receiving a REG-REQ from that same CM	15 min	15 min	
CM CMTS	T10	Wait for transaction end timeout			3 s
CMTS	T11	Wait for a DCC response on the old channel			300 ms
CM	T12	Wait for a DCC acknowledge			300 ms
CMTS	T13	Maximum holding time for QoS resources for DCC on the old channel			1 s
CM	T14	Minimum time after a DSx reject-temp-DCC and the next retry of DSx command	2 s		
CMTS	T15	Maximum holding time for QoS resources for DCC on the new channel	2 s		35 s
CM	T16	Maximum length of time CM remains in test mode after receiving TST-REQ message			30 min
CM	T17	Maximum time that CM MUST inhibit transmissions on a channel in response to its ranging class ID matching a bit value in the ranging hold-off priority field			300 s
CMTS	DCC-REQ retries	Number of retries on dynamic channel change request	3		
CM	DCC-RSP retries	Number of retries on dynamic channel change response	3		
CM	Lost DCI-REQ interval	Time from sending DCI-REQ and not receiving a DCI-RSP			2 s
CM	DCI-REQ retry	Number of retries of DCI-REQ before rebooting			16
CM	DCI backoff start	Initial value for DCI backoff	1 s		
CM	DCI backoff end	Last value for DCI backoff	16 s		
CMTS	CM UCD processing time	Time between the transmission of the last bit of a UCD with a new change count and the transmission time of the first bit of the first MAP using the new UCD (see clause 11.3.2).	1 ms		

Annex C

Common radio frequency interface encodings

(This annex forms an integral part of this Recommendation)

C.1 Encodings for configuration and MAC-layer messaging

The following type/length/value encodings MUST be used in both the configuration file (see Annex D), in CM registration requests and in dynamic service messages. All multi-octet quantities are in network-byte order, i.e., the octet containing the most-significant bits is the first transmitted on the wire.

The following configuration settings MUST be supported by all CMs which are compliant with this Recommendation.

C.1.1 Configuration file and registration settings

These settings are found in the configuration file and, if present, MUST be forwarded by the CM to the CMTS in its registration request.

C.1.1.1 Downstream frequency configuration setting

The receive frequency to be used by the CM unless a downstream channel list is present. It is an override for the channel selected during scanning. This is the centre frequency of the downstream channel in Hz stored as a 32-bit binary number.

Type	Length	Value
1	4	Rx frequency

Valid range: The receive frequency MUST be a multiple of 62'500 Hz.

C.1.1.2 Upstream channel ID configuration setting

The Upstream Channel ID that the CM MUST use. The CM MUST listen on the defined downstream channel until an upstream channel description message with this ID is found. It is an override for the channel selected during initialization.

Type	Length	Value
2	1	Channel ID

C.1.1.3 Network access control object

If the value field is a 1, CPEs attached to this CM are allowed access to the network, based on CM provisioning. If the value of this field is a 0, the CM MUST NOT forward traffic from attached CPE to the RF MAC network, but MUST continue to accept and generate traffic from the CM itself. The value of this field does not affect CMTS service flow operation and does not affect CMTS data forwarding operation.

Type	Length	On/off
3	1	1 or 0

The intent of "NACO = 0" is that the CM does not forward traffic from any attached CPE onto the cable network (a CPE is any client device attached to that CM, regardless of how that attachment is implemented). However, with "NACO = 0", management traffic to the CM is not restricted. Specifically, with NACO off, the CM remains manageable, including sending/receiving management traffic such as (but not limited to):

- ARP: Allow the modem to resolve IP addresses, so it can respond to queries or send traps.
- DHCP: Allow the modem to renew its IP address lease.
- ICMP: Enable network troubleshooting for tools such as "ping" and "traceroute."
- ToD: Allow the modem to continue to synchronize its clock after boot.
- TFTP: Allow the modem to download either a new configuration file or a new software image.
- SYSLOG: Allow the modem to report network events.
- SNMP: Allow management activity

In DOCS v1.1, with NACO off, the primary upstream and primary downstream service flows of the CM remain operational only for management traffic to and from the CM. With respect to DOCS v1.1 provisioning, a CMTS should ignore the NACO value and allocate any service flows that have been authorized by the provisioning server.

C.1.1.4 DOCS 1.0 class-of-service configuration setting

This field defines the parameters associated with a DOCS 1.0 class of service. Any CM registering with a DOCS 1.0 class-of-service configuration setting MUST be treated as a DOCS 1.0 CM. Refer to clause 8.3.8, "registration response (REG-RSP)".

This field defines the parameters associated with a class of service. It is somewhat complex in that it is composed from a number of encapsulated type/length/value fields. The encapsulated fields define the particular class-of-service parameters for the class of service in question. Note that the type fields defined are only valid within the encapsulated class-of-service configuration setting string. A single class-of-service configuration setting is used to define the parameters for a single service class. Multiple class definitions use multiple class-of-service configuration setting sets.

Type	Length	Value
4	n	

C.1.1.4.1 Class ID

The value of the field specifies the identifier for the Class of Service to which the encapsulated string applies.

Type	Length	Value
4.1	1	

Valid range: The class ID MUST be in the range 1 to 16.

C.1.1.4.2 Maximum downstream rate configuration setting

For a single-SID modem, the value of this field specifies the maximum downstream rate in bits per second that the CMTS is permitted to forward to CPE unicast MAC addresses learned or configured as mapping to the registering modem.

For a multiple-SID modem, the aggregate value of these fields specifies the maximum downstream rate in bits per second that the CMTS is permitted to forward to CPE unicast MAC addresses learned or configured as mapping to the registering modem.

This is the peak data rate for packet PDU data (including destination MAC address and the CRC) over a one-second interval. This does not include MAC packets addressed to broadcast or multicast MAC addresses. The CMTS MUST limit downstream forwarding to this rate. The CMTS MAY delay, rather than drop, over-limit packets.

Type	Length	Value
4.2	4	

NOTE – This is a limit, not a guarantee that this rate is available.

C.1.1.4.3 Maximum upstream rate configuration setting

The value of this field specifies the maximum upstream rate in bits per second that the CM is permitted to forward to the RF network.

This is the peak data rate for packet PDU data (including destination address and the CRC) over a one-second interval. The CM MUST limit all upstream forwarding (both contention and reservation-based), for the corresponding SID, to this rate. The CM MUST include packet PDU data packets addressed to broadcast or multicast addresses when calculating this rate.

The CM MUST enforce the maximum upstream rate. It SHOULD NOT discard upstream traffic simply because it exceeds this rate.

The CMTS MUST enforce this limit on all upstream data transmissions, including data sent in contention. The CMTS SHOULD generate an alarm if a modem exceeds its allowable rate.

Type	Length	Value
4.3	4	

NOTE 1 – The purpose of this parameter is for the CM to perform traffic shaping at the input to the RF network and for the CMTS to perform traffic policing to ensure that the CM does not exceed this limit.

The CMTS could enforce this limit by any of the following methods:

a) discarding over-limit requests;

b) deferring (through zero-length grants) the grant until it is conforming to the allowed limit;

c) discarding over-limit data packets;

d) reporting to a policy monitor (for example, using the alarm mechanism) that is capable of incapacitating errant CMs.

NOTE 2 – This is a limit, not a guarantee that this rate is available.

C.1.1.4.4 Upstream channel priority configuration setting

The value of the field specifies the relative priority assigned to this service class for data transmission in the upstream channel. Higher numbers indicate higher priority.

Type	Length	Value
4.4	1	

Valid range: 0-7.

C.1.1.4.5 Guaranteed minimum upstream channel data rate configuration setting

The value of the field specifies the data rate in bit/s which will be guaranteed to this service class on the upstream channel.

Type	Length	Value
4.5	4	

C.1.1.4.6 Maximum upstream channel transmit burst configuration setting

The value of the field specifies the maximum transmit burst (in bytes) that this service class is allowed on the upstream channel. A value of zero means there is no limit.

NOTE – This value does not include any physical layer overhead.

Type	Length	Value
4.6	2	

C.1.1.4.7 Class-of-service privacy enable

This configuration setting enables/disables baseline privacy on a provisioned CoS.

Type	Length	Enable/Disable
4.7 (= CoS_BP_ENABLE)	1	1 or 0

Table C.1 – Sample DOCS 1.0 class-of-service encoding

Type	Length	Value (sub)type	Length	Value	
4	28				Class-of-service configuration setting
		1	1	1	Service class
		2	4	10'000'000	Max. downstream rate of 10 Mbit/s
		3	4	300'000	Max. upstream rate of 300 kbit/s
		4	1	5	Return path priority of 5
		5	4	64'000	Min. guaranteed 64 kbit/s
		6	2	1518	Max. Tx burst of 1518 bytes
4	28				Class-of-service configuration setting
		1	1	2	Service class 2
		2	4	5'000'000	Max. forward rate of 5 Mbit/s
		3	4	300'000	Max. return rate of 300 Mbit/s
		4	1	3	Return path priority of 3
		5	4	32'000	Min. guaranteed 32 kbit/s
		6	2	1518	Max. Tx burst of 1518 bytes

C.1.1.5 CM message integrity check (MIC) configuration setting

The value field contains the CM message integrity check code. This is used to detect unauthorized modification or corruption of the configuration file.

Type	Length	Value
6	16	d1, d2,... d16

C.1.1.6 CMTS message integrity check (MIC) configuration setting

The value field contains the CMTS message integrity check code. This is used to detect unauthorized modification or corruption of the configuration file.

Type	Length	Value
7	16	d1, d2,... d16

C.1.1.7 Maximum number of CPEs

The maximum number of CPEs that can be granted access through a CM during a CM epoch. The CM epoch is (from clause 5.1.2.3.1) the time between start-up and hard reset of the modem. The maximum number of CPEs MUST be enforced by the CM.

NOTE 1 – This parameter should not be confused with the number of CPE addresses a CM may learn. A modem may learn Ethernet MAC addresses up to its maximum number of CPE addresses (from clause 5.1.2.3.1). The maximum number of CPEs that are granted access through the modem is governed by this configuration setting.

Type	Length	Value
18	1	

The CM MUST interpret this value as an unsigned integer. The non-existence of this option, or the value 0, MUST be interpreted as the default value of 1.

NOTE 2 – This is a limit on the maximum number of CPEs a CM will grant access to. Hardware limitations of a given modem implementation may require the modem to use a lower value.

C.1.1.8 TFTP server timestamp

The sending time of the configuration file in seconds. The definition of time is as in [RFC 868].

Type	Length	Value
19	4	Number of seconds since 00:00 1 Jan 1900

NOTE – The purpose of this parameter is to prevent replay attacks with old configuration files.

C.1.1.9 TFTP server provisioned modem address

The IP address of the modem requesting the configuration file.

Type	Length	Value
20	4	IP address

NOTE – The purpose of this parameter is to prevent IP spoofing during registration.

C.1.1.10 Upstream packet classification configuration setting

This field defines the parameters associated with one entry in an upstream traffic classification list. Refer to clause C.2.1.1.

Type	Length	Value
22	n	

C.1.1.11 Downstream packet classification configuration setting

This field defines the parameters associated with one classifier in a downstream traffic classification list. Refer to clause C.2.1.2.

Type	Length	Value
23	n	

C.1.1.12 Upstream service flow encodings

This field defines the parameters associated with upstream scheduling for one service flow. Refer to clause C.2.2.1.

Type	Length	Value
24	n	

C.1.1.13 Downstream service flow encodings

This field defines the parameters associated with downstream scheduling for one service flow. Refer to clause C.2.2.2.

Type	Length	Value
25	n	

C.1.1.14 Payload header suppression

This field defines the parameters associated with payload header suppression.

Type	Length	Value
26	n	

C.1.1.15 Maximum number of classifiers

This is the maximum number of classifiers associated with admitted or active upstream service flows that the CM is allowed to have. Both active and inactive classifiers are included in the count.

This is useful when using deferred activation of provisioned resources. The number of provisioned service flows may be high and each service flow might support multiple classifiers. Provisioning represents the set of service flows the CM can choose between. The CMTS can control the QoS resources committed to the CM by limiting the number of service flows that are admitted. However, it may still be desirable to limit the number of classifiers associated with the committed QoS resources. This parameter provides that limit.

Type	Length	Value
28	2	Maximum number of active and inactive classifiers associated with admitted or active upstream service flows

The default value MUST be 0: No limit.

C.1.1.16 Privacy enable

This configuration setting enables/disables baseline privacy [DOCS8] on the primary service flow and all other service flows for this CM. If a DOCS 2.0 CM receives this setting in a configuration file, the CM is required to forward this setting as part of the registration request (REG-REQ) as specified in clause 8.3.7, regardless of whether the configuration file is DOCS 1.1-style or not, while this setting is usually contained only in a DOCS 1.1-style configuration file with DOCS 1.1 service flow TLVs.

Type	Length	Value
29	1	0: Disable 1: Enable

The default value of this parameter MUST be 1: privacy enabled.

C.1.1.17 DOCSIS extension field

The DOCSIS extension field is used to extend the capabilities of the DOCSIS specification, through the use of new and/or vendor-specific features.

The DOCSIS extension field MUST be encoded using TLV 43 and MUST include the vendor ID field (refer to clause C.1.3.2) to indicate whether the DOCSIS extension field applies to all devices, or only to devices from a specific vendor. The vendor ID MUST be the first TLV embedded inside the DOCSIS extension field. If the first TLV inside the DOCSIS extension field is not a vendor ID, then the TLV MUST be discarded. In this context, the vendor ID of 0xFFFFFF is reserved to signal that this DOCSIS extension field contains general extension information (see clause C.1.1.17.1), otherwise the DOCSIS extension field contains vendor-specific information (see clause C.1.1.17.2).

This configuration setting MAY appear multiple times. This configuration setting MAY be nested inside a packet classification configuration setting, a service flow configuration setting, or a service flow response. The same vendor ID MAY appear multiple times. However, there MUST NOT be more than one vendor ID TLV inside a single TLV 43.

The CM MUST ignore any DOCSIS extension field that it cannot interpret, but MUST still include the TLV in the REG-REQ message. The CM MUST NOT initiate the DOCSIS extension field TLVs.

Type	Length	Value
43	n	

C.1.1.17.1 General extension information

When using the DOCSIS extension field (TLV 43) to encode general extension information, the vendor ID of 0xFFFFFF MUST be used as the first sub-TLV inside TLV 43.

Type	Length	Value
43	n	8, 3, 0xFFFFFF, followed by general extension information

The following sub-TLVs are defined only as part of the general extension information. The type values may be re-defined for any purpose as part of a vendor-specific information encoding.

C.1.1.17.1.1 CM load balancing policy ID

The CMTS load balancing algorithm uses this config file setting as the CM load balancing policy ID. If present, this value overrides the default group policy assigned by the CMTS (see clause 11.4.5.6). This configuration setting should only appear once in a configuration file. This configuration setting MUST only be used in configuration files and REG-REQ messages, and MUST NOT be nested inside a packet classification configuration setting, a service flow configuration setting, or a service flow response.

Type	Length	Value
43.1	4	Policy ID

C.1.1.17.1.2 CM load balancing priority

This config file setting is the CM load balancing priority to be used by the CMTS load balancing algorithm. If present, this value overrides the default priority assigned by the CMTS (see clause 11.4.5.6). This configuration setting should only appear once in a configuration file. This configuration setting MUST only be used in configuration files and REG-REQ messages, and MUST NOT be nested inside a packet classification configuration setting, a service flow configuration setting, or a service flow response.

Type	Length	Value
43.2	4	Priority

C.1.1.17.1.3 CM load balancing group ID

This config file setting is the restricted load balancing group ID defined at the CMTS. If present, this value overrides the general load balancing group. If no restricted load balancing group is defined that matches this group ID, the value is ignored by the CMTS (see clause 11.4.5.6). This configuration setting should only appear once in a configuration file. This configuration setting MUST only be used in configuration files and REG-REQ messages, and MUST NOT be nested inside a packet classification configuration setting, a service flow configuration setting, or a service flow response.

Type	Length	Value
43.3	4	Group ID

C.1.1.17.1.4 CM ranging class ID extension

This config file setting is the CM ranging class ID extension to be defined by the MSO. These bits will be prepended to the CM's default ranging class ID as the most significant bits of the 32-bit ranging class ID value. These bits will be sent in the REG-REQ as part of the CM's ranging class ID in the modem capabilities field. If the TLV is not included in the configuration file, the CM will use zero for this value. These bits allow the user to define special device classes that could be used to give those devices, or service types, preferential treatment with respect to ranging after a massive outage. After successful registration, the CM MUST store the entire 32-bit value in a non-volatile memory and use it for ranging decisions after a reboot or a re-init MAC event.

Type	Length	Value
43.4	2	Extended ID

C.1.1.17.1.5 L2VPN encoding

The L2VPN encoding parameter is a multi-part encoding that configures how the CMTS performs layer 2 virtual private network bridging for CPE packets. The subtypes of the L2VPN encoding are specified in [L2VPN]. The CMTS MAY support the DOCSIS Layer 2 Virtual Private Network feature as defined in [L2VPN]. If the L2VPN feature is not supported, the CMTS MUST ignore the information in the L2VPN configuration setting.

GEI-Encapsulated L2VPN encoding:

Type	Length	Value
43.5	n	L2VPN encoding subtype/length/value tuples

C.1.1.17.2 Vendor-specific information

Vendor-specific configuration information, if present, is encoded in the DOCSIS extension field (code 43) using the vendor ID field (refer to clause C.1.3.2) to specify which TLV tuples apply to which vendor's products.

Type	Length	Value
43	n	Per vendor definition

Example:

Configuration with vendor A specific fields and vendor B specific fields:

VSIF (43) + n (number of bytes inside this VSIF)

8 (vendor ID type) + 3 (length field) + vendor ID of vendor A

Vendor A – specific type #1 + length of the field + value #1

Vendor A – specific type #2 + length of the field + value #2

VSIF (43) + m (number of bytes inside this VSIF)

8 (vendor ID type) + 3 (length field) + vendor ID of vendor B

Vendor B – specific type + length of the field + value

C.1.1.18 Subscriber management TLVs

The information in these TLVs is not used by the CM; rather, the information is used by the CMTS to populate the subscriber management MIB for this CM.

If present in the configuration file, the CM MUST include these TLVs in the subsequent REG-REQ to be used by the CMTS to populate the subscriber management MIB for this CM. If present in the configuration file, the CM MUST include these TLVs in the CMTS MIC.

C.1.1.18.1 Subscriber management control

This three-byte field provides control information to the CMTS for the subscriber management MIB. The first two bytes represent the number of IP addresses permitted behind the CM. The third byte is used for control fields.

Type	Length	Value
35	3	Byte 1, 2: docsSubMgtCpeControlMaxCpeIP (low-order 10 bits) Byte 3, bit 0: docsSubMgtCpeControlActive Byte 3, bit 1: docsSubMgtCpeControlLearnable Byte 3, bits #2-7: reserved, must be set to zero

C.1.1.18.2 Subscriber management CPE IP table

This field lists the IP addresses used to populate docsSubMgtCpeIpTable in the subscriber management MIB at the CMTS.

Type	Length	Value
36	N (multiple of 4)	Ipa1, Ipa2, Ipa3, Ipa4

C.1.1.18.3 Subscriber management filter groups

The subscriber management MIB allows an upstream and downstream filter group to be assigned to a CM and its associated CPEs and one or more embedded service/application functional entities (eSAFEs). These filter groups are encoded in the configuration file in a single TLV as follows:

Type	Length	Value
37	N (multiples of 4 bytes)	Bytes 1, 2: docsSubMgtSubFilterDownstream group Bytes 3, 4: docsSubMgtSubFilterUpstream group Bytes 5, 6: docsSubMgtCmFilterDownstream group Bytes 7, 8: docsSubMgtCmFilterUpstream group Bytes 9, 10: docsSubMgtPsFilterDownstream group Bytes 11, 12: docsSubMgtPsFilterUpstream group Bytes 13, 14: docsSubMgtMtaFilterDownstream group Bytes 15, 16: docsSubMgtMtaFilterUpstream group Bytes 17, 18: docsSubMgtStbFilterDownstream group Bytes 19, 20: docsSubMgtStbFilterUpstream group

C.1.1.19 Enable 2.0 mode

This configuration setting enables/disables DOCS 2.0 mode for a CM registering with a DOCS 2.0 CMTS.

The default value of this parameter MUST be 1: 2.0 mode enabled.

Type	Length	Value
39	1	0: Disable 1: Enable

C.1.1.20 Enable test modes

This configuration setting enables/disables certain test modes for a CM which supports test modes. The definition of the test modes is beyond the scope of this Recommendation.

If this TLV is not present, the default value MUST be 0: Test modes disabled.

Type	Length	Value
40	1	0: Disable 1: Enable

C.1.1.21 Downstream channel list

A list of receive frequencies to which the CM is allowed to tune during scanning operations. When the downstream channel list is provided in a configuration file, the CM MUST NOT attempt to establish communications using a downstream channel that is absent from this list unless specifically directed to do so by the CMTS. When both the downstream channel list and the downstream frequency configuration setting (clause C.1.1.1) are included in the configuration file, the CM MUST ignore the downstream frequency configuration setting. This list can override the last operational channel stored in NVRAM as defined in clause 11.2.1. The CM MUST retain and employ this list of channels whenever the CM performs a re-initialize MAC or continue scanning operation. The CM MUST replace or remove the list by subsequent config file downloads. Upon power cycle, the CM MUST NOT enforce a previously learned downstream channel list. The CM MAY, however, remember this list as an aid to downstream channel acquisition.

Type	Length	Value
41	n	List of allowed Rx frequencies

The list of allowed downstream frequencies is composed of an ordered series of sub-TLVs (single downstream channel, downstream frequency range, and default scanning) as defined below. When scanning for a downstream channel (except after a power-cycle), the CM MUST scan through this ordered list and attempt to establish communications on the specified channel(s). The scanning is initialized as follows:

- If the CM is in an operational state, and then undergoes a re-initialize MAC operation (except due to a dynamic channel change), it MUST first scan the last operational frequency and then restart scanning at the beginning of the ordered list.
- If, while scanning this ordered list, the CM fails to become operational and is forced to re-initialize MAC, the CM MUST continue scanning from the next applicable frequency in the ordered list.
- If it reaches the default scanning TLV (TLV 41.3) in the config file, the CM begins its default scanning algorithm, completing initial ranging and DHCP and receiving a new config file via TFTP on the first valid frequency it sees. If the new config file does not contain TLV 41, the CM MUST continue with registration. If the new config file contains TLV 41, the CM MUST confirm that the current downstream frequency is explicitly listed in the downstream channel list. If the current downstream channel is not explicitly listed in the downstream channel list, the CM MUST NOT register on the current downstream channel and MUST restart scanning according to the downstream channel list contained in the configuration file.

Upon reaching the end of the list, the CM MUST begin again with the first sub-TLV in the list. The CM MUST be capable of processing a downstream channel list that contains up to 16 sub-TLVs.

This configuration setting MAY appear multiple times. If this configuration setting appears multiple times, all sub-TLVs MUST be considered part of a single downstream channel list in the order in which they appear in the configuration file. In other words, the sub-TLVs from the first instance of this configuration setting would comprise the first entries in the ordered series, the second instance would comprise the next entries, etc.

C.1.1.21.1 Single downstream channel

Upon reaching this sub-TLV in the downstream channel list, the CM MUST attempt to acquire a downstream signal on the specified frequency for a period of time specified by timeout. If a signal is acquired, and it is determined to be unusable by the CM, the CM MAY move on to the next sub-TLV in the downstream channel list without waiting for the timeout to expire.

The CM MUST be capable of processing a downstream channel list that contains multiple single downstream frequency TLVs.

Type	Length	Value
41.1	6 or 10	

C.1.1.21.1.1 Single downstream channel timeout

Timeout is specified in seconds (unsigned). A value of 0 for timeout means no timeout, i.e., the CM attempts to acquire a signal on the specified frequency and if unsuccessful moves immediately to the next sub-TLV in the downstream channel list. This is an optional parameter in a single downstream channel TLV. If the single downstream channel timeout is omitted, the CM MUST use a default time out of 0.

Type	Length	Value
41.1.1	2	Timeout

C.1.1.21.1.2 Single downstream channel frequency

Single downstream channel frequency is a required parameter in each single downstream channel TLV; the CM MUST ignore any single downstream channel TLV which lacks this parameter. The DSFrequency MUST be a multiple of 62'500 Hz.

Type	Length	Value
41.1.2	4	DSFrequency

C.1.1.21.2 Downstream frequency range

Upon reaching this sub-TLV in the downstream channel list, the CM MUST begin scanning with DSFrequencyStart and progress in steps as indicated by DSFrequencyStepSize until reaching DSFrequencyEnd, and then repeat for a period of time specified by timeout. If the value of timeout is less than the time necessary for the CM to complete one full scan of all channels in the downstream frequency range, the CM MUST complete one full scan and then move on to the next sub-TLV in the downstream channel list. Note, DSFrequencyEnd may be less than DSFrequencyStart, which indicates scanning in the downward direction. If a signal has been acquired on all available channels between DSFrequencyStart and DSFrequencyEnd (inclusive), and all channels have been determined to be unusable by the CM, the CM MAY move on to the next sub-TLV in the downstream channel list without waiting for the Timeout to expire.

The CM MUST be capable of processing a downstream channel list that contains multiple downstream frequency range TLVs.

Type	Length	Value
41.2	18 or 22	

C.1.1.21.2.1 Downstream frequency range timeout

Timeout is specified in seconds (unsigned). A value of 0 for timeout means no timeout, i.e., the CM attempts to acquire a signal once on each frequency within the defined range, and if unsuccessful moves immediately to the next sub-TLV in the downstream channel list. This is an optional parameter in a downstream frequency range TLV. If the downstream frequency range timeout is omitted, the CM MUST use a default time out of 0.

Type	Length	Value
41.2.1	2	Timeout

C.1.1.21.2.2 Downstream frequency range start

Downstream frequency range start is a required parameter in each downstream frequency range TLV; the CM MUST ignore any downstream frequency range TLV that lacks this parameter. Downstream frequency range start MUST be a multiple of 62'500 Hz.

Type	Length	Value
41.2.2	4	DSFrequencyStart

C.1.1.21.2.3 Downstream frequency range end

Downstream frequency range end is a required parameter in each downstream frequency range TLV; the CM MUST ignore any downstream frequency range TLV that lacks this parameter. Downstream frequency range end MUST be a multiple of 62'500 Hz.

Type	Length	Value
41.2.3	4	DSFrequencyEnd

C.1.1.21.2.4 Downstream frequency range step size

Downstream frequency range step size is a required parameter in each downstream frequency range TLV; the CM MUST ignore any downstream frequency range TLV which lacks this parameter. Downstream frequency range step size specifies the increments in Hz by which the CM MUST scan through the downstream frequency range.

The CM MUST support a minimum downstream frequency step size of 6'000'000 Hz. The CM MAY support downstream frequency step sizes less than 6'000'000 Hz.

Type	Length	Value
41.2.4	4	DSFrequencyStepSize

C.1.1.21.3 Default scanning

Upon reaching this sub-TLV in the downstream channel list, the CM MUST begin scanning according to its default scanning algorithm (which may be vendor-dependent), and repeat for a period of time specified by timeout. When the CM reaches a valid downstream frequency during default scanning, the CM completes initial ranging and DHCP, and receives a new config file, via TFTP. If the config file does not contain TLV 41, the CM continues with registration. If the config file contains TLV 41 and the current downstream channel is not explicitly listed in the downstream channel list, the CM restarts scanning according to the downstream channel list contained in the configuration file.

Timeout is specified in seconds (unsigned). If the value of timeout is less than the time necessary for the CM to complete one full scan of all channels in the default scanning algorithm, the CM MUST complete one full scan and move on to the next sub-TLV in the downstream channel list. A value of 0 for timeout means no timeout, i.e., the CM scans all available frequencies once, then moves to the next sub-TLV in the downstream channel list.

The CM MUST be capable of processing a downstream channel list that contains multiple default scanning TLVs.

Type	Length	Value
41.3	2	Timeout

C.1.1.21.4 Examples illustrating usage of the downstream channel list

Assume that a modem has been provisioned to receive a config file with a downstream channel list consisting of several single downstream channel (TLV 41.1) entries, a downstream frequency range (TLV 41.2) entry, a default scanning (TLV 41.3) entry, and no timeout entries.

When the modem first boots up, it locks onto the first downstream channel it can find and goes through initial ranging. After completing the ranging process, the modem downloads the config file with the downstream channel list. The modem then checks its current downstream frequency against the frequencies explicitly listed in the single downstream channel (TLV 41.1) entries and the downstream frequency range entry (TLV 41.2) of the downstream channel list, ignoring the default scan (TLV 41.3) entry at this point. If the current channel is not explicitly in the single downstream channel entries in the list or within the downstream frequency range entry in the list, the modem moves to the first sub-TLV in the TLV 41 list and attempts to lock onto that channel. If the modem is able to lock onto that frequency, it again tries to range and download a config file. Assuming that the modem receives the same config file, the modem would then proceed with registration.

If the modem is not able to lock on the first sub-TLV in the downstream channel list, it moves onto the next entry in the list and so on. If it reaches the downstream frequency range TLV, it will begin scanning at the downstream frequency range start, incrementing the frequency by the downstream frequency step size, and ending at the downstream frequency range end. If the CM finds a valid downstream frequency within the downstream frequency range, the modem ranges and downloads a config file. Assuming that the config file has not changed, the modem continues with registration on that channel.

However, if it reaches the default scanning sub-TLV without successfully registering, the modem starts its "default scan" process. If, during the course of its default scan, the modem finds a DS channel that it can lock onto, is able to complete ranging, and is able to download a config file, it will do so. However, at that point, the modem once again checks which downstream channels are explicitly in the list and acts accordingly.

As a second, less likely example, assume that a modem has been provisioned to receive a config file with a downstream channel list containing only a default scanning (TLV 41.3) entry. When the modem first boots up, it locks onto the first downstream channel it can find and goes through initial ranging. After completing the ranging process, the modem downloads the config file with the downstream channel list. Since the default scanning is the only parameter in the downstream channel list, the downstream frequency on which the CM locked is not explicitly included, so the CM continues to scan according to its algorithm. The CM will not register on a channel until it receives a config file with a downstream frequency explicitly listed in the downstream channel list or a config file with no downstream channel list.

C.1.1.22 Downstream unencrypted traffic (DUT) filtering encoding

This parameter enables the CM to perform downstream unencrypted traffic filtering as described in the DOCSIS layer 2 virtual private network specification [L2VPN]. If the CM does not support the DUT filtering capability, it MUST ignore the DUT filtering encoding TLV.

DUT filtering encoding:

Type	Length	Value
45	Length/value tuples are specified in [L2VPN]	

C.1.2 Configuration-file-specific settings

These settings are found only in the configuration file. They MUST NOT be forwarded to the CMTS in the registration request.

C.1.2.1 End-of-data marker

This is a special marker for end of data. It has no length or value fields.

Type
255

C.1.2.2 Pad configuration setting

This has no length or value fields and is only used following the end of data marker to pad the file to an integral number of 32-bit words.

Type
0

C.1.2.3 Software upgrade filename

The filename of the software upgrade file for the CM. The filename is a fully qualified directory-path name. The file is expected to reside on a TFTP server identified in a configuration setting option defined in clause D.2.2. See also clause 12.1.

Type	Length	Value
9	n	File name

C.1.2.4 SNMP write-access control

This object makes it possible to disable SNMP "set" access to individual MIB objects. Each instance of this object controls access to all of the writeable MIB objects whose object ID (OID) prefix matches. This object may be repeated to disable access to any number of MIB objects.

Type	Length	Value
10	n	OID prefix plus control flag

Where n is the size of the ASN.1 basic encoding rules [ITU-T X.690] encoding of the OID prefix plus one byte for the control flag.

The control flag may take values:

- 0: Allow write-access;
- 1: Disallow write-access.

Any OID prefix may be used. The null OID 0.0 may be used to control access to all MIB objects. (The OID 1.3.6.1 will have the same effect.)

When multiple instances of this object are present and overlap, the longest (most specific) prefix has precedence. Thus, one example might be

- someTable: Disallow write-access;
- someTable.1.3: Allow write-access.

This example disallows access to all objects in someTable except for some Table.1.3.

C.1.2.5 SNMP MIB object

This object allows arbitrary SNMP MIB objects to be set via the TFTP-Registration process.

Type	Length	Value
11	n	Variable binding

The value is an SNMP VarBind as defined in [RFC 1157]. The VarBind is encoded in ASN.1 basic encoding rules, just as it would be if part of an SNMP set request.

The cable modem MUST treat this object as if it were part of an SNMP set request with the following caveats:

- It MUST treat the request as fully authorized (it cannot refuse the request for lack of privilege).
- SNMP write-control provisions (see previous clause) do not apply.
- No SNMP response is generated by the CM.

This object MAY be repeated with different VarBinds to "set" a number of MIB objects. All such sets MUST be treated as if simultaneous.

Each VarBind MUST be limited to 255 bytes.

C.1.2.6 CPE Ethernet MAC address

This object configures the CM with the Ethernet MAC address of a CPE device (see clause 5.1.2.3.1). This object may be repeated to configure any number of CPE device addresses.

Type	Length	Value
14	6	Ethernet MAC address of CPE

C.1.2.7 Software upgrade TFTP server

The IP address of the TFTP server on which the software upgrade file for the CM resides. See clauses 12.1 and C.1.2.3.

Type	Length	Value
21	4	ip1, ip2, ip3, ip4

C.1.2.8 SnmpV3 kickstart value

Compliant CMs MUST understand the following TLV and its sub-elements and be able to kickstart SNMPv3 access to the CM regardless of whether the CMs are operating in 1.0 mode or 1.1 mode.

Type	Length	Value
34	n	Composite

Up to 5 of these objects may be included in the configuration file. Each results in an additional row being added to the usmDHKickstartTable and the usmUserTable and results in an agent public number being generated for those rows.

C.1.2.8.1 SnmpV3 kickstart security name

Type	Length	Value
34.1	2-16	UTF8 encoded security name

For the ASCII character set, the UTF8 and the ASCII encodings are identical. Normally, this will be specified as one of the DOCSIS built-in USM users, e.g., "docsisManager", "docsisOperator", "docsisMonitor", "docsisUser". The security name is NOT zero terminated. This is reported in the usmDHKickStartTable as usmDHKickStartSecurityName and in the usmUserTable as usmUserName and usmUserSecurityName.

C.1.2.8.2 SnmpV3 kickstart manager public number

Type	Length	Value
34.2	n	Manager's Diffie-Helman public number expressed as an octet string

This number is the Diffie-Helman public number derived from a privately generated (by the manager or operator) random number and transformed according to [RFC 2786]. This is reported in the usmDHKickStartTable as usmKickStartMgrPublic. When combined with the object reported in the same row as usmKickStartMyPublic, it can be used to derive the keys in the related row in the usmUserTable.

C.1.2.9 Manufacturer code verification certificate

The manufacturer code verification certificate (M-CVC) for secure software downloading specified by Appendix D of [DOCS8]. The CM config file MUST contain this M-CVC and/or C-CVC defined in clause C.1.2.10 in order to allow the 1.1-compliant CM to download the code file from the TFTP server whether or not the CM is provisioned to run with BPI, BPI+, or none of them. See Appendix D of [DOCS8] for details.

Type	Length	Value
32	n	Manufacturer CVC (DER-encoded ASN.1)

If the length of the M-CVC exceeds 254 bytes, the M-CVC MUST be fragmented into two or more successive type 32 elements. Each fragment, except the last, MUST be 254 bytes in length. The CM reconstructs the M-CVC by concatenating the contents (value of the TLV) of successive type 32 elements in the order in which they appear in the configuration file. For example, the first byte following the length field of the second type 32 element is treated as if it immediately follows the last byte of the first type 32 element.

C.1.2.10 Co-signer code verification certificate

The co-signer code verification certificate (C-CVC) for secure software downloading specified by Appendix D of [DOCS8]. The CM configuration file MUST contain this C-CVC and/or M-CVC defined in clause C.1.2.9 in order to allow the 1.1-compliant CM to download the code file from TFTP server whether or not the CM is provisioned to run with BPI, BPI+, or none of them. See Appendix D of [DOCS8] for details.

Type	Length	Value
33	n	Co-signer CVC (DER-encoded ASN.1)

If the length of the C-CVC exceeds 254 bytes, the C-CVC MUST be fragmented into two or more successive type 33 elements. Each fragment, except the last, MUST be 254 bytes in length. The CM reconstructs the C-CVC by concatenating the contents (value of the TLV) of successive type 33 elements in the order in which they appear in the config file. For example, the first byte following the length field of the second Type 33 element is treated as if it immediately follows the last byte of the first Type 33 element.

C.1.2.11 SNMPv3 notification receiver

This TLV specifies a network management station that will receive notifications from the modem when it is in co-existence mode. Up to 10 of these elements may be included in the configuration file.

Type	Length	Value
38	n	Composite

C.1.2.11.1 SNMPv3 notification receiver IP address

This sub-TLV specifies the IP address of the notification receiver.

Type	Length	Value
38.1	4	ip1, ip2, ip3, ip4

C.1.2.11.2 SNMPv3 notification receiver UDP port number

This sub-TLV specifies the UDP port number of the notification receiver. If this sub-TLV is not present, the default value of 162 should be used.

Type	Length	Value
38.2	2	UDP port number

C.1.2.11.3 SNMPv3 notification receiver trap type

Type	Length	Value
38.3	2	Trap type

This sub-TLV specifies the type of trap to send. The trap type may take values:

- 1 = SNMP v1 trap in an SNMP v1 packet;
- 2 = SNMP v2c trap in an SNMP v2c packet;
- 3 = SNMP inform in an SNMP v2c packet;
- 4 = SNMP v2c trap in an SNMP v3 packet;
- 5 = SNMP inform in an SNMP v3 packet.

C.1.2.11.4 SNMPv3 notification receiver timeout

This sub-TLV specifies the timeout value to use when sending an inform message to the notification receiver.

Type	Length	Value
38.4	2	Time in milliseconds

C.1.2.11.5 SNMPv3 notification receiver retries

This sub-TLV specifies the number of times to retry sending an inform message if an acknowledgement is not received.

Type	Length	Value
38.5	2	Number of retries

C.1.2.11.6 SNMPv3 notification receiver filtering parameters

This sub-TLV specifies the ASN.1 formatted object identifier of the snmpTrapOID value that identifies the notifications to be sent to the notification receiver. SNMP V3 allows the specification of which trap OIDs are to be sent to a trap receiver. This object specifies the OID of the root of a trap filter sub-tree. All traps with a trap OID contained in this trap filter sub-tree MUST be sent to the trap receiver. This object starts with the ASN.1 universal type 6 (object identifier) byte, then the ASN.1 length field, then the ASN.1 encoded object identifier components.

Type	Length	Value
38.6	n	Filter OID

C.1.2.11.7 SNMPv3 notification receiver security name

This sub-TLV specifies the V3 security name to use when sending a V3 notification. This sub-TLV is only used if trap type is set to 4 or 5. This name must be a name specified in a configuration file TLV Type 34 as part of the DH kickstart procedure. The notifications will be sent using the authentication and privacy keys calculated by the modem during the DH kickstart procedure.

This sub-TLV is not required for trap type = 1, 2 or 3 above. If it is not supplied for a trap type of 4 or 5, then the V3 notification will be sent in the noAuthNoPriv security level using the security name "@config".

Type	Length	Value
38.7	n	Security name

C.1.2.12 Multicast MAC address

This object configures the CM with a static multicast MAC address that is being provisioned into the CM. This object may be repeated to configure any number of static multicast MAC addresses. The CM MUST support a minimum of 10 static multicast MAC addresses. The CM MUST forward any multicast frames that match the static multicast MAC address from the cable network to the CMCI subject to the provisions of clause 5.1.2.3.2. IGMP has no impact on this forwarding.

Type	Length	Value
42	6	Static multicast MAC address

C.1.3 Registration request/response-specific encodings

These encodings are not found in the configuration file, but are included in the registration request and option 60 of the DHCP request. Some encodings are also used in the registration response.

The CM MUST include all modem capabilities encodings that are subject to negotiation with the CMTS in its registration request. Modem capabilities encodings that are not subject to negotiation with the CMTS are explicitly stated in the description of the particular modem capability. The CMTS MUST include modem capabilities in the registration response.

C.1.3.1 Modem capabilities encoding

The value field describes the capabilities of a particular modem, i.e., implementation-dependent limits on the particular feature or number of features that the modem can support. It is composed from a number of encapsulated type/length/value fields. The encapsulated subtypes define the specific capabilities for the modem in question. Note that the subtype fields defined are only valid within the encapsulated capabilities configuration setting string.

Type	Length	Value
5	n	

The set of possible encapsulated fields is described below.

All these capabilities are to be included in both the registration request and option 60 of the DHCP request unless the description of the capability explicitly prohibits this.

C.1.3.1.1 Concatenation support

If the value field is a 1, the CM requests concatenation support from the CMTS.

Type	Length	On/Off
5.1	1	1 or 0

C.1.3.1.2 DOCS version

DOCS version of this modem.

Type	Length	Value
5.2	1	0: DOCS v1.0 1: DOCS v1.1 2: DOCS v2.0 3-255: Reserved

If this tuple is absent, the CMTS MUST assume DOCS v1.0 operation. The absence of this tuple or the value 'DOCSIS 1.0' does not necessarily mean the CM only supports DOCS 1.0 functionality – the CM MAY indicate it supports other individual capabilities with other modem capability encodings (refer to clause G.2). This capability is provided by the CM for the benefit of the CMTS, the operation of the CM is not affected by the value returned by the CMTS.

C.1.3.1.3 Fragmentation support

If the value field is a 1, the CM requests fragmentation support from the CMTS.

Type	Length	Value
5.3	1	1 or 0

C.1.3.1.4 Payload header suppression support

If the value field is a 1, the CM requests payload header suppression support from the CMTS.

Type	Length	Value
5.4	1	1 or 0

C.1.3.1.5 IGMP support

If the value field is a 1, the CM supports DOCS 1.1-compliant IGMP.

Type	Length	Value
5.5	1	1 or 0

NOTE – This CM capability is not subject to negotiation with the CMTS. The CM MUST include this capability in the DHCP request, but MUST NOT include this capability in the registration request. If a CMTS does receive this capability with in a registration request it MUST return the capability with the same value in the registration response.

C.1.3.1.6 Privacy support

The value indicates the BPI support of the CM.

Type	Length	Value
5.6	1	0: BPI support 1: BPI plus support 2-255: reserved

C.1.3.1.7 Downstream SAID support

This field shows the number of downstream SAIDs the modem can support.

Type	Length	Value
5.7	1	Number of downstream SAIDs the CM can support

If the number of SAIDs is 0, the modem can support only 1 SAID.

C.1.3.1.8 Upstream SID support

This field shows the number of upstream SIDs the modem can support.

Type	Length	Value
5.8	1	Number of upstream SIDs the CM can support

If the number of SIDs is 0, that means the modem can support only 1 SID.

C.1.3.1.9 Optional filtering support

This field shows the optional filtering support in the modem.

Type	Length	Value
5.9	1	Packet filtering support array bit #0: 802.1P filtering bit #1: 802.1Q filtering bit #2-7: Reserved, MUST be set to zero

NOTE – This CM capability is not subject to negotiation with the CMTS. The CM MUST include this capability in the DHCP request, but MUST NOT include this capability in the registration request. If a

CMTS does receive this capability with in a registration request, it MUST return the capability with the same value in the registration response.

C.1.3.1.10 Transmit equalizer taps per modulation interval

This field shows the maximal number of pre-equalizer taps per modulation interval T supported by the CM.

NOTE – All CMs MUST support T-spaced equalizer coefficients. CM support of 2 or 4 taps per modulation interval is optional. If this tuple is missing, it is implied that the CM only supports T-spaced equalizer coefficients. A CM MUST include this capability in the registration request and its value MUST be 1.

Type	Length	Value
5.10	1	1, 2 or 4

C.1.3.1.11 Number of transmit equalizer taps

This field shows the number of equalizer taps that are supported by the CM.

NOTE – All CMs MUST support an equalizer length of at least 8 symbols. CM support of up to 64 T-spaced, T/2-spaced or T/4-spaced taps is optional. If this tuple is missing, it is implied that the CM only supports an equalizer length of 8 taps. A CM MUST include this capability in the registration request and its value MUST be 24.

Type	Length	Value
5.11	1	8 to 64

C.1.3.1.12 DCC support

The value is the DCC support of the CM.

Type	Length	Value
5.12	1	0: DCC is not supported 1: DCC is supported

C.1.3.1.13 IP filters support

This field shows the number of IP filters that are supported by the CM.

Type	Length	Value
5.13	2	16-65535

NOTE – This CM capability is not subject to negotiation with the CMTS. The CM MUST include this capability in the DHCP request, but MUST NOT include this capability in the registration request. If a CMTS does receive this capability in a registration request it MUST return the capability with the same value in the registration response.

C.1.3.1.14 LLC filters support

This field shows the number of LLC filters that are supported by the CM.

Type	Length	Value
5.14	2	10-65535

NOTE – This CM capability is not subject to negotiation with the CMTS. The CM MUST include this capability in the DHCP request, but MUST NOT include this capability in the registration request. If a

CMTS does receive this capability in a registration request it MUST return the capability with the same value in the registration response.

C.1.3.1.15 Expanded unicast SID space

Indicates if the CM can support the expanded unicast SID space.

Type	Length	Value
5.15	1	0: Expanded unicast SID space is not supported 1: Expanded unicast SID space is supported

C.1.3.1.16 Ranging hold-off support

The CM indicates support for the ranging hold-off feature by reporting its ranging class ID in the value field. The low order 16 bits of the ranging class ID is comprised of a static bit map that indicates the device type. The CM sets the bits of the devices to 1 in the bit map. Only a stand-alone cable modem will set Bit #0. For example, a standalone CM would report a value of 1; a CM with a CableHome PS would report a value of 2; a CM with a PacketCable MTA, and an ePS would report a value of 6; an eSTB would report a value of 8 although it contained an eCM. Bits 16 through 31 are derived from the configuration file as described in clause C.1.1.17.1.4. The ranging class ID is not negotiable. The value field in the REG-RSP is ignored by the CM.

Type	Length	Value
5.16	4	Ranging class ID (bitmap) Bit #0: CM Bit #1: ePS Bit #2: eMTA Bit #3: DSG/eSTB Bits 4 through 15: Reserved Bits 16 through 31: CM ranging class ID extension

C.1.3.1.17 L2VPN capability

This capability indicates whether the CM is compliant with the DOCSIS layer 2 virtual private network feature as defined in [L2VPN]. The CM MAY support the DOCSIS layer 2 virtual private network feature as defined in [L2VPN].

Type	Length	Value
5.17	Length/value tuples are specified in [L2VPN]	

C.1.3.1.18 L2VPN eSAFE host capability

This capability encoding informs the CMTS of the type and MAC address of an eSAFE host embedded with a CM that supports the L2VPN feature. A CM MUST NOT include L2VPN eSAFE host capability TLV in the registration request or DHCP option 60 if it does not indicate support for [L2VPN] via the L2VPN capability encoding.

Type	Length	Value
5.18	Length/value tuples are specified in [L2VPN]	

C.1.3.1.19 Downstream unencrypted traffic (DUT) filtering

This capability indicates whether the CM supports the DUT filtering feature as defined in the DOCSIS layer 2 virtual private network specification [L2VPN]. The CM MAY support DUT filtering.

Type	Length	Value
5.19	Length/value tuples are specified in [L2VPN]	

C.1.3.2 Vendor ID encoding

The value field contains the vendor identification specified by the three-byte vendor-specific organization unique identifier of the CM MAC address.

The vendor ID MUST be used in a registration request, but MUST NOT be used as a stand-alone configuration file element. It MAY be used as a sub-field of the vendor-specific information field in a configuration file. When used as a sub-field of the vendor-specific information field, this identifies the vendor ID of the CMs that are intended to use this information. When the vendor ID is used in a registration request, then it is the vendor ID of the CM sending the request.

Type	Length	Value
8	3	v1, v2, v3

C.1.3.3 Modem IP address

For backward compatibility with DOCS v 1.0. Replaced by 'TFTP server provisioned modem address'.

Type	Length	Value
12	4	IP address

C.1.3.4 Service(s) not available response

This configuration setting MUST be included in the registration response message if the CMTS is unable or unwilling to grant any of the requested classes of service that appeared in the registration request. Although the value applies only to the failed service class, the entire registration request MUST be considered to have failed (none of the class-of-service configuration settings are granted).

Type	Length	Value
13	3	Class ID, type, confirmation code

Class ID is the class-of-service class from the request that is not available.

Type is the specific class-of-service object within the class that caused the request to be rejected.

Confirmation code (refer to clause C.4).

C.1.3.5 Vendor-specific capabilities

Vendor-specific data about the CM that is to be included in the REG-REQ, but which is not part of the configuration file, if present, MUST be encoded in the vendor-specific capabilities (VSC) (code 44) using the vendor ID field (refer to clause C.1.3.2) to specify which TLV tuples apply to which vendors products. The vendor ID MUST be the first TLV embedded inside VSC. If the first TLV inside VSIF is not a vendor ID, then the TLV MUST be discarded.

This configuration setting MAY appear multiple times. The same vendor ID MAY appear multiple times. There MUST NOT be more than one vendor ID TLV inside a single VSC.

Type	Length	Value
44	n	Per vendor definition

Example:

Configuration with vendor A-specific fields and vendor B-specific fields:

VSC (44) + n (number of bytes inside this VSC)

8 (vendor ID type) + 3 (length field) + vendor ID of vendor

Vendor-specific type #1 + length of the field + value #1

Vendor-specific type #2 + length of the field + value #2

C.1.4 Dynamic-service-message-specific encodings

These encodings are not found in the configuration file, nor in the registration request/response signalling. They are only found in DSA-REQ, DSA-RSP, DSA-ACK, DSC-REQ, DSC-RSP, DSC-ACK and DSD-REQ messages (see clauses 8.3.12 through 8.3.18).

C.1.4.1 HMAC-Digest

The HMAC-Digest setting is a keyed message digest. If privacy is enabled, the HMAC-digest attribute MUST be the final attribute in the dynamic service message's attribute list. The message digest is performed over the all of the dynamic service parameters (starting immediately after the MAC management message header and up to, but not including, the HMAC digest setting), other than the HMAC-digest, in the order in which they appear within the packet.

Inclusion of the keyed digest allows the receiver to authenticate the message. The HMAC-Digest algorithm, and the upstream and downstream key generation requirements are documented in [DOCS8].

This parameter contains a keyed hash used for message authentication. The HMAC algorithm is defined in [RFC 2104]. The HMAC algorithm is specified using a generic cryptographic hash algorithm. Baseline privacy uses a particular version of HMAC that employs the secure hash algorithm (SHA-1), defined in [FIPS 180-1].

A summary of the HMAC-digest attribute format is shown below. The fields are transmitted from left to right.

Type	Length	Value
27	20	A 160-bit (20-octet) keyed SHA hash

C.1.4.2 Authorization block

The authorization block contains an authorization "hint". The specifics of the contents of this "hint" are beyond the scope of this Recommendation, but include [ITU-T J.163].

The authorization block MAY be present in CM-initiated DSA-REQ and DSC-REQ messages, and CMTS-initiated DSA-RSP and DSC-RSP messages. This parameter MUST NOT be present in CMTS-initiated DSA-REQ and DSC-REQ messages, nor CM-initiated DSA-RSP and DSC-RSP messages.

The authorization block information applies to the entire content of the message. Thus, only a single authorization block per message MAY be present. The authorization block, if present, MUST be passed to the authorization module in the CMTS. The authorization block information is only processed by the authorization module.

Type	Length	Value
30	n	Sequence of n octets

C.1.4.3 Key sequence number

The value shows the key sequence number of the BPI+ authorization key that is used to calculate the HMAC-digest in case that the privacy is enabled.

Type	Length	Value
31	1	Auth key sequence number (0-15)

C.2 Quality-of-service-related encodings

C.2.1 Packet classification encodings

The following type/length/value encodings MUST be used in both the configuration file, registration messages and dynamic service messages to encode parameters for packet classification and scheduling. All multi-octet quantities are in network-byte order, i.e., the octet containing the most-significant bits is the first transmitted on the wire.

A classifier MUST contain at least one encoding from clause C.2.1.5, "IP packet classification encodings", clause C.2.1.6, "Ethernet LLC packet classification encodings", or C.2.1.7, "IEEE 802.1P/Q packet classification encodings".

The following configuration settings MUST be supported by all CMs that are compliant with this Recommendation. All CMTSs MUST support classification of downstream packets based on IP header fields (see clause C.2.1.5).

C.2.1.1 Upstream packet classification encoding

This field defines the parameters associated with an upstream classifier.

Note that the same subtype fields defined are valid for both the encapsulated upstream and downstream packet classification configuration setting string. These type fields are not valid in other encoding contexts.

Type	Length	Value
22	n	

C.2.1.2 Downstream packet classification encoding

This field defines the parameters associated with a downstream classifier.

Note that the same subtype fields defined are valid for both the encapsulated upstream and downstream flow classification configuration setting string. These type fields are not valid in other encoding contexts.

Type	Length	Value
23	n	

C.2.1.3 General packet classifier encodings

C.2.1.3.1 Classifier reference

The value of the field specifies a reference for the classifier. This value is unique per dynamic service message, configuration file or registration request message.

Type	Length	Value
[22/23].1	1	1-255

C.2.1.3.2 Classifier identifier

The value of the field specifies an identifier for the classifier. This value is unique to per service flow. The CMTS assigns the packet classifier identifier.

Type	Length	Value
[22/23].2	2	1-65'535

C.2.1.3.3 Service flow reference

The value of the field specifies a service flow reference that identifies the corresponding service flow.

In all packet classifier TLVs that occur in any message where the service flow ID is not known (e.g., CM-initiated DSA-REQ and REG-REQ), this TLV MUST be included. In all packet classifier TLVs that occur in a DSC-REQ and CMTS-initiated DSA-REQ messages, the service flow reference MUST NOT be specified.

Type	Length	Value
[22/23].3	2	1-65'535

C.2.1.3.4 Service flow identifier

The value of this field specifies the service flow ID that identifies the corresponding service flow.

In packet classifier TLVs where the service flow ID is not known, this TLV MUST NOT be included (e.g., CM-initiated DSA-REQ and REG-REQ). In packet classifier TLVs that occur in a DSC-REQ and CMTS-initiated DSA-REQ message, the service flow ID MUST be specified.

Type	Length	Value
[22/23].4	4	1-4'294'967'295

C.2.1.3.5 Rule priority

The value of the field specifies the priority for the classifier, which is used for determining the order of the classifier. A higher value indicates higher priority.

Classifiers that appear in configuration files and registration messages MAY have priorities in the range 0-255 with the default value 0. Classifiers that appear in DSA/DSC message MUST have priorities in the range 64-191, with the default value 64.

Type	Length	Value
[22/23].5	1	

C.2.1.3.6 Classifier activation state

The value of this field specifies whether this classifier should become active in selecting packets for the service flow. An inactive classifier is typically used with an AdmittedQosParamSet to ensure resources are available for later activation. The actual activation of the classifier depends both on this attribute and on the state of its service flow. If the service flow is not active, then the classifier is not used, regardless of the setting of this attribute.

Type	Length	Value
[22/23].6	1	0: Inactive 1: Active

The default value is 1: Activate the classifier.

C.2.1.3.7 Dynamic service change action

When received in a dynamic service change request, this indicates the action that should be taken with this classifier.

Type	Length	Value
[22/23].7	1	0: DSC add classifier 1: DSC replace classifier 2: DSC delete classifier

C.2.1.4 Classifier error encodings

This field defines the parameters associated with classifier errors.

Type	Length	Value
[22/23].8	n	

A classifier error encoding consists of a single classifier error parameter set, which is defined by the following individual parameters: errored parameter, confirmation code and error message.

The classifier error encoding is returned in REG-RSP, DSA-RSP and DSC-RSP messages to indicate the reason for the recipient's negative response to a classifier establishment request in a REG-REQ, DSA-REQ or DSC-REQ message.

On failure, the REG-RSP, DSA-RSP or DSC-RSP MUST include one classifier error encoding for at least one failed classifier requested in the REG-REQ, DSA-REQ or DSC-REQ message. A classifier error encoding for the failed classifier MUST include the confirmation code and errored parameter and MAY include an error message. If some classifier sets are rejected but other classifier sets are accepted, then classifier error encodings MUST be included for only the rejected classifiers. On success of the entire transaction, the RSP or ACK message MUST NOT include a classifier error encoding.

Multiple classifier error encodings may appear in a REG-RSP, DSA-RSP or DSC-RSP message, since multiple classifier parameters may be in error. A message with even a single classifier error encoding MUST NOT contain any other protocol classifier encodings (e.g., IP, 802.1P/Q).

A classifier error encoding MUST NOT appear in any REG-REQ, DSA-REQ or DSC-REQ messages.

C.2.1.4.1 Errored parameter

The value of this parameter identifies the subtype of a requested classifier parameter in error in a rejected classifier request. A classifier error parameter set MUST have exactly one errored parameter TLV within a given classifier error encoding.

Subtype	Length	Value
[22/23].8.1	n	Classifier encoding subtype in error

If the length is one, then the value is the single-level subtype where the error was found, e.g., 7 indicates an invalid change action. If the length is two, then the value is the multi-level subtype where the error was found, e.g., 9-2 indicates an invalid IP protocol value.

C.2.1.4.2 Error code

This parameter indicates the status of the request. A non-zero value corresponds to the confirmation code as described in clause C.4. A classifier error parameter set MUST have exactly one error code within a given classifier error encoding.

Subtype	Length	Value
[22/23].8.2	1	Confirmation code

A value of okay (0) indicates that the classifier request was successful. Since a classifier error parameter set applies only to errored parameters, this value MUST NOT be used.

C.2.1.4.3 Error message

This subtype is optional in a classifier error parameter set. If present, it indicates a text string to be displayed on the CM console and/or log that further describes a rejected classifier request. A classifier error parameter set MAY have zero or one error message subtypes within a given classifier error encoding.

Subtype	Length	Value
[22/23].8.3	n	Zero-terminated string of ASCII characters

NOTE 1 – The length n includes the terminating zero.

NOTE 2 – The entire classifier encoding message MUST have a total length of less than 256 characters.

C.2.1.5 IP packet classification encodings

This field defines the parameters associated with IP packet classification.

Type	Length	Value
[22/23].9	n	

C.2.1.5.1 IP type of service range and mask

The values of the field specify the matching parameters for the IP ToS byte range and mask. An IP packet with IP ToS byte value "ip-tos" matches this parameter if tos-low ≤ (ip-tos AND tos-mask) ≤ tos-high. If this field is omitted, then comparison of the IP packet ToS byte for this entry is irrelevant.

Type	Length	Value
[22/23].9.1	3	tos-low, tos-high, tos-mask

C.2.1.5.2 IP protocol

The value of the field specifies the matching value for the IP protocol field [b-RFC 1700]. If this parameter is omitted, then comparison of the IP header protocol field for this entry is irrelevant.

There are two special IP protocol field values: "256" matches traffic with any IP protocol value, and "257" matches both TCP and UDP traffic. An entry that includes an IP protocol field value greater than 257 MUST be invalidated for comparisons (i.e., no traffic can match this entry).

Type	Length	Value
[22/23].9.2	2	prot1, prot2

Valid range: 0-257

C.2.1.5.3 IP source address

The value of the field specifies the matching value for the IP source address. An IP packet with IP source address "ip-src" matches this parameter if src = (ip-src AND smask), where "smask" is the parameter from clause C.2.1.5.4. If this parameter is omitted, then comparison of the IP packet source address for this entry is irrelevant.

Type	Length	Value
[22/23].9.3	4	src1, src2, src3, src4

C.2.1.5.4 IP source mask

The value of the field specifies the mask value for the IP source address, as described in clause C.2.1.5.3. If this parameter is omitted, then the default IP source mask is 255.255.255.255.

Type	Length	Value
[22/23].9.4	4	smask1, smask2, smask3, smask4

C.2.1.5.5 IP destination address

The value of the field specifies the matching value for the IP destination address. An IP packet with IP destination address "ip-dst" matches this parameter if dst = (ip-dst AND dmask), where "dmask" is the parameter from clause C.2.1.5.6. If this parameter is omitted, then comparison of the IP packet destination address for this entry is irrelevant.

Type	Length	Value
[22/23].9.5	4	dst1, dst2, dst3, dst4

C.2.1.5.6 IP destination mask

The value of the field specifies the mask value for the IP destination address, as described in clause C.2.1.5.5, "IP destination address". If this parameter is omitted, then the default IP destination mask is 255.255.255.255.

Type	Length	Value
[22/23].9.6	4	dmask1, dmask2, dmask3, dmask4

C.2.1.5.7 TCP/UDP source port start

The value of the field specifies the low-end TCP/UDP source port value. An IP packet with TCP/UDP port value "src-port" matches this parameter if sportlow ≤ src-port ≤ sporthigh. If this parameter is omitted, then the default value of sportlow is 0. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23].9.7	2	sportlow1, sportlow2

C.2.1.5.8 TCP/UDP source port end

The value of the field specifies the high-end TCP/UDP source port value. An IP packet with TCP/UDP port value "src-port" matches this parameter if sportlow ≤ src-port ≤ sporthigh. If this parameter is omitted, then the default value of sporthigh is 65'535. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23].9.8	2	sporthigh1, sporthigh2

C.2.1.5.9 TCP/UDP destination port start

The value of the field specifies the low-end TCP/UDP destination port value. An IP packet with TCP/UDP port value "dst-port" matches this parameter if dportlow ≤ dst-port ≤ dporthigh. If this parameter is omitted, then the default value of dportlow is 0. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23].9.9	2	dportlow1, dportlow2

C.2.1.5.10 TCP/UDP destination port end

The value of the field specifies the high-end TCP/UDP destination port value. An IP packet with TCP/UDP port value "dst-port" matches this parameter if dportlow ≤ dst-port ≤ dporthigh. If this parameter is omitted, then the default value of dporthigh is 65'535. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23].9.10	2	dporthigh1, dporthigh2

C.2.1.6 Ethernet LLC packet classification encodings

This field defines the parameters associated with Ethernet LLC packet classification.

Type	Length	Value
[22/23].10	n	

C.2.1.6.1 Destination MAC address

The values of the field specifies the matching parameters for the MAC destination address. An Ethernet packet with MAC destination address "etherdst" matches this parameter if dst = (etherdst AND msk). If this parameter is omitted, then comparison of the Ethernet MAC destination address for this entry is irrelevant.

Type	Length	Value
[22/23].10.1	12	dst1, dst2, dst3, dst4, dst5, dst6, msk1, msk2, msk3, msk4, msk5, msk6

C.2.1.6.2 Source MAC address

The value of the field specifies the matching value for the MAC source address. If this parameter is omitted, then comparison of the Ethernet MAC source address for this entry is irrelevant.

Type	Length	Value
[22/23].10.2	6	src1, src2, src3, src4, src5, src6

C.2.1.6.3 Ethertype/DSAP/MacType

Type, eprot1 and eprot2 indicate the format of the layer 3 protocol ID in the Ethernet packet as follows:

If type = 0, the rule does not use the layer 3 protocol type as a matching criterion. If type = 0, then eprot1, eprot2 are ignored when considering whether a packet matches the current rule.

If type = 1, the rule applies only to frames which contain an Ethertype value. Ethertype values are contained in packets using the DEC-Intel-Xerox (DIX) encapsulation or the RFC 1042 subnetwork access protocol (SNAP) encapsulation formats. If type = 1, then eprot1, eprot2 gives the 16-bit value of the Ethertype that the packet must match in order to match the rule.

If type = 2, the rule applies only to frames using the IEEE 802.2 encapsulation format with a destination service (DSAP) other than 0xAA (which is reserved for SNAP). If type = 2, the lower 8 bits of the eprot1, eprot2, MUST match the DSAP byte of the packet in order to match the rule.

If type = 3, the rule applies only to MAC management messages (FC field 1100001x) with a "type" field of its MAC management message header (see clause 6.3.1) between the values of eprot1 and eprot2, inclusive. As exceptions, the following MAC management message types MUST NOT be classified, and are always transmitted on the primary service flow:

- Type 4: RNG-REQ;
- Type 6: REG-REQ;
- Type 7: REG-RSP;
- Type 14: REG-ACK.

If type = 4, the rule is considered a "catch-all" rule that matches all data PDU packets. The rule does not match MAC management messages. The values of eprot1 and eprot2 are ignored in this case.

If the Ethernet frame contains an 802.1P/Q tag header (i.e., Ethertype 0x8100), this object applies to the embedded Ethertype field within the 802.1P/Q header.

Other values of type are reserved. If this TLV is omitted, then comparison of either the Ethertype or IEEE 802.2 DSAP for this rule is irrelevant.

Type	Length	Value
[22/23].10.3	3	type, eprot1, eprot2

C.2.1.7 IEEE 802.1P/Q packet classification encodings

This field defines the parameters associated with IEEE 802.1P/Q packet classification.

Type	Length	Value
[22/23].11	n	

C.2.1.7.1 IEEE 802.1P user_priority

The values of the field specify the matching parameters for the IEEE 802.1P user_priority bits. An Ethernet packet with IEEE 802.1P user_priority value "priority" matches these parameters if pri-low ≤ priority ≤ pri-high. If this field is omitted, then comparison of the IEEE 802.1P user_priority bits for this entry is irrelevant.

If this parameter is specified for an entry, then Ethernet packets without IEEE 802.1Q encapsulation MUST NOT match this entry. If this parameter is specified for an entry on a CM that does not support forwarding of IEEE 802.1Q-encapsulated traffic, then this entry MUST NOT be used for any traffic.

Type	Length	Value
[22/23].11.1	2	pri-low, pri-high

Valid range: 0-7 for pri-low and pri-high.

C.2.1.7.2 IEEE 802.1Q vlan_ID

The value of the field specify the matching value for the IEEE 802.1Q vlan_id bits. Only the first (i.e., most-significant) 12 bits of the specified vlan_id field are significant; the final four bits MUST be ignored for comparison. If this field is omitted, then comparison of the IEEE 802.1Q vlan_id bits for this entry is irrelevant.

If this parameter is specified for an entry, then Ethernet packets without IEEE 802.1Q encapsulation MUST NOT match this entry. If this parameter is specified for an entry on a CM that does not support forwarding of IEEE 802.1Q-encapsulated traffic, then this entry MUST NOT be used for any traffic.

Type	Length	Value
[22/23].11.2	2	vlan_id1, vlan_id2

C.2.1.7.3 Vendor-specific classifier parameters

This allows vendors to encode vendor-specific classifier parameters using the DOCSIS extension field. The vendor ID MUST be the first TLV embedded inside vendor-specific classifier parameters. If the first TLV inside vendor-specific classifier parameters is not a vendor ID, then the TLV MUST be discarded (refer to clause C.1.1.17).

Type	Length	Value
[22/23].43	n	

C.2.2 Service flow encodings

The following type/length/value encodings MUST be used in the configuration file, registration messages, and dynamic service messages to encode parameters for service flows. All multi-octet quantities are in network-byte order, i.e., the octet containing the most-significant bits is the first transmitted on the wire.

The following configuration settings MUST be supported by all CMs which are compliant with this Recommendation.

C.2.2.1 Upstream service flow encodings

This field defines the parameters associated with upstream scheduling for a service flow. It is somewhat complex in that it is composed from a number of encapsulated type/length/value fields.

Note that the encapsulated upstream and downstream service flow configuration setting strings share the same subtype field numbering plan, because many of the subtype fields defined are valid for both types of configuration settings. These type fields are not valid in other encoding contexts.

Type	Length	Value
24	n	

C.2.2.2 Downstream service flow encodings

This field defines the parameters associated with downstream scheduling for a service flow. It is somewhat complex in that it is composed from a number of encapsulated type/length/value fields.

Note that the encapsulated upstream and downstream flow classification configuration setting strings share the same subtype field numbering plan, because many of the subtype fields defined are valid for both types of configuration settings except service flow encodings. These type fields are not valid in other encoding contexts.

Type	Length	Value
25	n	

C.2.2.3 General service flow encodings

C.2.2.3.1 Service flow reference

The service flow reference is used to associate a packet classifier encoding with a service flow encoding. A service flow reference is only used to establish a service flow ID. Once the service flow exists and has assigned a service flow ID, the service flow reference MUST no longer be used. The service flow reference is unique per configuration file, registration message exchange, or dynamic service add message exchange.

Type	Length	Value
[24/25].1	2	1-65'535

C.2.2.3.2 Service flow identifier

The service flow identifier is used by the CMTS as the primary reference of a service flow. Only the CMTS can issue a service flow identifier. It uses this parameterization to issue service flow identifiers in CMTS-initiated DSA-Requests and in its REG/DSA-Response to CM-initiated REG/DSA-Requests. The CM specifies the SFID of a service flow using this parameter in a DSC-REQ message. Both the CM and CMTS MAY use this TLV to encode service flow IDs in a DSD-REQ.

The configuration file MUST NOT contain this parameter.

Type	Length	Value
[24/25].2	4	1-4'294'967'295

C.2.2.3.3 Service identifier

The value of this field specifies the service identifier assigned by the CMTS to a service flow with a non-null AdmittedQosParamSet or ActiveQosParamSet. This is used in the bandwidth allocation

MAP to assign upstream bandwidth. This field MUST be present in CMTS-initiated DSA-REQ or DSC-REQ message related to establishing an admitted or active upstream service flow. This field MUST also be present in REG-RSP, DSA-RSP and DSC-RSP messages related to the successful establishment of an admitted or active upstream service flow. This field MUST NOT be present in settings related to downstream service flows; the service identifier only applies to upstream service flows.

Even though a service flow has been successfully admitted or activated (i.e., has an assigned service ID) the service flow ID MUST be used for subsequent DSx message signalling as it is the primary handle for a service flow. If a service flow is no longer admitted or active (via DSC-REQ) its service ID MAY be reassigned by the CMTS.

Subtype	Length	Value
[24/25].3	2	SID (low-order 14 bits)

C.2.2.3.4 Service class name

The value of the field refers to a predefined CMTS service configuration to be used for this service flow.

Type	Length	Value
[24/25].4	2 to 16	Zero-terminated string of ASCII characters

NOTE – The length includes the terminating zero.

When the service class name is used in a service flow encoding, it indicates that all the unspecified QoS parameters of the service flow need to be provided by the CMTS. It is up to the operator to synchronize the definition of service class names in the CMTS and in the configuration file.

C.2.2.3.5 Quality of service parameter set type

This parameter MUST appear within every service flow encoding, with the exception of service flow encodings in the DSD-REQ where the quality of service parameter set type has no value. It specifies the proper application of the QoS parameter set or service class name: to the provisioned set, the admitted set, and/or the active set. When two QoS parameter sets are the same, a multi-bit value of this parameter MAY be used to apply the QoS parameters to more than one set. A single message MAY contain multiple QoS parameter sets in separate type 24/25 service flow encodings for the same service flow. This allows specification of the QoS parameter sets when their parameters are different. Bit 0 is the LSB of the value field.

For every service flow that appears in a registration request or registration response message, there MUST be a service flow encoding that specifies a ProvisionedQosParamSet. This service flow encoding, or other service flow encoding(s), MAY also specify an admitted and/or active set.

Any service flow encoding that appears in a dynamic service message MUST NOT specify the ProvisionedQosParamSet.

Type	Length	Value
[24/25].6	1	Bit #0: Provisioned set Bit #1: Admitted set Bit #2: Active set

Table C.2 – Values used in REG-REQ and REG-RSP messages

Value	Messages
001	Apply to provisioned set only
011	Apply to provisioned and admitted set; and perform admission control
101	Apply to provisioned and active sets; perform admission control on admitted set in separate service flow encoding; and activate the service flow
111	Apply to provisioned, admitted and active sets; perform admission control; and activate this service flow

Table C.3 – Values used in REG-REQ, REG-RSP and dynamic service messages

Value	Messages
010	Perform admission control and apply to admitted set
100	Check against admitted set in separate service flow encoding; perform admission control if needed; activate this service flow; and apply to active set
110	Perform admission control and activate this service flow; apply parameters to both admitted and active sets

The value 000 is used only in dynamic service change messages. It is used to set the active and admitted sets to Null (see clause 10.1.7.4).

A CMTS MUST handle a single update to each of the active and admitted QoS parameter sets. The ability to process multiple service flow encodings that specify the same QoS parameter set is NOT required, and is left as a vendor-specific function. If a DSA/DSC contains multiple updates to a single QoS parameter set and the vendor does not support such updates, then the CMTS MUST reply with error code 2, reject-unrecognized-configuration-setting.

C.2.2.4 Service flow error encodings

This field defines the parameters associated with service flow errors.

Type	Length	Value
[24/25].5	n	

A service flow error encoding consists of a single service flow error parameter set that is defined by the following individual parameters: errored parameter, confirmation code and error message.

The service flow error encoding is returned in REG-RSP, DSA-RSP and DSC-RSP messages to indicate the reason for the recipient's negative response to a service flow establishment request in a REG-REQ, DSA-REQ or DSC-REQ message.

The service flow error encoding is returned in REG-ACK, DSA-ACK and DSC-ACK messages to indicate the reason for the recipient's negative response to the expansion of a service class name in a corresponding REG-RSP, DSA-RSP or DSC-RSP.

On failure, the REG-RSP, DSA-RSP or DSC-RSP MUST include one service flow error encoding for at least one failed service flow requested in the REG-REQ, DSA-REQ or DSC-REQ message. On failure, the REG-ACK, DSA-ACK or DSC-ACK MUST include one service flow error encoding for at least one failed service class name expansion in the REG-RSP, DSA-RSP or DSC-RSP message. A service flow error encoding for the failed service flow MUST include the confirmation code and errored parameter and MAY include an error message. If some service flow parameter sets are rejected but other service flow parameter sets are accepted, then service flow error encodings MUST be included for only the rejected service flow.

On success of the entire transaction, the RSP or ACK message MUST NOT include a service flow error encoding.

Multiple service flow error encodings MAY appear in a REG-RSP, DSA-RSP, DSC-RSP, REG-ACK, DSA-ACK or DSC-ACK message, since multiple service flow parameters may be in error. A message with even a single service flow error encoding MUST NOT contain any QoS parameters.

A service flow error encoding MUST NOT appear in any REG-REQ, DSA-REQ or DSC-REQ messages.

C.2.2.4.1 Errored parameter

The value of this parameter identifies the subtype of a requested service flow parameter in error in a rejected service flow request or service class name expansion response. A service flow error parameter set MUST have exactly one errored parameter TLV within a given service flow error encoding.

Subtype	Length	Value
[24/25].5.1	1	Service flow encoding subtype in error

C.2.2.4.2 Error code

This parameter indicates the status of the request. A non-zero value corresponds to the confirmation code as described in clause C.4. A service flow error parameter set MUST have exactly one error code within a given service flow error encoding.

Subtype	Length	Value
[24/25].5.2	1	Confirmation code

A value of okay (0) indicates that the service flow request was successful. Since a service flow error parameter set only applies to errored parameters, this value MUST NOT be used.

C.2.2.4.3 Error message

This subtype is optional in a service flow error parameter set. If present, it indicates a text string to be displayed on the CM console and/or log that further describes a rejected service flow request. A service flow error parameter set MAY have zero or one error message subtypes within a given service flow error encoding.

Subtype	Length	Value
[24/25].5.3	n	Zero-terminated string of ASCII characters

NOTE 1 – The length n includes the terminating zero.

NOTE 2 – The entire service flow encoding message MUST have a total length of less than 256 characters.

C.2.2.5 Common upstream and downstream quality-of-service parameter encodings

The remaining type 24 and 25 parameters are QoS parameters. Any given QoS parameter type MUST appear zero or one times per service flow encoding.

C.2.2.5.1 Traffic priority

The value of this parameter specifies the priority assigned to a service flow. Given two service flows identical in all QoS parameters besides priority, the higher priority service flow SHOULD be given lower delay and higher buffering preference. For otherwise non-identical service flows, the priority parameter SHOULD NOT take precedence over any conflicting service flow QoS parameter. The specific algorithm for enforcing this parameter is not mandated here.

For upstream service flows, the CMTS SHOULD use this parameter when determining precedence in request service and grant generation, and the CM MUST preferentially select contention request opportunities for priority request service IDs (refer to clause A.2.3) based on this priority and its request/transmission policy (refer to clause C.2.2.6.3).

Type	Length	Value
[24/25].7	1	0 to 7 – Higher numbers indicate higher priority

NOTE – The default priority is 0.

C.2.2.5.2 Maximum sustained traffic rate

This parameter is the rate parameter R of a token-bucket-based rate limit for packets. R is expressed in bits per second, and MUST take into account all MAC frame data PDU of the service flow from the byte following the MAC header HCS to the end of the CRC[28]. The number of bytes forwarded (in bytes) is limited during any time interval T by Max(T), as described in equation C.1:

$$\text{Max}(T) = T \times (R/8) + B \quad \text{(C.1)}$$

where the parameter B (in bytes) is the maximum traffic burst configuration setting (refer to clause C.2.2.5.3).

NOTE 1 – This parameter does not limit the instantaneous rate of the service flow.

NOTE 2 – The specific algorithm for enforcing this parameter is not mandated here. Any implementation that satisfies the above equation is conformant.

NOTE 3 – If this parameter is omitted or set to zero, then there is no explicitly enforced traffic rate maximum. This field specifies only a bound, not a guarantee that this rate is available.

C.2.2.5.2.1 Upstream maximum sustained traffic rate

For an upstream service flow, the CM MUST NOT request bandwidth exceeding the Max(T) requirement in equation C.1 during any interval T because this could force the CMTS to fill MAPs with deferred grants.

The CM MUST defer upstream packets that violate equation C.1 and "rate shape" them to meet the expression, up to a limit as implemented by vendor buffering restrictions.

The CMTS MUST enforce equation C.1 on all upstream data transmissions, including data sent in contention. The CMTS MAY consider unused grants in calculations involving this parameter. The CMTS MAY enforce this limit by any of the following methods:

a) discarding over-limit requests;

b) deferring (through zero-length grants) the grant until it is conforming to the allowed limit; or

c) discarding over-limit data packets.

A CMTS MUST report this condition to a policy module. If the CMTS is policing by discarding either packets or requests, the CMTS MUST allow a margin of error between the CM and CMTS algorithms.

Type	Length	Value
24.8	4	R (in bits per second)

[28] The payload size includes every PDU in a concatenated MAC frame.

C.2.2.5.2.2 Downstream maximum sustained traffic rate

For a downstream service flow, this parameter is only applicable at the CMTS. The CMTS MUST enforce equation C.1 on all downstream data transmissions. The CMTS MUST NOT forward downstream packets that violates equation C.1 in any interval T. The CMTS SHOULD "rate shape" the downstream traffic by enqueueing packets arriving in excess of equation C.1, and delay them until the equation can be met.

This parameter is not intended for enforcement on the CM.

Type	Length	Value
25.8	4	R (in bits per second)

C.2.2.5.3 Maximum traffic burst

The value of this parameter specifies the token bucket size B (in bytes) for this service flow as described in equation C.1. This value is calculated from the byte following the MAC header HCS to the end of the CRC[28].

The minimum value of B is 1522 bytes. If this parameter is omitted, the default value for B is 3044 bytes. This parameter has no effect unless a non-zero value has been provided for the maximum sustained traffic rate parameter. For an upstream service flow, if B is sufficiently less than the maximum concatenated burst parameter, then enforcement of the rate limit equation will limit the maximum size of a concatenated burst.

Type	Length	Value
[24/25].9	4	B (bytes)

The specific algorithm for enforcing this parameter is not mandated here. Any implementation that satisfies the above equation is conformant.

NOTE – The value of this parameter effects the trade-off between the data latency perceived by an individual application and the traffic engineering requirements of the network. A large value will tend to reduce the latency introduced by rate limiting for applications with bursty traffic patterns. A small value will tend to spread out the bursts of data generated by such applications, which may benefit traffic engineering within the network.

C.2.2.5.4 Minimum reserved traffic rate

This parameter specifies the minimum rate, in bit/s, reserved for this service flow. The CMTS SHOULD be able to satisfy bandwidth requests for a service flow up to its minimum reserved traffic rate. If less bandwidth than its minimum reserved traffic rate is requested for a service flow, the CMTS MAY reallocate the excess reserved bandwidth for other purposes. The aggregate minimum reserved traffic rate of all service flows MAY exceed the amount of available bandwidth. This value of this parameter is calculated from the byte following the MAC header HCS to the end of the CRC[28]. If this parameter is omitted, then it defaults to a value of 0 bit/s (i.e., no bandwidth is reserved for the flow by default).

This field is only applicable at the CMTS and MUST be enforced by the CMTS.

Type	Length	Value
[24/25].10	4	

[28] The payload size includes every PDU in a concatenated MAC frame.

NOTE – The specific algorithm for enforcing the value specified in this field is not mandated here.

C.2.2.5.5 Assumed minimum reserved rate packet size

The value of this field specifies an assumed minimum packet size (in bytes) for which the minimum reserved traffic rate will be provided. This parameter is defined in bytes and is specified as the bytes following the MAC header HCS to the end of the CRC[28]. If the service flow sends packets of a size smaller than this specified value, such packets will be treated as being of the size specified in this parameter for calculating the minimum reserved traffic rate.

The CMTS MUST apply this parameter to its minimum reserved traffic rate algorithm. This parameter is used by the CMTS to estimate the per-packet overhead of each packet in the service flow.

If this parameter is omitted, then the default value is CMTS implementation-dependent.

Type	Length	Value
[24/25].11	2	

C.2.2.5.6 Timeout for active QoS parameters

The value of this parameter specifies the maximum duration resources remain unused on an active service flow. If there is no activity on the service flow within this time interval, the CMTS MUST change the active and admitted QoS parameter sets to null. The CMTS MUST signal this resource change with a DSC-REQ to the CM.

Type	Length	Value
[24/25].12	2	Seconds

This parameter MUST be enforced at the CMTS and SHOULD NOT be enforced at the CM. The parameter is processed by the CMTS for every QoS set contained in registration messages and dynamic service messages. If the parameter is omitted, the default of 0 (i.e., infinite time-out) is assumed. The value specified for the active QoS set must be less than or equal to the corresponding value in the admitted QoS set, which must be less than or equal to the corresponding value in the provisioned/authorized QoS set. If the requested value is too large, the CMTS MAY reject the message or respond with a value less than that requested. If the registration or dynamic service message is accepted by the CMTS and acknowledged by the CM, the active MQoS time-out timer is loaded with the new value of the time-out. The timer is activated if the message activates the associated service flow. The timer is deactivated if the message sets the active QoS set to null.

C.2.2.5.7 Timeout for admitted QoS parameters

The value of this parameter specifies the duration that the CMTS MUST hold resources for a service flow's admitted QoS parameter set while they are in excess of its active QoS parameter set. If there is no DSC-REQ to activate the admitted QoS parameter set within this time interval, and there is no DSC to refresh the QoS parameter sets and restart the timeout (see clause 10.1.5.2), the resources that are admitted but not activated MUST be released, and only the active resources retained. The CMTS MUST set the admitted QoS parameter set equal to the active QoS parameter set for the service flow and initiate a DSC-REQ exchange with the CM to inform it of the change.

Type	Length	Value
[24/25].13	2	Seconds

[28] The payload size includes every PDU in a concatenated MAC frame.

This parameter MUST be enforced at the CMTS and SHOULD NOT be enforced at the CM. The parameter is processed by the CMTS for every QoS set contained in registration messages and dynamic service messages. If the parameter is omitted, the default of 200 seconds is assumed. A value of 0 means that the service flow can remain in the admitted state for an infinite amount of time and MUST NOT be timed out due to inactivity. However, this is subject to policy control by the CMTS. The value specified for the active QoS set must be less than or equal to the corresponding value in the admitted QoS set, which must be less than or equal to the corresponding value in the provisioned/authorized QoS set. If the requested value is too large, the CMTS MAY reject the message or respond with a value less than that requested. If the registration or dynamic service message containing this parameter is accepted by the CMTS and acknowledged by the CM, the admitted QoS timeout timer is loaded with the new value of the timeout. The timer is activated if the message admits resources greater than the active set. The timer is deactivated if the message sets the active QoS set and admitted QoS set equal to each other.

C.2.2.5.8 Vendor-specific QoS parameters

This allows vendors to encode vendor-specific QoS parameters using the DOCSIS extension field. The vendor ID MUST be the first TLV embedded inside vendor-specific QoS parameters. If the first TLV inside vendor-specific QoS parameters is not a vendor ID, then the TLV MUST be discarded (refer to clause C.1.1.17).

Type	Length	Value
[24/25].43	n	

C.2.2.6 Upstream-specific QoS parameter encodings

C.2.2.6.1 Maximum concatenated burst

The value of this parameter specifies the maximum concatenated burst (in bytes) that a service flow is allowed. This parameter is calculated from the FC byte of the concatenation MAC header to the last CRC in the concatenated MAC frame.

A value of 0 means there is no limit. If this parameter is omitted, the default value is 1522.

This field is only applicable at the CM. If defined, this parameter MUST be enforced at the CM.

NOTE 1 – This value does not include any physical layer overhead.

Type	Length	Value
24.14	2	

NOTE 2 – This applies only to concatenated bursts. It is legal and, in fact it may be useful, to set this smaller than the maximum Ethernet packet size. Of course, it is also legal to set this equal to or larger than the maximum Ethernet packet size.

NOTE 3 – The maximum size of a concatenated burst can also be limited by the enforcement of a rate limit, if the maximum traffic burst parameter is small enough, and by limits on the size of data grants in the UCD message.

C.2.2.6.2 Service flow scheduling type

The value of this parameter specifies which upstream scheduling service is used for upstream transmission requests and packet transmissions. If this parameter is omitted, then best effort service MUST be assumed.

This parameter is only applicable at the CMTS. If defined, this parameter MUST be enforced by the CMTS.

Type	Length	Value
24.15	1	0: Reserved 1: For undefined (CMTS implementation-dependent[a]) 2: For best effort 3: For non-real-time polling service 4: For real-time polling service 5: For unsolicited grant service with activity detection 6: For unsolicited grant service 7-255: Reserved for future use
[a] The specific implementation-dependent scheduling service type could be defined in the 24.43 vendor-specific information field.		

C.2.2.6.3 Request/transmission policy

The value of this parameter specifies which IUC opportunities the CM uses for upstream transmission requests and packet transmissions for this service flow, whether requests for this service flow may be piggybacked with data and whether data packets transmitted on this service flow can be concatenated, fragmented or have their payload headers suppressed. For UGS, it also specifies how to treat packets that do not fit into the UGS grant. See clause 10.2 for requirements related to settings of the bits of this parameter for each service flow scheduling type.

This parameter is required for all service flow scheduling types except best effort. If omitted in a best effort service flow QoS parameter set, the default value of zero MUST be used. Bit #0 is the LSB of the value field. Bits are set to 1 to select the behaviour defined below:

Type	Length	Value
24.16	4	Bit #0: The service flow MUST NOT use "all CMs" broadcast request opportunities. Bit #1: The service flow MUST NOT use priority request multicast request opportunities (refer to clause A.2.3). Bit #2: The service flow MUST NOT use request/data opportunities for requests. Bit #3: The service flow MUST NOT use request/data opportunities for data. Bit #4: The service flow MUST NOT piggyback requests with data. Bit #5: The service flow MUST NOT concatenate data. Bit #6: The service flow MUST NOT fragment data. Bit #7: The service flow MUST NOT suppress payload headers Bit #8: The service flow MUST drop packets that do not fit in the unsolicited grant size[a) b)] All other bits are reserved.

a) This bit only applies to service flows with the unsolicited grant service flow scheduling type, if this bit is set on any other service flow scheduling type it MUST be ignored.

b) Packets that classify to an unsolicited grant service flow and are larger than the grant size associated with that service flow are normally transmitted on the primary service flow. This parameter overrides that default behaviour.

NOTE – Data grants include both short and long data grants.

C.2.2.6.4 Nominal polling interval

The value of this parameter specifies the nominal interval (in units of microseconds) between successive unicast request opportunities for this service flow on the upstream channel. This parameter is typically suited for real-time and non-real-time polling service.

The ideal schedule for enforcing this parameter is defined by a reference time, t_0, with the desired transmission times $t_i = t_0 + i \times$ interval. The actual poll times, t'_i, MUST be in the range $t_i \leq t'_i \leq (t_i + \text{jitter})$, where interval is the value specified with this TLV, and jitter is tolerated poll jitter. The accuracy of the ideal poll times, t_i, are measured relative to the CMTS master clock used to generate timestamps (refer to clause 9.3).

This field is only applicable at the CMTS. If defined, this parameter MUST be enforced by the CMTS.

Type	Length	Value
24.17	4	µs

C.2.2.6.5 Tolerated poll jitter

The value in this parameter specifies the maximum amount of time that the unicast request interval may be delayed from the nominal periodic schedule (measured in microseconds) for this service flow.

The ideal schedule for enforcing this parameter is defined by a reference time, t_0, with the desired poll times $t_0 + i \times$ interval. The actual poll, t'_i, MUST be in the range $t_i \leq t'_i \leq t_i + \text{jitter}$, where jitter is the value specified with this TLV and interval is the nominal poll interval. The accuracy of the ideal poll times, t_i, are measured relative to the CMTS master clock used to generate timestamps (refer to clause 9.3).

This parameter is only applicable at the CMTS. If defined, this parameter represents a service commitment (or admission criteria) at the CMTS.

Type	Length	Value
24.18	4	µs

C.2.2.6.6 Unsolicited grant size

The value of this parameter specifies the unsolicited grant size in bytes. The grant size includes the entire MAC frame data PDU from the frame control byte to the end of the MAC frame.

This parameter is applicable at the CMTS and MUST be enforced at the CMTS.

Type	Length	Value
24.19	2	

NOTE – For UGS, this parameter should be used by the CMTS to compute the size of the unsolicited grant in mini-slots.

C.2.2.6.7 Nominal grant interval

The value of this parameter specifies the nominal interval (in units of microseconds) between successive data grant opportunities for this service flow. This parameter is required for unsolicited grant and unsolicited grant with activity detection service flows.

The ideal schedule for enforcing this parameter is defined by a reference time, t_0, with the desired transmission times $t_i = t_0 + i \times$ interval. The actual grant times, t'_i, MUST be in the range $t_i \le t'_i \le t_i +$ jitter, where interval is the value specified with this TLV, and jitter is the tolerated grant jitter. When multiple grants per interval are requested, all grants MUST be within this interval, thus the nominal grant interval and tolerated grant jitter MUST be maintained by the CMTS for all grants in this service flow. The accuracy of the ideal grant times, t_i, are measured relative to the CMTS master clock used to generate timestamps (refer to clause 9.3).

This field is mandatory for unsolicited grant and unsolicited grant with activity detection scheduling types. This field is only applicable at the CMTS, and MUST be enforced by the CMTS.

Type	Length	Value
24.20	4	µs

C.2.2.6.8 Tolerated grant jitter

The values in this parameter specifies the maximum amount of time that the transmission opportunities may be delayed from the nominal periodic schedule (measured in microseconds) for this service flow.

The ideal schedule for enforcing this parameter is defined by a reference time, t_0, with the desired transmission times $t_i = t_0 + i \times$ interval. The actual transmission opportunities, t'_i, MUST be in the range $t_i \le t'_i \le t_i +$ jitter, where jitter is the value specified with this TLV and interval is the nominal grant interval. The accuracy of the ideal grant times, t_i, are measured relative to the CMTS master clock used to generate timestamps (refer to clause 9.3).

This field is mandatory for unsolicited grant and unsolicited grant with activity detection scheduling types. This field is only applicable at the CMTS, and MUST be enforced by the CMTS.

Type	Length	Value
24.21	4	μsec

C.2.2.6.9 Grants per interval

For unsolicited grant service, the value of this parameter indicates the actual number of data grants per nominal grant interval. For unsolicited grant service with activity detection, the value of this parameter indicates the maximum number of active grants per nominal grant interval. This is intended to enable the addition of sessions to an existing unsolicited grant service flow via the dynamic service change mechanism, without negatively impacting existing sessions.

The ideal schedule for enforcing this parameter is defined by a reference time, t_0, with the desired transmission times $t_i = t_0 + i \times \text{interval}$. The actual grant times, t'_i, MUST be in the range $t_i \leq t'_i \leq t_i + \text{jitter}$, where interval is the nominal grant interval, and jitter is the tolerated grant jitter. When multiple grants per interval are requested, all grants MUST be within this interval, thus the nominal grant interval and tolerated grant jitter MUST be maintained by the CMTS for all grants in this service flow.

This field is mandatory for unsolicited grant and unsolicited grant with activity detection scheduling types. This field is only applicable at the CMTS, and MUST be enforced by the CMTS.

Type	Length	Value	Valid range
24.22	1	# of grants	0-127

C.2.2.6.10 IP type of service overwrite

The CMTS MUST overwrite IP packets with IP ToS byte value "orig-ip-tos" with the value "new-ip-tos", where new-ip-tos = ((orig-ip-tos AND tos-and-mask) OR tos-or-mask). If this parameter is omitted, then the IP packet ToS byte is not overwritten.

This parameter is only applicable at the CMTS. If defined, this parameter MUST be enforced by the CMTS.

Type	Length	Value
24.23	2	tos-and-mask, tos-or-mask

C.2.2.6.11 Unsolicited grant time reference

For unsolicited grant service and unsolicited grant service with activity detection, the value of this parameter specifies a reference time, t_0, from which can be derived the desired transmission times $t_i = t_0 + i \times \text{interval}$, where interval is the nominal grant interval (refer to clause C.2.2.6.7). This parameter is applicable only for messages transmitted from the CMTS to the CM, and only when a UGS or UGS-AD service flow is being made active. In such cases, this is a mandatory parameter.

Type	Length	Value	Valid range
24.2	4	CMTS timestamp	0-4'294'967'295

The timestamp specified in this parameter represents a count state of the CMTS 10.24 MHz master clock. Since a UGS or UGS-AD service flow is always activated before transmission of this parameter to the modem, the reference time, t_0, is to be interpreted by the modem as the ideal time of the next grant only if t_0 follows the current time. If t_0 precedes the current time, the modem can calculate the offset from the current time to the ideal time of the next grant according to:

$$\text{interval} - (((\text{current time} - t_0)/10.24) \text{ modulus interval})$$

where interval is in units of microseconds, and current time and t_0 are in 10.24-MHz units.

C.2.2.7 Downstream-specific QoS parameter encodings

C.2.2.7.1 Maximum downstream latency

The value of this parameter specifies the maximum latency between the reception of a packet by the CMTS on its NSI and the forwarding of the packet to its RF interface.

If defined, this parameter represents a service commitment (or admission criteria) at the CMTS and MUST be guaranteed by the CMTS. A CMTS does not have to meet this service commitment for service flows that exceed their minimum downstream reserved rate.

Type	Length	Value
25.14	4	µs

C.2.2.8 Payload header suppression

This field defines the parameters associated with payload header suppression.

Type	Length	Value
26	n	

NOTE – The entire payload header suppression TLV MUST have a length of less than 255 characters.

C.2.2.8.1 Classifier reference

The value of the field specifies a classifier reference that identifies the corresponding classifier (refer to clause C.2.1.3.1).

Type	Length	Value
26.1	1	1-255

C.2.2.8.2 Classifier identifier

The value of the field specifies a classifier identifier that identifies the corresponding classifier (refer to clause C.2.1.3.2).

Type	Length	Value
26.2	2	1-65'535

C.2.2.8.3 Service flow reference

The value of the field specifies a service flow reference that identifies the corresponding service flow (refer to clause C.2.2.3.1).

Type	Length	Value
26.3	2	1-65'535

C.2.2.8.4 Service flow identifier

The value of this field specifies the service flow identifier that identifies the service flow to which the PHS rule applies.

Type	Length	Value
26.4	4	1-4'294'967'295

C.2.2.8.5 Dynamic service change action

When received in a dynamic service change request, this indicates the action that MUST be taken with this payload header suppression byte string.

Type	Length	Value
26.5	1	0: Add PHS rule 1: Set PHS rule 2: Delete PHS rule 3: Delete all PHS rules

The "set PHS rule" command is used to add specific TLVs to a partially defined payload header suppression rule. A PHS rule is partially defined when the PHSF and PHSS values are not both known. A PHS rule becomes fully defined when the PHSF and PHSS values are both known. Once a PHS rule is fully defined, "set PHS rule" MUST NOT be used to modify existing TLVs.

The "Delete all PHS Rules" command is used to delete all PHS rules for a specified service flow. See clause 8.3.15 for details on DSC-REQ required PHS parameters when using this option.

NOTE – An attempt to add a PHS rule that already exists is an error condition.

C.2.2.9 Payload header suppression error encodings

This field defines the parameters associated with payload header suppression errors.

Type	Length	Value
26.6	n	

A payload header suppression error encoding consists of a single payload header suppression error parameter set which is defined by the following individual parameters: Errored parameter, confirmation code and error message.

The payload header suppression error encoding is returned in REG-RSP, DSA-RSP and DSC-RSP messages to indicate the reason for the recipient's negative response to a payload header suppression rule establishment request in a REG-REQ, DSA-REQ or DSC-REQ message.

On failure, the REG-RSP, DSA-RSP or DSC-RSP MUST include one payload header suppression error encoding for at least one failed payload header suppression rule requested in the REG-REQ, DSA-REQ or DSC-REQ message. A payload header suppression error encoding for the failed payload header suppression rule MUST include the confirmation code and errored parameter and MAY include an error message. If some payload header suppression rule sets are rejected but other payload header suppression rule sets are accepted, then payload header suppression error encodings MUST be included for only the rejected payload header suppression rules. On success of the entire transaction, the RSP or ACK message MUST NOT include a payload header suppression error encoding.

Multiple payload header suppression error encodings MAY appear in a REG-RSP, DSA-RSP or DSC-RSP message, since multiple payload header suppression parameters may be in error. A message with even a single payload header suppression error encoding MUST NOT contain any other protocol payload header suppression encodings (e.g., IP, 802.1P/Q).

A payload header suppression error encoding MUST NOT appear in any REG-REQ, DSA-REQ or DSC-REQ messages.

C.2.2.9.1 Errored parameter

The value of this parameter identifies the subtype of a requested payload header suppression parameter in error in a rejected payload header suppression request. A payload header suppression error parameter set MUST have exactly one errored parameter TLV within a given payload header suppression error encoding.

Subtype	Length	Value
26.6.1	1	Payload header suppression encoding subtype in error

C.2.2.9.2 Error code

This parameter indicates the status of the request. A non-zero value corresponds to the confirmation code as described in clause C.4. A payload header suppression error parameter set MUST have exactly one error code within a given payload header suppression error encoding.

Subtype	Length	Value
26.6.2	1	Confirmation code

A value of okay (0) indicates that the payload header suppression request was successful. Since a payload header suppression error parameter set only applies to errored parameters, this value MUST NOT be used.

C.2.2.9.3 Error message

This subtype is optional in a payload header suppression error parameter set. If present, it indicates a text string to be displayed on the CM console and/or log that further describes a rejected payload header suppression request. A payload header suppression error parameter set MAY have zero or one error message subtypes within a given payload header suppression error encoding.

Subtype	Length	Value
26.6.3	n	Zero-terminated string of ASCII characters

NOTE 1 – The length n includes the terminating zero.

NOTE 2 – The entire payload header suppression encoding message MUST have a total length of less than 256 characters.

C.2.2.10 Payload header suppression rule encodings

C.2.2.10.1 Payload header suppression field (PHSF)

The value of this field are the bytes of the headers that MUST be suppressed by the sending entity, and MUST be restored by the receiving entity. In the upstream, the PHSF corresponds to the string of PDU bytes starting with the first byte after the MAC header checksum. For the downstream, the PHSF corresponds to the string of PDU bytes starting with the 13th byte after the MAC header checksum. This string of bytes is inclusive of both suppressed and unsuppressed bytes of the PDU header. The value of the unsuppressed bytes within the PHSF is implementation-dependent.

The ordering of the bytes in the value field of the PHSF TLV string MUST follow the sequence:

Upstream

MSB of PHSF value = 1st byte of PDU

2nd MSB of PHSF value = 2nd byte of PDU

...

nth byte of PHSF (LSB of PHSF value) = nth byte of PDU

Downstream

MSB of PHSF value = 13th byte of PDU

2nd MSB of PHSF value = 14th byte of PDU

...

nth byte of PHSF (LSB of PHSF value) = (n + 13)th byte of PDU

Type	Length	Value
26.7	n	String of bytes suppressed

The length n MUST always be the same as the value for PHSS.

C.2.2.10.2 Payload header suppression index (PHSI)

The payload header suppression index (PHSI) has a value between 1 and 255 that uniquely references the suppressed byte string. The index is unique per service flow in the upstream direction and unique per CM in the downstream direction. The upstream and downstream PHSI values are independent of each other.

Type	Length	Value
26.8	1	Index value

C.2.2.10.3 Payload header suppression mask (PHSM)

The value of this field is used to interpret the values in the payload header suppression field. It is used at both the sending and receiving entities on the link. The PHSM allows fields such as sequence numbers or checksums, that vary in value, to be excluded from suppression with the constant bytes around them suppressed.

Type	Length	Value
26.9	n	bit 0: 0 = Do not suppress first byte of the suppression field. 1 = Suppress first byte of the suppression field. bit 1: 0 = Do not suppress second byte of the suppression field. 1 = Suppress second byte of the suppression field. bit x: 0 = Do not suppress (x + 1) byte of the suppression field. 1 = Suppress (x + 1) byte of the suppression field.

The length n is ceiling (PHSS/8). Bit 0 is the MSB of the value field. The value of each sequential bit in the PHSM is an attribute for the corresponding sequential byte in the PHSF.

If the bit value is a "1" (and verification passes or is disabled), the sending entity MUST suppress the byte, and the receiving entity MUST restore the byte from its cached PHSF. If the bit value is a "0", the sending entity MUST NOT suppress the byte, and the receiving entity MUST restore the byte by using the next byte in the packet.

If this TLV is not included, the default is to suppress all bytes.

C.2.2.10.4 Payload header suppression size (PHSS)

The value of this field is the total number of bytes in the payload header suppression field (PHSF) for a service flow that uses payload header suppression.

Type	Length	Value
26.10	1	Number of bytes in the suppression string

This TLV is used when a service flow is being created. For all packets that get classified and assigned to a service flow with payload header suppression enabled, suppression MUST be performed over the specified number of bytes as indicated by the PHSS and according to the PHSM. If this TLV is included in a service flow definition with a value of 0 bytes, then payload header suppression is disabled. A non-zero value indicates payload header suppression is enabled. Until the PHSS value is known, the PHS rule is considered partially defined, and suppression will not be performed. A PHS rule becomes fully defined when both PHSS and PHSF are known.

C.2.2.10.5 Payload header suppression verification (PHSV)

The value of this field indicates to the sending entity whether or not the packet header contents are to be verified prior to performing suppression. If PHSV is enabled, the sender MUST compare the bytes in the packet header with the bytes in the PHSF that are to be suppressed as indicated by the PHSM.

Type	Length	Value
26.11	1	0: Verify 1: Do not verify

If this TLV is not included, the default is to verify. Only the sender MUST verify suppressed bytes. If verification fails, the payload header MUST NOT be suppressed (refer to clause 10.4.3).

C.2.2.10.6 Vendor-specific PHS parameters

This allows vendors to encode vendor-specific PHS parameters using the DOCSIS extension field. The vendor ID MUST be the first TLV embedded inside vendor-specific PHS parameters. If the first TLV inside vendor-specific PHS parameters is not a vendor ID, then the TLV MUST be discarded (refer to clause C.1.1.17).

Type	Length	Value
26.43	n	

C.3 Encodings for other interfaces

C.3.1 Telephone settings option

This configuration setting describes parameters that are specific to telephone return systems. It is composed from a number of encapsulated type/length/value fields (see [DOCS6]).

Type	Length	Value
15 (= TRI_CFG01)	n	

C.3.2 Baseline privacy configuration settings option

This configuration setting describes parameters that are specific to baseline privacy. It is composed from a number of encapsulated type/length/value fields (see [DOCS8]).

Type	Length	Value
17 (= BP_CFG)	n	

C.4 Confirmation code

The confirmation code (CC) provides a common way to indicate failures for registration response, registration ack, dynamic service addition response, dynamic service addition ack, dynamic service delete response, dynamic service change response, dynamic service change ack and dynamic channel change response MAC management messages. The confirmation codes in this clause are used both as message confirmation codes and as error codes in error set encodings that may be carried in these messages.

Confirmation code is one of the following:

– okay/success (0);
– reject-other (1);
– reject-unrecognized-configuration-setting (2);
– reject-temporary/reject-resource (3);
– reject-permanent/reject-admin (4);
– reject-not-owner (5);
– reject-service-flow-not-found (6);
– reject-service-flow-exists (7);
– reject-required-parameter-not-present (8);
– reject-header-suppression (9);
– reject-unknown-transaction-id (10);
– reject-authentication-failure (11);
– reject-add-aborted (12);
– reject-multiple-errors (13);
– reject-classifier-not-found (14);
– reject-classifier-exists (15);
– reject-PHS-rule-not-found (16);
– reject-PHS-rule-exists (17);
– reject-duplicate-reference-ID-or-index-in-message (18);
– reject-multiple-upstream-service-flows (19);
– reject-multiple-downstream-service-flows (20);
– reject-classifier-for-another-service-flow (21);
– reject-PHS-for-another-service-flow (22);
– reject-parameter-invalid-for-context (23);
– reject-authorization-failure (24);
– reject-temporary-DCC (25).

The confirmation codes MUST be used in the following way:

- Okay or success (0) means the message was received and successful.
- Reject-other (1) is used when none of the other reason codes apply.
- Reject-unrecognized-configuration-setting (2) is used when a configuration setting is not recognized or when its value is outside of the specified range.
- Reject-temporary (3), also known as reject-resource, indicates that the current loading of the CMTS or CM prevents granting the request, but that the request might succeed at another time

- Reject-permanent (4), also known as reject-admin, indicates that, for policy, configuration or capabilities reasons, the request would never be granted unless the CMTS or CM were manually reconfigured or replaced.
- Reject-not-owner (5) the requester is not associated with this service flow.
- Reject-service-flow-not-found (6) the service flow indicated in the request does not exist.
- Reject-service-flow-exists (7) the service flow to be added already exists.
- Reject-required-parameter-not-present (8) a required parameter has been omitted.
- Reject-header-suppression (9) the requested header suppression cannot be supported for whatever reason.
- Reject-unknown-transaction-id (10) the requested transaction continuation is invalid because the receiving end-point does not view the transaction as being 'in process' (i.e., the message is unexpected or out of order).
- Reject-authentication-failure (11) the requested transaction was rejected because the message contained an invalid HMAC-Digest, CMTS-MIC, provisioned IP address or timestamp.
- Reject-add-aborted (12) the addition of a dynamic service flow was aborted by the initiator of the dynamic service addition.
- Reject-multiple-errors (13) is used when multiple errors have been detected.
- Reject-classifier-not-found (14) is used when the request contains an unrecognized classifier ID.
- Reject-classifier-exists (15) indicates that the ID of a classifier to be added already exists.
- Reject-PHS-rule-not-found (16) indicates that the request contains an SFID/classifier ID pair for which no PHS rule exists.
- Reject-PHS-rule-exists (17) indicates that the request to add a PHS rule contains an SFID/classifier ID pair for which a PHS rule already exists.
- Reject-duplicate-reference-ID-or-index-in-message (18) indicates that the request used an SFR, classifier reference, SFID or classifier ID twice in an illegal way.
- Reject-multiple-upstream-service-flows (19) is used when DSA/DSC contains parameters for more than one upstream flow.
- Reject-multiple-downstream-service-flows (20) is used when DSA/DSC contains parameters for more than one downstream flow.
- Reject-classifier-for-another-service-flow (21) is used in DSA-RSP when the DSA-REQ includes classifier parameters for a SF other than the new SF(s) being added by the DSA.
- Reject-PHS-for-another-service-flow (22) is used in DSA-RSP when the DSA-REQ includes a PHS rule for a SF other than the new SF(s) being added by the DSA.
- Reject-parameter-invalid-for-context (23) indicates that the parameter supplied cannot be used in the encoding in which it was included, or that the value of a parameter is invalid for the encoding in which it was included.
- Reject-authorization-failure (24) the requested transaction was rejected by the authorization module.
- Reject-temporary-DCC (25) indicates that the requested resources are not available on the current channels at this time, and the CM should re-request them on new channels after completing a channel change in response to a DCC command which the CMTS will send. If no DCC is received, the CM must wait for a time of at least T14 before re-requesting the resources on the current channels.

C.4.1 Confirmation codes for dynamic channel change

The CM may return in the DCC-RSP message an appropriate rejection code from clause C.4. It may also return one of the following confirmation codes, which are unique to DCC-RSP. The confirmation codes MUST be used in the following way:

- depart (180) indicates the CM is on the old channel and is about to perform the jump to the new channel.
- arrive (181) indicates the CM has performed the jump and has arrived at the new channel.
- reject-already-there (182) indicates that the CMTS has asked the CM to move to a channel that it is already occupying.
- reject-20-disable (183) indicates that the CMTS has asked a CM with 2.0 mode disabled to move to a type 3 channel that it cannot use and a UCD substitution was sent in the corresponding DCC-REQ.

C.4.2 Confirmation codes for major errors

These confirmation codes MUST be used only as message confirmation codes in REG-ACK, DSA-RSP, DSA-ACK, DSC-RSP or DSC-ACK messages, or as the response code in REG-RSP messages for 1.1 CMs. In general, the errors associated with these confirmation codes make it impossible either to generate an error set that can be uniquely associated with a parameter set in the REG-REQ, DSA-REQ or DSC-REQ message, or to generate a full RSP message.

– reject-major-service-flow-error (200);
– reject-major-classifier-error (201);
– reject-major-PHS-rule-error (202);
– reject-multiple-major-errors (203);
– reject-message-syntax-error (204);
– reject-primary-service-flow-error (205);
– reject-message-too-big (206);
– reject-invalid-modem-capabilities (207).

The confirmation codes MUST be used only in the following way:

- Reject-major-service-flow-error (200) indicates that the REQ message did not have either a SFR or SFID in a service flow encoding, and that service flow major errors were the only major errors.
- Reject-major-classifier-error (201) indicates that the REQ message did not have a classifier reference, or did not have both a classifier ID and a service flow ID, and that classifier major errors were the only major errors.
- Reject-major-PHS-rule-error (202) indicates that the REQ message did not have a both a service flow reference/identifier and a classifier reference/identifier, and that PHS rule major errors were the only major errors.
- Reject-multiple-major-errors (203) indicates that the REQ message contained multiple major errors of types 200, 201, 202.
- Reject-message-syntax-error (204) indicates that the REQ message contained syntax error(s) (e.g., a TLV length error) resulting in parsing failure.
- Reject-primary-service-flow-error (205) indicates that a REG-REQ or REG-RSP message did not define a required primary service flow, or a required primary service flow was not specified active.
- Reject-message-too-big (206) is used when the length of the message needed to respond exceeds the maximum allowed message size.

- Reject-invalid-modem-capabilities (207) indicates that the REG-REQ contained either an invalid combination of modem capabilities or modem capabilities that are inconsistent with the services in the REG-REQ.

Annex D

CM configuration interface specification

(This annex forms an integral part of this Recommendation)

D.1 CM IP addressing

D.1.1 DHCP fields used by the CM

The following fields MUST be present in the DHCP request from the CM and MUST be set as described below:

- The hardware type (htype) MUST be set to 1 (Ethernet).
- The hardware length (hlen) MUST be set to 6.
- The client hardware address (chaddr) MUST be set to the 48-bit MAC address associated with the RF interface of the CM.
- The "client identifier" option MUST be included, with the hardware type set to 1, and the value set to the same 48 bit MAC address as the chaddr field.
- Option code 60 (vendor class identifier) – To allow for the differentiation between DOCS 2.0 and DOCS 1.x CM requests, a compliant CM MUST send the following ASCII-coded string in option code 60, "docsis2.0:xxxxxxx". Where xxxxxxx MUST be an ASCII representation of the hexadecimal encoding of the modem capabilities; refer to clause C.1.3.1. For example, the ASCII encoding for the first two TLVs (concatenation and DOCS version) of a DOCS 2.0 modem would be 05nn010101020102. Note that many more TLVs are required for a DOCS2.0 modem and the field "nn" will contain the length of all the TLVs. This example shows only two TLVs for simplicity.
- The "parameter request list" option MUST be included. The option codes that MUST be included in the list are:
  - Option code 1 (subnet mask);
  - Option code 2 (time offset);
  - Option code 3 (router option);
  - Option code 4 (time server option);
  - Option code 7 (log server option).

The following fields are expected in the DHCP response returned to the CM. Fields identified as critical MUST be present in the DHCP response, and fields identified as non-critical SHOULD be present. The CM MUST configure itself with the critical fields from the DHCP response, and, if present, with the non-critical fields.

- The IP address to be used by the CM (yiaddr) (critical).
- The IP address of the TFTP server for use in the next phase of the bootstrap process (siaddr) (critical).
- If the DHCP server is on a different network (requiring a relay agent), then the IP address of the relay agent (giaddr).

  NOTE – This may differ from the IP address of the first hop router (non-critical).
- The name of the CM configuration file to be read from the TFTP server by the CM (file) (critical).
- The subnet mask to be used by the CM (subnet mask, option 1) (non-critical).

- The time offset of the CM from universal coordinated time (UTC) (time offset, option 2). This is used by the CM to calculate the local time for use in timestamping error logs (non-critical).
- A list of addresses of one or more routers to be used for forwarding CM-originated IP traffic (router option, option 3). The CM is not required to use more than one router IP address for forwarding, but MUST use at least one (non-critical).
- A list of [RFC 868] time-servers from which the current time may be obtained (time server option, option 4) (non-critical).
- A list of SYSLOG servers to which logging information may be sent (log server option, option 7); see [DOCS5] (non-critical).

If a critical field is missing, or is invalid in the DHCP response during initialization, the CM MUST log an error and re-initialize its MAC and continue channel scanning.

If a non-critical field is missing or is invalid in the DHCP response during initialization, the CM MUST log a warning, ignore the field and go operational, with the following considerations:

- If the subnet mask is missing or is invalid, the CM MUST use the default for the IP of class A, B or C as defined in [RFC 791].
- If the time server is missing or is invalid, the CM MUST initialize the time for the events to Jan 1, 1970, 0h00.

If the IP address field is missing or is invalid in the DHCP response during renew or rebind, the CM MUST log an error and re-initialize its MAC and continue channel scanning.

If any other critical or non-critical field is missing or is invalid in the DHCP response during renew or rebind, the CM MUST log a warning, ignore the field if it is valid, and stay operational.

To assist the DHCP server in differentiating a CM discovery request from a CPE-side LAN discovery request, a CMTS MUST implement the following:

- All CMTSs MUST support the DHCP relay agent information option [RFC 3046]. Specifically, the CMTS MUST include the 48-bit MAC address of the RF side interface of the CM generating or bridging the DHCP discovery request in the agent remote ID sub-option field before relaying the discovery to a DHCP server.
- If the CMTS is a router, it MUST use a giaddr field to differentiate between CM and CPE side stations if they are provisioned to be in different IP subnets. Bridging CMTSs SHOULD also provide this functionality.

D.2 CM configuration

D.2.1 CM binary configuration file format

The CM-specific configuration data MUST be contained in a file which is downloaded to the CM via TFTP. This is a binary file in the same format defined for DHCP vendor extension data [RFC 2132].

It MUST consist of a number of configuration settings (1 per parameter) each of the form:

Type/Length/Value

Type is a single-octet identifier which defines the parameter.

Length is a single octet containing the length of the value field in octets (not including type and length fields).

Value is from one to 254 octets containing the specific value for the parameter.

The configuration settings MUST follow each other directly in the file, which is a stream of octets (no record markers).

Configuration settings are divided into three types:

- standard configuration settings which MUST be present;
- standard configuration settings which MAY be present;
- DOCSIS extension field configuration settings.

CMs MUST be capable of processing all standard configuration settings. CMs MUST ignore any configuration setting present in the configuration file that it cannot interpret. To allow uniform management of CMs conformant to this Recommendation, conformant CMs MUST support a 8192-byte configuration file at a minimum.

Authentication of the provisioning information is provided by two message integrity check (MIC) configuration settings, CM MIC and CMTS MIC.

- CM MIC is a digest which ensures that the data sent from the provisioning server were not modified en route. This is NOT an authenticated digest (it does not include any shared secret).
- CMTS MIC is a digest used to authenticate the provisioning server to the CMTS during registration. It is taken over a number of fields, one of which is a shared secret between the CMTS and the provisioning server.

Use of the CM MIC allows the CMTS to authenticate the provisioning data without needing to receive the entire file.

Thus the file structure is of the form shown in Figure D.1:

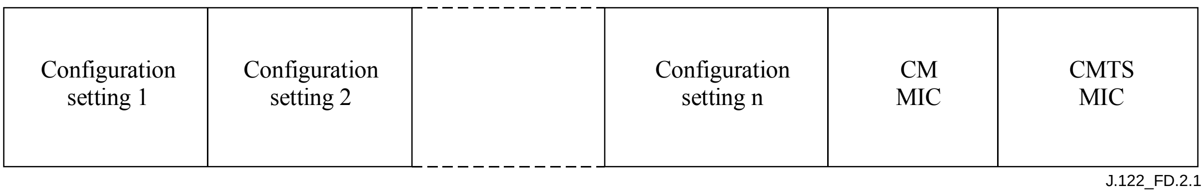

Figure D.1 – Binary configuration file format

D.2.2 Configuration file settings

The following configuration settings MUST be included in the configuration file and MUST be supported by all CMs. The CM MUST NOT send a REG-REQ based on a configuration file that lacks these mandatory items:

- Network access configuration setting;
- CM MIC configuration setting;
- CMTS MIC configuration setting;
- End configuration setting;
- DOCS 1.0 class of service configuration setting.

or:

- Upstream service flow configuration setting;
- Downstream service flow configuration setting.

NOTE 1 – A DOCS 1.0 CM must be provided with a DOCS 1.0 class of service configuration. A CM conformant with this Recommendation should only be provisioned with DOCS 1.0 class of service configuration information if it is to behave as a DOCS 1.0 CM; otherwise, it should be provisioned with service flow configuration settings.

The following configuration settings MAY be included in the configuration file and if present MUST be supported by all CMs:

- downstream frequency configuration setting;
- upstream channel id configuration setting;
- baseline privacy configuration setting;
- software upgrade filename configuration setting;
- upstream packet classification setting;
- downstream packet classification setting;
- SNMP write-access control;
- SNMP MIB object;
- software server IP address;
- CPE Ethernet MAC address;
- maximum number of CPEs;
- maximum number of classifiers;
- privacy enable configuration setting;
- payload header suppression;
- TFTP server timestamp;
- TFTP server provisioned modem address;
- pad configuration setting;
- SNMPv3 notification receiver;
- enable 2.0 mode;
- enable test modes;
- static multicast MAC address.

The following configuration setting MAY be included in the configuration file and, if present and applicable to this type of modem, MUST be supported:

- telephone settings option.

The following configuration settings MAY be included in the configuration file and, if present, MAY be supported by a CM.

- DOCSIS extension field configuration settings.

NOTE 2 – There is a limit on the size of registration request and registration response frames (see clause 8.2.5.2). The configuration file should not be so large as to require the CM or CMTS to exceed that limit.

D.2.3 Configuration file creation

The sequence of operations required to create the configuration file is as shown in Figures D.2 through D.5.

1) Create the type/length/value entries for all the parameters required by the CM.

Type length value for parameter 1
Type length value for parameter 2
Type length value for parameter n

Figure D.2 – Create TLV entries for parameters required by the CM

2) Calculate the CM message integrity check (MIC) configuration setting as defined in clause D.2.3.1 and add to the file following the last parameter using code and length values defined for this field.

Type length value for parameter 1
Type length value for parameter 2
Type length value for parameter n
Type length value for CM MIC

Figure D.3 – Add CM MIC

3) Calculate the CMTS message integrity check (MIC) configuration setting as defined in clause D.3.1 and add to the file following the CM MIC using code and length values defined for this field.

Type length value for parameter 1
Type length value for parameter 2
Type length value for parameter n
Type length value for CM MIC
Type length value for CMTS MIC

Figure D.4 – Add CMTS MIC

4) Add the end-of-data marker.

Type length value for parameter 1
Type length value for parameter 2
Type length value for parameter n
Type length value for CM MIC
Type length value for CMTS MIC
End-of-data marker

Figure D.5 – Add end-of-data marker

D.2.3.1 CM MIC calculation

The CM message integrity check configuration setting MUST be calculated by performing an MD5 digest over the bytes of the configuration setting fields. It is calculated over the bytes of these settings as they appear in the TFTPed image, without regard to TLV ordering or contents. There are two exceptions to this disregard of the contents of the TFTPed image:

1) The bytes of the CM MIC TLV itself are omitted from the calculation. This includes the type, length, and value fields.
2) The bytes of the CMTS MIC TLV are omitted from the calculation. This includes the type, length, and value fields.

On receipt of a configuration file, the CM MUST recompute the digest and compare it to the CM MIC configuration setting in the file. If the digests do not match then the configuration file MUST be discarded.

D.3 Configuration verification

It is necessary to verify that the CM's configuration file has come from a trusted source. Thus, the CMTS and the configuration server share an authentication string that they use to verify portions of the CM's configuration in the registration request.

D.3.1 CMTS MIC calculation

The CMTS message integrity check configuration setting MUST be calculated by performing an MD5 digest over the following configuration setting fields, when present in the configuration file, in the order shown:

- downstream frequency configuration setting;
- upstream channel id configuration setting;
- network access configuration setting;
- DOCS 1.0 class-of-service configuration setting;
- baseline privacy configuration setting;
- DOCSIS extension field configuration settings;
- CM MIC configuration setting;
- maximum number of CPEs;
- TFTP server timestamp;
- TFTP server provisioned modem address;
- upstream packet classification setting;
- downstream packet classification setting;
- upstream service flow configuration setting;
- downstream service flow configuration setting;
- maximum number of classifiers;
- privacy enable configuration setting;
- payload header suppression;
- subscriber management control;
- subscriber management CPE IP table;
- subscriber management filter groups;
- enable test modes.

The bulleted list specifies the order of operations when calculating the CMTS MIC over configuration setting type fields. The CMTS MUST calculate the CMTS MIC over TLVs of the same type in the order they were received. Within type fields, the CMTS MUST calculate the CMTS MIC over the subtypes in the order they were received. To allow for correct CMTS MIC calculation by the CMTS, the CM MUST NOT reorder configuration file TLVs of the same type or subtypes within any given type in its registration request message.

All configuration setting fields MUST be treated as if they were contiguous data when calculating the CM MIC.

The digest MUST be added to the configuration file as its own configuration setting field using the CMTS MIC configuration setting encoding.

The authentication string is a shared secret between the provisioning server (which creates the configuration files) and the CMTS. It allows the CMTS to authenticate the CM provisioning. The authentication string is to be used as the key for calculating the keyed CMTS MIC digest as stated in clause D.3.1.1.

The mechanism by which the shared secret is managed is up to the system operator.

On receipt of a configuration file, the CM MUST forward the CMTS MIC as part of the registration request (REG-REQ).

On receipt of a REG-REQ, the CMTS MUST recompute the digest over the included fields and the authentication string and compare it to the CMTS MIC configuration setting in the file. If the digests do not match, the registration request MUST be rejected by setting the authentication failure result in the registration response status field.

D.3.1.1 Digest calculation

The CMTS MIC digest field MUST be calculated using HMAC-MD5 as defined in [RFC 2104].

Annex E

The data-over-cable spanning tree protocol

(This annex forms an integral part of this Recommendation)

Clause 5.1.2.1 requires the use of the spanning tree protocol on CMs that are intended for commercial use and on bridging CMTSs. This annex describes how the IEEE 802.1D spanning tree protocol is adapted to work for data-over-cable systems.

E.1 Background

A spanning tree protocol is frequently employed in a bridged network in order to deactivate redundant network connections; i.e., to reduce an arbitrary network mesh topology to an active topology that is a rooted tree that spans all of the network segments. The spanning tree algorithm and protocol should not be confused with the data-forwarding function itself; data forwarding may follow transparent learning bridge rules, or may employ any of several other mechanisms. By deactivating redundant connections, the spanning tree protocol eliminates topological loops, which would otherwise cause data packets to be forwarded forever for many kinds of forwarding devices.

A standard spanning tree protocol [b-ISO/IEC 10038] is employed in most bridged local area networks. This protocol was intended for private LAN use and requires some modification for cable data use.

E.2 Public spanning tree

To use a spanning tree protocol in a public-access network such as data-over-cable, several modifications are needed to the basic IEEE 802.1D process. Primarily, the public spanning tree must be isolated from any private spanning tree networks to which it is connected. This is to protect both the public cable network and any attached private networks. Figure E.1 illustrates the general topology.

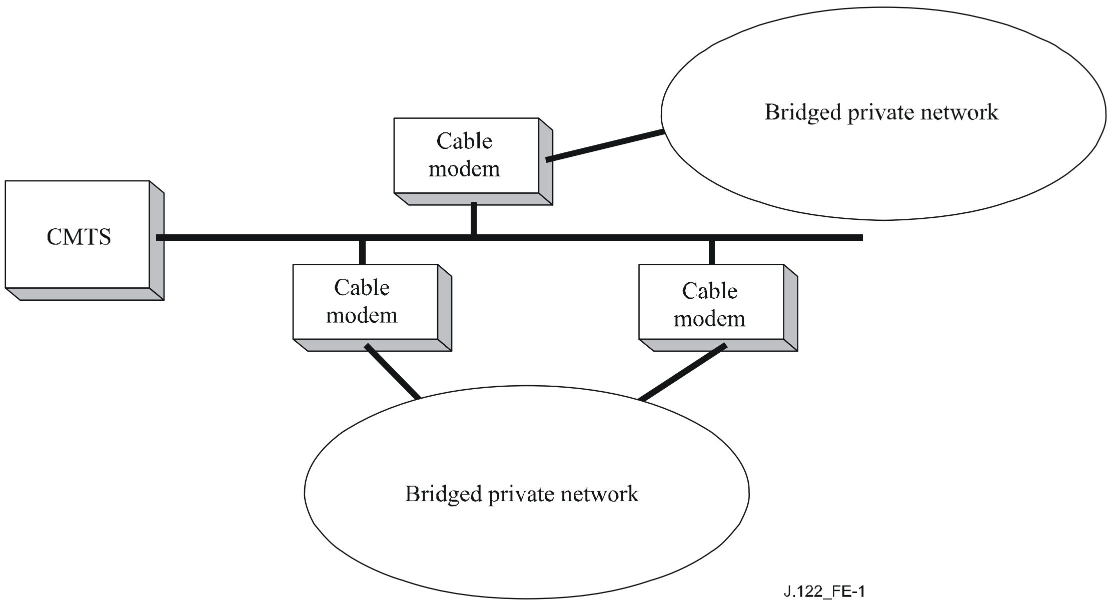

Figure E.1 – Spanning tree topology

The task for the public spanning tree protocol, with reference to Figure E.1, is to:

- Isolate the private bridged networks from each other. If the two private networks merge spanning trees then each is subject to instabilities in the other's network. Also, the combined tree may exceed the maximum allowable bridging diameter.
- Isolate the public network from the private networks' spanning trees. The public network must not be subject to instabilities induced by customers' networks; nor should it change the spanning tree characteristics of the customers' networks.
- Disable one of the two redundant links into the cable network, so as to prevent forwarding loops. This should occur at the cable modem, rather than at an arbitrary bridge within the customer's network.

The spanning tree protocol must also serve the topology illustrated in Figure E.2:

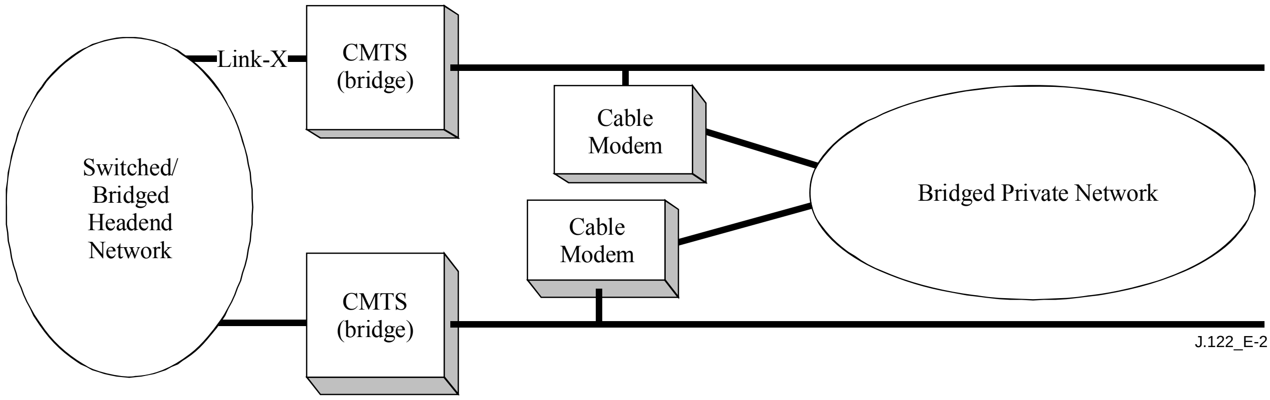

Figure E.2 – Spanning tree across CMTSs

In Figure E.2, in normal operation, the spanning tree protocol should deactivate a link at one of the two cable modems. It should not divert traffic across the private network. Note that in some circumstances, such as deactivation of link-X, spanning tree will divert traffic onto the private network (although limits on learned MAC addresses will probably throttle most transit traffic). If this diversion is undesirable, then it must be prevented by means external to spanning tree; for example, by using routers.

E.3 Public spanning tree protocol details

The data-over-cable spanning tree algorithm and protocol is identical to that defined in [b-ISO/IEC 10038], with the following exceptions:

- When transmitting configuration bridge protocol data units (BPDUs), the data-over-cable spanning tree multicast address 01-E0-2F-00-00-03 MUST be used rather than that defined in IEEE 802.1D. These BPDUs will be forwarded rather than recalculated by ordinary IEEE 802.1D bridges.
- When transmitting configuration BPDUs, the SNAP header AA-AA-03-00-E0-2F-73-74 MUST be used rather than the LLC 42-42-03 header employed by IEEE 802.1D. This is to further differentiate these BPDUs from those used by IEEE 802.1D bridges, in the event that some of those bridges do not correctly identify multicast MAC addresses[29].
- IEEE 802.1D BPDUs MUST be ignored and silently discarded.

[29] It is likely that there are a number of spanning tree bridges deployed that rely solely on the LSAPs to distinguish 802.1D packets. Such devices would not operate correctly if the data-over-cable BPDUs also used LSAP = 0x42.

- Topology change notification (TCN) PDUs MUST NOT be transmitted (or processed). TCNs are used in IEEE networks to accelerate the aging of the learning database when the network topology may have changed. Since the learning mechanism within the cable network typically differs, this message is unnecessary and may result in unnecessary flooding.
- CMTSs operating as bridges must participate in this protocol and must be assigned higher priorities (more likely to be root) than cable modems. The NSI interface on the CMTS SHOULD be assigned a port cost equivalent to a link speed of at least 100 Mbit/s. These two conditions, taken together, should ensure that:
  1) a CMTS is the root; and
  2) any other CMTS will use the headend network rather than a customer network to reach the root.
- The MAC forwarder of the CMTS MUST forward BPDUs from upstream to downstream channels, whether or not the CMTS is serving as a router or a bridge.

Note that CMs with this protocol enabled will transmit BPDUs onto subscriber networks in order to identify other CMs on the same subscriber network. These public spanning tree BPDUs will be carried transparently over any bridged private subscriber network. Similarly, bridging CMTSs will transmit BPDUs on the NSI as well as on the RFI interface. The multicast address and SNAP header defined above are used on all links.

E.4 Spanning tree parameters and defaults

Clause 4.10.2 of [b-ISO/IEC 10038] specifies a number of recommended parameter values. Those values should be used, with the exceptions listed below:

E.4.1 Path cost

In [b-ISO/IEC 10038], the following formula is used:

Path_Cost = 1000/Attached_LAN_speed_in_Mbit/s

For CMs, this formula is adapted as:

Path_Cost = 1000/(Upstream_modulation_rate × bits_per_symbol_for_long_data_grant)

That is, the modulation type (QPSK or 16QAM) for the long data grant IUC is multiplied by the raw modulation rate to determine the nominal path cost. Table E.1 provides the derived values.

Table E.1 – CM path cost

Modulation rate	Default path cost	
kHz	**QPSK**	**16QAM**
160	3125	1563
320	1563	781
640	781	391
1280	391	195
2560	195	98

For CMTSs, this formula is:

Path_Cost = 1000/(Downstream_symbol_rate × bits_per_symbol)

E.4.2 Bridge priority

The bridge priority for CMs SHOULD default to 36 864 (0x9000). This is to bias the network so that the root will tend to be at the CMTS. The CMTS SHOULD default to 32 768, as per [b-ISO/IEC 10038].

Note that both of these recommendations affect only the *default* settings. These parameters, as well as others defined in [b-ISO/IEC 10038], SHOULD be manageable throughout their entire range through the bridge MIB ([b-RFC 1493]) or other means.

Annex F

European specification additions

(This annex forms an integral part of this Recommendation)

This annex applies to the second technology option referred to in clause 1.1, "scope".

This annex describes the physical layer specifications required for what are generally called Euro-DOCSIS cable modems. This is an optional annex and in no way affects certification of North American DOCSIS 1.x and DOCSIS 2.0 modems.

The numbering of the clauses has been made so that the suffix after the F refers to the part of the specification that has changed. As a consequence, some clauses are missing in this annex because no change is required.

F.1 Scope and purpose

Refer to clause 1.

F.2 References

Refer to clause 2.

F.3 Glossary

Refer to clause 3.

F.4 Functional assumptions

This clause describes the characteristics of cable television plants to be assumed 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 with the environment described in this clause.

F.4.1 Broadband access network

A coaxial-based broadband access network is assumed. This may take the form of either an all-coax or hybrid-fibre/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 analogue transmission. The key functional characteristics assumed in this annex are the following:

- two-way transmission;
- a maximum optical/electrical spacing between the CMTS and the most distant customer terminal of 160 km (route metres);
- a maximum differential optical/electrical spacing between the CMTS and the closest and most distant modems of 160 km (route metres).

F.4.2 Equipment assumptions

F.4.2.1 Frequency plan

In the downstream direction, the cable system is assumed to have a passband with a typical lower edge between 47 and 87.5 MHz and an upper edge which is implementation-dependent but is typically in the range of 300 to 862 MHz. Within that passband, PAL/SECAM analogue television signals in 7/8 MHz channels, FM-radio signals, as well as other narrow-band and wideband digital signals are assumed to be present.

In the upstream direction, the cable system is assumed to have a passband with a lower edge at 5 MHz and an upper edge which is implementation-dependent but is typically in the range of 25 to 65 MHz.

F.4.2.2 Compatibility with other services

Refer to clause 4.2.2.

F.4.2.3 Fault isolation impact on other users

Refer to clause 4.2.3.

F.4.2.4 Cable system terminal devices

Compliance with EMC requirements is not covered by this Recommendation. The protection requirements with respect to electromagnetic compatibility are contained in harmonized standards published in the official journal of the European communities.

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

Requirements for safety are outside the scope of this Recommendation. Safety standards for European applications are published by CENELEC.

NOTE 1 – Examples of such CENELEC product safety standards are [b-EN 60950] and [b-EN 50083-1].

NOTE 2 – For CENELEC safety categories of interfaces, see [b-EG 201 212].

F.4.3 RF channel assumptions

Refer to 4.3.

F.4.3.1 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 F.1. This assumes nominal analogue video carrier level (peak envelope power) in a 7/8 MHz channel bandwidth. All conditions are present concurrently.

Table F.1 – Assumed downstream RF channel transmission characteristics for analogue TV and sound signals

Parameter	Value
Frequency range	Cable system normal downstream operating range is from 47 MHz to as high as 862 MHz. However, the operating range for data communication is from 108 to 862 MHz. 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. Only 8-MHz channels are used for data communication
Transit delay from headend to most distant customer	≤0.800 ms (typically much less)
Carrier-to-noise ratio in a 8 MHz band (analogue video level)	Not less than 44 dB (Note 4)
Carrier-to-interference ratio for total power (discrete and broadband ingress signals)	Not less than 52 dB within the design bandwidth

Table F.1 – Assumed downstream RF channel transmission characteristics for analogue TV and sound signals

Parameter	Value
Composite triple beat distortion for analogue modulated carriers	Not greater than –57 dBc within the design bandwidth (Note 6 a)
Composite second-order distortion for analogue modulated carriers	Not greater than –57 dBc within the design bandwidth (Note 6 b)
Cross-modulation level	Under consideration
Amplitude ripple (maximum)	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, –30 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 analogue video carrier level at the system outlet, inclusive of above signal level variation	77 dBµV (Note 6 c)
Lowest analogue video carrier level at the system outlet, inclusive of above signal level variation	60 dBµV (Note 6 d)

NOTE 1 – Transmission is from the headend combiner to the CM input at the customer location.

NOTE 2 – For measurements above, the normal downstream operating frequency band (except hum), impairments are referenced to the highest-frequency PAL/SECAM carrier level.

NOTE 3 – For hum measurements above, the normal downstream operating frequency band, a continuous-wave carrier is sent at the test frequency at the same level as the highest-frequency PAL/SECAM carrier.

NOTE 4 – This presumes that the average digital carrier is operated at the analogue peak carrier power level. When the digital carrier is operated below the analogue peak carrier level, this C/N may be less.

NOTE 5 – Measurements methods are defined in [b-EN 50083-7]

NOTE 6 – For SECAM systems, the following values apply:

a) Not greater than –52 dBc within the design bandwidth.

b) Not greater than –52 dBc within the design bandwidth.

c) 74 dBµV.

d) 57 dBµV.

F.4.3.2 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 F.2. All conditions are present concurrently.

Table F.2 – Assumed upstream RF channel transmission characteristics

Parameter	Value
Frequency range	5 up to 65 MHz edge to edge
Transit delay from the most distant CM to the nearest CM or CMTS	≤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 (Note 2)
Carrier-to-interference (the sum of noise, distortion, common-path distortion and cross-modulation) ratio in active channel	Not less than 22 dB
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 (Notes 3 and 4)
Amplitude ripple (maximum)	5-65 MHz: 2.5 dB in 2 MHz
Group delay ripple (maximum)	5-65 MHz: 300 ns in 2 MHz
Micro-reflections (maximum) single echo	–10 dBc @ ≤0.5 µs –20 dBc @ ≤1.0 µs –30 dBc @ >1.0 µs
Seasonal and diurnal signal level variation	Not greater than 12 dB min. to max.
NOTE 1 – Transmission is from the CM output at the customer location to the headend. NOTE 2 – 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. NOTE 3 – Amplitude and frequency characteristics sufficiently strong to partially or wholly mask the data carrier. NOTE 4 – Impulse noise levels more prevalent at lower frequencies (<15 MHz).	

F.4.3.2.1 Availability

Refer to clause 4.3.2.1.

F.4.4 Transmission levels

The nominal average power level of the downstream CMTS QAM signal(s) within an 8-MHz channel is targeted to be in the range –13 dBc to 0 dBc relative to the analogue peak video carrier level and will normally not exceed the analogue peak video carrier level (typically between –10 to –6 dBc for 64QAM, and between –6 to –4 dBc for 256QAM). 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.

F.4.5 Frequency inversion

Refer to clause 4.5.

F.5 Communication protocols

Refer to clause 5.

F.6 Physical media dependent sublayer specification

F.6.1 Scope

This clause applies to the second technology option referred to in clause 1.1, "scope". In those cases where the requirement for both technology options are identical, a reference is provided to the main text.

Whenever any reference in this clause to spurious emissions conflicts with any legal requirement for the area of operation, the latter shall take precedence.

F.6.2 Upstream

F.6.2.1 Overview

Refer to clause 6.2.1.

F.6.2.2 Signal processing requirements

Refer to clause 6.2.2.

F.6.2.3 Modulation formats

Refer to clause 6.2.3.

F.6.2.4 R-S encode

F.6.2.4.1 R-S encode modes

Refer to clause 6.2.4.1.

F.6.2.4.2 R-S bit-to-symbol ordering

Refer to clause 6.2.4.2.

F.6.2.5 R-S frame structure

Refer to clause 6.2.5.

F.6.2.5.1 R-S codeword length

Refer to clause 6.2.5.1.

F.6.2.5.1.1 Burst size

Refer to clause 6.2.5.1.1.

F.6.2.5.1.2 Fixed codeword length

Refer to clause 6.2.5.1.2.

F.6.2.5.1.3 Shortened last codeword

Refer to clause 6.2.5.1.3.

F.6.2.5.2 R-S FEC disabled

Refer to clause 6.2.5.2.

F.6.2.6 TDMA byte interleaver

Refer to clause 6.2.6.

F.6.2.6.1 Byte interleaver parameters

Refer to clause 6.2.6.1.

F.6.2.6.2 Interleaver operating modes

Refer to clause 6.2.6.2.

F.6.2.6.2.1 Fixed mode

Refer to clause 6.2.6.2.1.

F.6.2.6.2.2 Dynamic mode

Refer to clause 6.2.6.2.2.

F.6.2.6.2.2.1 Dynamic mode calculations

Refer to clause 6.2.6.2.2.1.

F.6.2.7 Scrambler (randomizer)

Refer to clause 6.2.7.

F.6.2.8 TCM encoder

Refer to clause 6.2.8.

F.6.2.8.1 Bytes to TCM symbol mapping

Refer to clause 6.2.8.1.

F.6.2.9 Preamble prepend

Refer to clause 6.2.9.

F.6.2.10 Modulation rates

Refer to clause 6.2.10.

F.6.2.11 S-CDMA framer and interleaver

Refer to clause 6.2.11.

F.6.2.11.1 S-CDMA framing considerations

Refer to clause 6.2.11.1.

F.6.2.11.2 Mini-slot numbering

Refer to clause 6.2.11.2.

F.6.2.11.2.1 Mini-slot numbering parameters in UCD

Refer to clause 6.2.11.2.1.

F.6.2.11.2.2 Mini-slot numbering examples

Refer to clause 6.2.11.2.2.

F.6.2.11.3 Transmission time

Refer to clause 6.2.11.3.

F.6.2.11.4 Latency considerations

Refer to clause 6.2.11.4.

F.6.2.11.5 Spreader-off bursts for maintenance on S-CDMA channel

Refer to clause 6.2.11.5.

F.6.2.11.6 Limiting the number of codes assigned to a CM

Refer to clause 6.2.11.6.

F.6.2.12 S-CDMA framer

Refer to clause 6.2.12.

F.6.2.12.1 Subframe definition

Refer to clause 6.2.12.1.

F.6.2.12.2 Framer operation

Refer to clause 6.2.12.2.

F.6.2.12.2.1 Rules for preamble and coded TCM symbols

Refer to clause 6.2.12.2.1.

F.6.2.12.2.2 Rules for uncoded symbols and the uncoded TCM subsymbols

Refer to clause 6.2.12.2.2.

F.6.2.12.2.3 Subframe example

Refer to clause 6.2.12.2.3.

F.6.2.12.2.4 Frame transmission

Refer to clause 6.2.12.2.4.

F.6.2.13 Symbol mapping

Refer to clause 6.2.13.

F.6.2.14 S-CDMA spreader

Refer to clause 6.2.14.

F.6.2.14.1 Code hopping

Refer to clause 6.2.14.1.

F.6.2.15 Transmit pre-equalizer

Refer to clause 6.2.15.

F.6.2.16 Spectral shaping

Refer to clause 6.2.16.

F.6.2.16.1 Upstream frequency agility and range

The upstream PMD sublayer MUST support operation over the frequency range of 5-65 MHz edge to edge.

Offset frequency resolution MUST be supported, having a range of ±32 kHz (increment = 1 Hz; implement within ±10 Hz).

F.6.2.16.2 Spectrum format

Refer to clause 6.2.16.2.

F.6.2.17 Relative processing delays

Refer to clause 6.2.17.

F.6.2.18 Transmit power requirements

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

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

The protocol by which power adjustments are performed is defined in clause 11.2.4. Such adjustments by the CM MUST be within the ranges of tolerances described below. A CM MUST confirm that the transmit power limits are met after a RNG-RSP is received or after a UCD change.

Transmit power is defined as the average RF power in the occupied bandwidth (channel width) transmitted in the data symbols of a burst, assuming equally likely QAM symbols, measured at the F-connector of the CM.

Maximum and minimum transmit power level requirements refer to the CM's target transmit power level, defined as the CM's estimate of its actual transmit power. The actual transmitted power MUST be within ±2 dB of the target power. The target transmit power MUST be variable over the range specified in Table F.5.

Transmit power as reported by the CM in the MIB is referenced to the 64QAM constellation. When transmitting with other constellations, a slightly different transmit power will result, depending on the constellation gain in Table F.3 (see clause 6.2.13). As an example, if the reported power is 90 dBµV, 64QAM will be transmitted with a target power of 90 dBµV, while QPSK will be transmitted with 88.82 dBµV.

Table F.3 – Constellation gains and power limits

Constellation	Constellation Gain G_{const} relative to 64QAM (dB)	P_{min} (dBµV)	P_{max} (dBµV) TDMA	P_{max} (dBµV) S-CDMA	$P_{min} - G_{const}$ (dBµV)	$P_{max} - G_{const}$ (dBµV) TDMA	$P_{max} - G_{const}$ (dBµV) S-CDMA
QPSK	–1.18	68	118	113	69.18	119.18	114.18
8QAM	–0.21	68	115	113	68.21	115.21	113.21
16QAM	–0.21	68	115	113	68.21	115.21	113.21
32QAM	0.00	68	114	113	68.00	114.00	113.00
64QAM	0.00	68	114	113	68.00	114.00	113.00
128QAM	0.05	68	N/A	113	67.95	N/A	112.95

The actual transmitted power within a burst MUST be constant to within 0.1 dB peak to peak. 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 clauses F.6.2.18.1 and F.6.2.18.2

F.6.2.18.1 TDMA transmit power calculations

In TDMA mode, the CM determines its target transmit power P_t as follows. Define:

P_r = Reported power level (dBµV) of CM in MIB (refers to 64QAM constellation)

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

G_{const} = Constellation gain (dB) relative to 64QAM constellation (see Table F.3)

P_{min} = Minimum target transmit power permitted for the CM per clause 6.2.21.1 (see Table F.3)

P_{max} = Maximum target transmit power permitted for the CM per clause 6.2.21.1 (see Table F.3)

P_{hi} = $\min(P_{max} – G_{const})$ over all burst profiles used by the CM (see Table F.3)

P_{low} = $\max(P_{min} – G_{const})$ over all burst profiles used by the CM (see Table F.3)

P_t = Target transmit power level (dBµV) of CM (actual transmitted power as estimated by CM)

The CM updates its reported power by the following steps:

1) $P_r = P_r + \Delta P$ //Add power level adjustment to reported power level

2) $P_r = \min[P_r, P_{hi}]$ //Clip at max power limit

3) $P_r = \max[P_r, P_{low}]$ //Clip at min power limit

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

Usually the reported power level is a relatively constant quantity, while the transmitted power level varies dynamically as different burst profiles, with different constellation gains, are transmitted. A CM's target transmit power MUST never be below P_{min} or above P_{max}. This implies that, in some cases, the extreme transmit power levels (e.g., 118 dBµV for QPSK and 68 dBµV) may not be permitted if burst profiles with multiple constellations are active. Also, if only QPSK is used, the reported power may be greater than 118 dBµV, although the target transmit power will not exceed 118 dBµV.

For example, if only QPSK and 64QAM burst profiles are active, P_{hi} = 114 dBµV and P_{low} = 69.2 dBµV. The maximum permitted QPSK transmitted power is 114 – 1.2 = 112.8 dBµV, the minimum QPSK power is 69.2 – 1.2 = 68 dBµV, the maximum 64QAM power is 114 dBµV, and the minimum 64QAM power is 69.2 dBµV.

F.6.2.18.2 S-CDMA transmit power calculations

Refer to clause 6.2.18.2.

F.6.2.18.2.1 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 P_t as follows. Define:

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

P_{hi} = $\min[P_{max} – G_{const}]$ over all burst profiles used by the CM (see Table F.3)

P_{low} = $\max[P_{min} – G_{const}] + 10 \log(number\_active\_codes/number\_of\_codes\_per\_mini\text{-}slot)$ where the maximum is over all burst profiles used by the CM (see Table F.3)

The CM updates its reported power by the following steps:

1) $P_r = P_r + \Delta P$ //Add power level adjustment to reported power level

2) $P_r = \min[P_r, P_{hi}]$ //Clip at max power limit

3) $P_r = \max[P_r, P_{low}]$ //Clip at min power limit

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

$$P_{t,i} = P_r + G_{const,i} – 10 \log(number\_active\_codes)$$

i.e., the reported power 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 P_t in a frame is the sum of the individual transmit powers $P_{t,i}$ of each code, where the sum is performed using absolute power quantities (non-dB domain).

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

The transmitted power level 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 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 (e.g., 68 and 113 dBμV) may not be permitted. Also if, for example, only QPSK is used, the reported power may be greater than 113 dBμV, although the target transmit power will not exceed 113 dBμV.

If, for example, QPSK and 64QAM burst profiles are active, the number of active codes is 128 and the number of codes per mini-slot is 2, then P_{hi} = 113 dBμV and P_{low} = 87.24 dBμV. The maximum permitted QPSK transmitted power is 113 – 1.18 = 111.82 dBμV when all active codes are transmitted, the minimum QPSK power is 87.24 – 1.18 – 10log(128) + 10log(2) = 68 dBμV 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 64QAM power is 113 dBμV when all active codes are transmitted, and the minimum 64QAM power is 87.24 – 10 log(128) + 10 log(2) = 69.18 dBμV with one mini-slot transmitted. The minimum QPSK power permitted while transmitting, for example, 2 mini-slots is 71 dBμV, and the minimum 64QAM power permitted while transmitting 2 mini-slots is 72.2 dBμV.

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.

F.6.2.18.2.2 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 P_t as follows. Define:

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

P_{hi_S} = min[113 – G_{const}] over all spreader-on burst profiles used by the CM (see Table F.3)

P_{low_S} = max[68 – G_{const}] + 10 log(*number_active_codes*/*number_of_codes_per_mini-slot*) where the maximum is over all burst profiles used by the CM (see Table F.3)

P_{max_T} = maximum target transmit power permitted for the CM in TDMA mode (see Table F.3) 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 (see Table F.3)

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

P_{sf} = CM power shortfall

P_{hr} = S-CDMA power headroom

ΔP = power level adjustment in dB sent from CMTS to CM

The CM updates its reported power by the following steps:

1) $P_r = P_r + \Delta P$ //Add power level adjustment to reported power level

2) $P_r = \min[P_r, P_{hi\_T}]$ //Clip at max TDMA power limit

3) $P_r = \max[P_r, P_{low\_S}]$ //Clip at min S-CDMA power limit

4) $P_{on} = \min[P_r, P_{hi\_S}]$ //Clip at max S-CDMA power limit

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

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

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

1) $P_{sf} = P_r + \max[G_{const,i}] - 113$ //Difference between spreader-off reported and max spreader-on target powers

2) $P_{sf} = \max[P_{sf}, 0]$ //Set P_{sf} to 0 if P_t is less than 113 dBμV

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

$$P_{t,i} = P_{on} + G_{const,i} - 10 \log(number\_active\_codes) + P_{hr}$$

i.e., the clipped reported power 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 codes limit and can transmit additional power per code. The total transmit power P_t in a frame is the sum of the individual transmit powers $P_{t,i}$ of each code, 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, the burst profile contains QPSK for IUCs 1, 2, 3 and 4 and 64QAM 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} = 113 dBμV, P_{low_S} = 87.24 dBμV, and P_{hi_T} = 118 dBμV. Assume the CM ranges at spreader-off target transmit power of 117 dBμV. The CM reports P_{sf} = 117 dBμV – 113 dBμV = 4 dB. The CMTS uses P_{sf} to set (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 to:

$P_{ti} = P_{on} + G_{const,i} - 10 \log(number\_active\_codes) + P_{hr}$

= 113 dBμV + 0 dB – 21 dB + 6 dB //For a code with 64QAM modulation

= 98 dBμV

A parameter that may be used to illustrate the effect of increased power per code is the effective transmit power, P_{eff}, the power that would result hypothetically if all N_{act} active codes were transmitted. It is computed as:

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

$$= P_{on} + P_{hr} + 10 \log \frac{1}{N_{act}} \sum_{i=1}^{N_{act}} 10^{G_{const,i}/10}$$

where the last term is the average constellation gain.

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

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

Continuing the above example, the result is:

$$P_{eff} = 113 \text{ dB}\mu\text{V} + 6 \text{ dB}$$
$$= 119 \text{ dB}\mu\text{V}$$

Limiting the number of codes has given the CM an enhanced effective power of 119 dBμV, which is 6 dB above the normal maximum of 113 dBμV, and 2 dB above the ranging power of 117 dBμV. In this example, the CMTS used its discretion to ask for 2 dB more enhancement 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. 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 that resulted in the SNR measurement. In the MIB, *effective_SNR* corresponds to a reference case with 64QAM 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 the last station maintenance was 17 dB using QPSK modulation ($G_{const,ranging}$ = –1.18 dB), then the effective SNR referenced to 64QAM modulation is:

$$effective\_SNR = 17 \text{ dB} - 4 \text{ dB} + 6 \text{ dB} + 1.18 \text{ dB} = 20.18 \text{ dB}$$

F.6.2.18.3 Transmit power step size

The step resolution in transmit power MUST be 1 dB or less. When a CM is commanded with finer resolution than it can implement, it MUST round to the nearest supported step size. If the commanded step is halfway between two supported step sizes, the CM MUST choose the smaller step. For example, with a supported step resolution of 1 dB, a command to step ±0.5 dB would result in no step, while a command to step ±0.75 dB would result in a ±1 dB step.

The step size accuracy MUST be within ±0.4 dB. For example, the actual power increase resulting from a command to increase the power level by 1 dB in a CM's next transmitted burst MUST be between 0.6 dB and 1.4 dB.

A relaxation in step size accuracy to ±1.4 dB is allowed for one gain change when changing the power throughout the full power control range in either direction (from low-end to high-end power and vice versa). The locations of these two gain changes with relaxed accuracy MUST be at least 2 dB apart, thus enabling the use of large step attenuators in the coverage of the full power control range (hysteresis effect).

F.6.2.19 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:

i) the modulation rate (six rates from 160 ksymb/s to 5.12 Msymb/s in octave steps);

ii) the centre frequency (Hz);

iii) the 1536-bit preamble superstring; and

iv) the S-CDMA channel parameters.

The channel parameters are further described in Table 8-18; these characteristics are shared by all users on a given channel. The burst profile attributes are listed in Table F.4 and are further described in Table 8-19; these parameters are the shared attributes corresponding to a burst type. 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 F.5.

Table F.4 – Burst profile attributes

Burst profile attributes	Configuration settings
Modulation	QPSK, 8QAM, 16QAM, 32QAM, 64QAM, 128QAM (TCM only)
Differential encoding	On/off
TCM encoding	On/off
Preamble length	0-1536 bits (see clause F.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)[a)]	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
Interleaver width (N_r) (R-S codeword length, k + 2 × T)	18 to 255
Byte interleaver depth (I_r)[b)]	0 to floor ($2048/N_r$)[c)]
Byte interleaver block size (B_r)[d)]	2 × N_r to 2048
Interleaver packet size (N_f), (in bytes, including FEC)	≥18 bytes
Preamble type	QPSK0/QPSK1
S-CDMA spreader[e)]	On/off
S-CDMA codes per subframe[e)]	1 to 128

Table F.4 – Burst profile attributes

Burst profile attributes	Configuration settings
S-CDMA interleaver step[e)]	1 to (spreading intervals per frame – 1)

a) 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.

b) If depth = 1, no interleaving; if depth = 0, dynamic mode.

c) N_r is the R-S codeword size k + 2T as defined in clause F.6.2.1.

d) Used only in dynamic mode.

e) Used only for S-CDMA channels.

Table F.5 – User-unique burst parameters

User-unique parameter	Adjustment command	
Power level	8-bit two's complement, resolution = 0.25 dB	TDMA: +68 to +114 dBµV (32QAM, 64QAM) +68 to +115 dBµV (8QAM, 16QAM) +68 to +118 dBµV (QPSK) S-CDMA: +68 to +113 dBµV (all modulations) Resolution + 1 dB or better
Offset frequency (Note)	Range = ±32 kHz, resolution = 1 Hz	Frequency range per clause F.6.2.16.1 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 clause F.4.1, broadband access network 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 clause F.6.2.15, transmit pre-equalizer)	DOCSIS 2.0 mode: 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 2.0 mode: 24 complex coefficients DOCSIS 1.1 mode: up to 64 complex coefficients

The CM MUST generate each burst at the appropriate time as conveyed in the mini-slot grants provided by the CMTS MAPs (see clause 8.3.4).

The CM MUST support all burst profiles commanded by the CMTS via the burst descriptors in the UCD (see clause 8.3.3), and subsequently assigned for transmission in a MAP (see clause 8.3.4).

F.6.2.19.1 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 ±0.25 µs plus ±1/2 symbol owing to resolution. As an example, for the maximum symbol rate of 5.12 Msymb/s, the corresponding symbol period would be 195 ns, the corresponding maximum resolution for the timing adjust MUST be 195 ns, and the corresponding minimum accuracy MUST be ±348 ns. The accuracy of CM burst timing of ±0.25 µs plus ±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 MHz) = 6.25 µs/64 ~= 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 µs/(64*256) ~= 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 Mchips/s, the corresponding maximum resolution for implementation of the timing correction would be (±0.01)*195 ns or roughly ±2 ns.

F.6.2.19.2 TDMA reconfiguration times

Refer to clause 6.2.19.2.

F.6.2.19.3 S-CDMA reconfiguration times

Refer to clause 6.2.19.3.

F.6.2.19.4 CM timing offsets when changing modulation rate

Refer to clause 6.2.19.4.

F.6.2.20 Burst timing convention

Refer to clause 6.2.20.

F.6.2.21 Fidelity requirements

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

F.6.2.21.1 Spurious emissions

The spurious emissions specifications are separated into two regions based on the transmit power. Region 1 is defined to have a power range of +74 dBµV to (P_{max} – 3), i.e., the central region. Region 2 is defined from +68 dBµV to +74 dBµV and (P_{max} – 3) to P_{max}, i.e., the low and high ends of the transmit power. P_{max} depends on the modulation order, per Table F.5, as follows: for TDMA, +118 dBµV for QPSK, +115 dBµV for 8QAM/16QAM, +114 dBµV for 32QAM/64QAM, and +113 dBµV for all modulations of S-CDMA.

For S-CDMA mode, when a modem is transmitting fewer than 4 spreading codes, the region 2 specifications are used for all transmit power levels. Otherwise, for all other numbers of spreading codes (e.g., 4 to 128) or for TDMA mode, the spurious emissions specifications are used according to the power ranges defined for regions 1 and 2 above.

In addition, for S-CDMA, the spurious emission specifications for S-CDMA MUST be met for any *number_allocated_codes*, as defined in clause F.6.2.19.

The noise and spurious power MUST NOT exceed the levels given in Tables F.6, F.7 and F.8.

Table F.6 – Spurious emissions

Parameter	Transmitting Burst	Between Bursts
Inband	–40 dBc	The greater of –72 dBc or +1 dBμV
Adjacent band	See Table F.7	The greater of –72 dBc or +1 dBμV
3 or fewer carrier-related frequency bands (such as second harmonic, if <65 MHz)	Region 1: –50 dBc for transmitted modulation rate = 320 kHz and above; –47 dBc for transmitted modulation rate = 160 kHz Region 2: –47 dBc	The greater of –72 dBc or +1 dBμV
Bands within 5 to 65 MHz (excluding assigned channel, adjacent channels and carrier-related channels)	See Table F.8	The greater of –72 dBc or +1 dBμV
CM integrated spurious emissions limits (all in 250 kHz, includes discretes) 87.5 to 108 MHz	30 dBμV	1 dBμV
CM integrated spurious emissions limits (all in 4.75 MHz, includes discretes[a]) 65 to 87.5 MHz 108 to 136 MHz 136 to 862 MHz	 max –40 dBc, 34 dBμV 20 dBμV 15 dBμV	 34 dBμV 15 dBμV max (15 dBμV, –40 dB ref d/s[b])
CM discrete spurious emissions limits[c]) 65 to 87.5 MHz 108 to 862 MHz	 max –50 dBc, 24 dBμV 10 dBμV	 24 dBμV 10 dBμV

a) The frequencies from 108 to 136 MHz may be forbidden due to national regulations.

b) "dB ref d/s" is relative to the received downstream signal level. Some spurious outputs are proportional to the receive signal level.

c) These specification limits exclude a single discrete spur related to the tuned received channel; this single discrete spur MUST NOT be greater than 20 dBμV.

Table F.7 – Adjacent channel spurious emissions relative to the transmitted burst power level

Transmitted carrier modulation rate	Specification in the interval, region 1	Specification in the interval, region 2	Measurement interval and distance from carrier edge	Adjacent channel carrier modulation rate
All modulation rates	–47 dBc	–45 dBc	20 kHz to 180 kHz	160 ksymb/s
	–47 dBc	–45 dBc	40 kHz to 360 kHz	320 ksymb/s
	–46 dBc	–45 dBc	80 kHz to 720 kHz	640 ksymb/s
	–45 dBc	–44 dBc	160 kHz to 1440 kHz	1280 ksymb/s
	–44 dBc	–41 dBc	320 kHz to 2880 kHz	2560 ksymb/s
	–42 dBc	–38 dBc	640 kHz to 5760 kHz	5120 ksymb/s

Table F.8 – Spurious emissions in 5 to 65 MHz relative to the transmitted burst power level

Possible modulation rate in this interval	Specification in the interval, region 1	Specification in the interval, region 2	Initial measurement interval and distance from carrier edge
160 ksymb/s	–54 dBc	–53 dBc	220 kHz to 380 kHz
320 ksymb/s	–52 dBc	–50 dBc	240 kHz to 560 kHz
640 ksymb/s	–50 dBc	–47 dBc	280 kHz to 920 kHz
1280 ksymb/s	–48 dBc	–44 dBc	360 kHz to 1640 kHz
2560 ksymb/s	–46 dBc	–41 dBc	520 kHz to 3080 kHz
5120 ksymb/s	–44 dBc	–38 dBc	840 kHz to 5960 kHz

In Table F.6, 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., 160 to 5120 kHz). All requirements expressed in dBc are relative to the actual transmit power that the CM emits.

The measurement bandwidth for the three (or fewer) carrier-related frequency bands (below 65 MHz) is 160 kHz, with up to three 160 kHz bands, each with no more than the value given in Table F.6, allowed to be excluded from the bands within 5 to 65 MHz transmitting burst specifications of Table F.8. Carrier-related spurious emissions include all products whose frequency is a function of the carrier frequency of the upstream transmission such as, but not limited to, carrier harmonics.

The measurement bandwidth is also 160 kHz for the between bursts specifications of Table F.6 below 65 MHz.

The transmitting burst specifications 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 specifications apply except during a used grant of mini-slots, 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 microseconds at the 5.12 Msymb/s rate, or as short as 200 microseconds at the 160 ksymb/s rate.

These specification limits exclude three or fewer discrete spurs. Such spurs must not be greater than 20 dBμV.

F.6.2.21.1.1 Adjacent channel spurious emissions

Spurious emissions from a transmitted carrier may occur in an adjacent channel that could be occupied by a carrier of the same or different modulation rate. Table F.7 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 that is 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.

F.6.2.21.1.2 Spurious emissions in 5 to 65 MHz

Spurious emissions, other than those in an adjacent channel or carrier-related emissions listed above, may occur in intervals (frequency bands) that could be occupied by other carriers of the same or different modulation rates. To accommodate these different modulation rates and associated bandwidths, the spurious emissions are measured in an interval equal to the bandwidth corresponding to the modulation rate of the carrier that could be transmitted in that interval. This interval is independent of the current transmitted modulation rate.

Table F.8 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, 5 MHz or 65 MHz, is reached. Measurement intervals should not include the three or fewer carrier related emission bands excluded above.

F.6.2.21.2 Spurious emissions during burst on/off transients

Each transmitter MUST control spurious emissions, prior to and during ramp-up and during and following ramp-down, 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 μs of constant slewing. This requirement applies when the CM is transmitting at +115 dBμV or more; 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 from +115 dBμV, down to a maximum change of 7 mV at +91 dBμV and below. This requirement does not apply to CM power-on and power-off transients.

F.6.2.21.3 Modulation error ratio (MER)

Refer to clause 6.2.21.3.

F.6.2.21.3.1 Definitions

Refer to clause 6.2.21.3.1.

F.6.2.21.3.2 Requirements

Unless otherwise stated, the MER MUST meet or exceed the following limits over the full transmit power range of Table 6-8 for each modulation, each 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 MHz + modulation rate × 1.25/2] to [65 MHz – 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:

$MER_{symb} \geq 30$ dB over 15 to 47 MHz carrier frequency

$MER_{symb} \geq 27$ dB over 10 to 15 MHz and 47 to 54 MHz carrier frequency

$MER_{symb} \geq 23$ dB over 5 to 10 MHz and 54 to 65 MHz carrier frequency

Case 1b: For modulation rate 5.12 MHz:

$MER_{symb} \geq 27$ dB over 15 to 47 MHz carrier frequency

$MER_{symb} \geq 24$ dB over 10 to 15 MHz and 47 to 54 MHz carrier frequency

$MER_{symb} \geq 20$ dB over 5 to 10 MHz and 54 to 65 MHz carrier frequency

Case 2: Flat channel, transmit equalization ON

Case 2a: For TDMA/QPSK, $MER_{symb} \geq 30$ dB.

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

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

Case 3: Echo channel, transmit equalization ON

Case 3a: In the presence of a single echo selected from the channel micro-reflections defined in Table 4-2, the measured MER_{symb} MUST be ≥ 33 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 4-2 (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.

F.6.2.21.4 Filter distortion

Refer to clause 6.2.21.4.

F.6.2.21.4.1 Amplitude

Refer to clause 6.2.21.4.1.

F.6.2.21.4.2 Phase

Refer to clause 6.2.21.4.2.

F.6.2.21.5 Carrier phase noise

Refer to clause 6.2.21.5.

F.6.2.21.6 Channel frequency accuracy

Refer to clause 6.2.21.6.

F.6.2.21.7 Modulation rate accuracy

Refer to clause 6.2.21.7.

F.6.2.21.8 Modulation timing jitter

F.6.2.21.8.1 Symbol timing jitter for asynchronous operation

Refer to clause 6.2.21.8.1.

F.6.2.21.8.2 Chip timing jitter for synchronous operation

All jitter specifications assume a downstream input to the CM per clauses F.6.3.5, F.6.3.6, F.6.3.7.2, F.6.3.7.3, F.6.3.9 and F.6.3.10.

For S-CDMA mode, upstream chip clock timing error (with the mean error subtracted out) relative to the CMTS master clock MUST be less than 0.005 RMS of the chip period over a 35-second measurement interval. This applies:

1) to 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 chip 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 chip 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-bursted) upstream signal is transmitted at the commanded carrier frequency, modulation rate and level.
- The chip sequence at the spreader output 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 upstream symbol clock as in normal synchronous operation.

F.6.2.22 Upstream demodulator input power characteristics

The maximum total input power to the upstream demodulator MUST NOT exceed 95 dBμV in the 5-65 MHz frequency range of operation.

The intended received power in each carrier MUST be within the values shown in Table F.9.

Table F.9 – Maximum range of commanded nominal receive power in each carrier

Modulation rate (ksymb/s)	Maximum range (dBμV)
160	+44 to +74
320	+47 to +77
640	+50 to +80
1280	+53 to +83
2560	+56 to +86
5120	+59 to +89

The demodulator MUST operate within its defined performance specifications with received bursts within ±6 dB of the nominal commanded received power.

F.6.2.23 Upstream electrical output from the CM

The CM MUST output an RF modulated signal with the characteristics delineated in Table F.10.

Table F.10 – Electrical output from CM

Parameter	Value
Frequency	5 to 65 MHz edge to edge
Level range (one channel)	TDMA: +68 to +114 dBµV (32QAM, 64QAM) +68 to +115 dBµV (8QAM, 16QAM) +68 to +118 dBµV (QPSK) S-CDMA: +68 to +113 dBµV (all modulations of S-CDMA)
Modulation type	QPSK, 8QAM, 16QAM, 32QAM, 64QAM and 128QAM
Modulation rate (nominal)	TDMA: 160, 320, 640, 1280, 2560 and 5120 ksymb/s S-CDMA: 1280, 2560 and 5120 ksymb/s
Channel bandwidth	TDMA: 200, 400, 800, 1600, 3200 and 6400 kHz S-CDMA: 1600, 3200 and 6400 kHz
Nominal output impedance	75 ohms
Output return loss	>6 dB (5-65 MHz)
Connector	F connector per [IEC 60169-24] (common with the input)

F.6.3 Downstream

This DOCSIS 2.0 downstream specification applies only to a CMTS supporting exactly one QAM channel per RF output port. A CMTS supporting more than one QAM channel per RF output port would instead conform to the European technology option of the DOCSIS downstream radio frequency interface as specified in Annex A of [DOCS9] in lieu of this clause.

F.6.3.1 Downstream protocol

The downstream PMD sublayer MUST conform to [b-EN 300 429].

F.6.3.2 Interleaving

The downstream PMD sublayer MUST support the interleaver with the characteristics defined in Table F.11. This interleaver mode fully complies with [b-EN 300 429].

Table F.11 – Interleaver characteristics

I (number of taps)	J (increment)	Burst protection 64QAM/256QAM	Latency 64QAM/256QAM
12	17	18 µs/14 µs	0.43 ms/0.32 ms

F.6.3.3 Downstream frequency plan

The downstream frequency plan will include all centre frequencies between 112 and 858 MHz on 250 kHz increments. It is up to the operator to decide which frequencies to use to meet national and network requirements.

F.6.3.4 CMTS electrical output

The CMTS MUST output an RF modulated signal with the following characteristics defined in Table F.12.

Table F.12 – CMTS output

Parameter	Value
Centre frequency (f_c)	112 to 858 MHz ± 30 kHz
Level	Adjustable over the range 110 to 121 dBμV
Modulation type	64QAM and 256QAM
Symbol rate (nominal) 64QAM 256QAM	 6.952 Msymb/s 6.952 Msymb/s
Nominal channel spacing	8 MHz
Frequency response 64QAM 256QAM	 ~0.15 square root raised cosine shaping ~0.15 square root raised cosine shaping
Total discrete spurious inband (f_c ± 4MHz) Inband spurious and noise (f_c ± 4 MHz) Adjacent channel (f_c ± 4.0 MHz) to (f_c ± 4.75 MHz) Adjacent channel (f_c ± 4.75 MHz) to (f_c ±12 MHz) Next adjacent channel (f_c ± 12 MHz) to (f_c ± 20 MHz) Other channels (80 MHz to 1000 MHz)	<–57 dBc <–46.7 dBc; where channel spurious and noise includes all discrete spurious, noise, carrier leakage, clock lines, synthesizer products and other undesired transmitter products. Noise within ±50 kHz of the carrier is excluded. <–58 dBc in 750 kHz <–60.6 dBc in 7.25 MHz, excluding up to three spurs, each of which must be <–60 dBc when each is measured with 10 kHz bandwidth. Less than the greater of –63.7 dBc or 49.3 dBμV in 8 MHz, excluding up to three discrete spurs. The total power in the spurs must be <–60 dBc when each is measured with 10 kHz bandwidth. <49.3 dBμV in each 8-MHz channel, excluding up to three discrete spurs. The total power in the spurs must be <–60 dBc when each is measured with 10 kHz bandwidth.
Phase noise	1 kHz-10 kHz: –33 dBc double-sided noise power 10 kHz-50 kHz: –51 dBc double-sided noise power 50 kHz-3 MHz: –51 dBc double-sided noise power
Output impedance	75 ohms
Output return loss	>14 dB within an output channel up to 750 MHz; >13 dB in an output channel above 750 MHz
Connector	F connector per [IEC 60169-24]

F.6.3.5 Downstream electrical input to CM

The CM MUST accept an RF modulated signal with the following characteristics (see Table F.13).

Table F.13 – Electrical input to CM

Parameter	Value
Centre frequency	112 to 858 MHz ± 30 kHz
Level range (one channel)	43 to 73 dBμV for 64QAM 47 to 77 dBμV for 256QAM
Modulation type	64QAM and 256QAM
Symbol rate (nominal)	6.952 Msymb/s (64QAM) and 6.952 Msymb/s (256QAM)
Bandwidth	8 MHz (alpha = 0.15 square root raised cosine shaping for 64QAM and alpha = 0.15 square root raised cosine shaping for 256QAM)
Total input power (80-862 MHz)	<90 dBμV
Input (load) impedance	75 ohms
Input return loss	>6 dB (85 to 862 MHz)
Connector	F connector per [IEC 60169-24] (common with the output)

F.6.3.6 CM BER performance

The bit-error-rate performance of a CM MUST be as described in this clause. The requirements apply to the I = 12, J = 17 mode of interleaving.

F.6.3.6.1 64QAM

F.6.3.6.1.1 64QAM 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) of 25.5 dB or greater.

F.6.3.6.1.2 64QAM image rejection performance

Performance as described in clause F.6.3.6.1.1 MUST be met with an analogue or digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

F.6.3.6.1.3 64QAM adjacent channel performance

Performance as described in clause F.6.3.6.1.1 MUST be met with a digital signal at 0 dBc in the adjacent channels.

Performance as described in clause F.6.3.6.1.1 MUST be met with an analogue signal at 10 dBc in the adjacent channels.

Performance as described in clause F.6.3.6.1.1, with an additional 0.2-dB allowance, MUST be met with a digital signal at +10 dBc in the adjacent channels.

F.6.3.6.2 256QAM

F.6.3.6.2.1 256QAM CM BER performance

Implementation loss of the CM MUST be 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 in Table F.14.

Table F.14 – 256QAM CM BER performance

Input receive signal level	E_s/N_o
47 dBµV to 54 dBµV	34.5 dB
>54 to +77 dBµV	31.5 dB

F.6.3.6.2.2 256QAM image rejection performance

Performance as described in clause F.6.3.6.2.1 MUST be met with an analogue or digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

F.6.3.6.2.3 256QAM adjacent channel performance

Performance as described in clause F.6.3.6.2.1 MUST be met with an analogue or digital signal at 0 dBc in the adjacent channels.

Performance as described in clause F.6.3.6.2.1, with an additional 0.5-dB allowance, MUST be met with an analogue signal at +10 dBc in the adjacent channels.

Performance as described in clause F.6.3.6.2.1, with an additional 1.0-dB allowance, MUST be met with a digital signal at +10 dBc in the adjacent channels.

F.6.3.6.2.4 Additional specifications for QAM

The following additional specifications are given for the QAM modulation.

Parameter	Specification
I/Q phase offset	<1.0º
I/Q crosstalk	–50 dB
I/Q amplitude imbalance	0.05 dB max
I/Q timing skew	<3.0 ns

F.6.3.7 CMTS timestamp jitter

Refer to clause 6.3.7.

F.6.3.7.1 CMTS master clock jitter for asynchronous operation

Refer to clause 6.3.7.1.

F.6.3.7.2 CMTS master clock jitter for synchronous operation

Refer to clause 6.3.7.2.

F.6.3.7.3 CMTS master clock frequency drift for synchronous operation

Refer to clause 6.3.7.3.

F.6.3.8 CMTS clock generation

The CMTS has the following three options related to the synchronization of the CMTS master clock and the downstream symbol clock:

1) Not locked.
2) Downstream symbol clock locked to CMTS master clock.
3) CMTS master clock locked to downstream symbol clock.

For S-CDMA operation, the master clock and the downstream symbol clock MUST be locked using either option 2 or 3.

Let f_b' represent the rate of the downstream symbol clock, which is locked to the CMTS master clock, and let f_m' represent the rate of the CMTS master clock locked to the downstream symbol clock. Let f_b represent the nominal specified downstream symbol rate and let f_m represent the nominal CMTS master clock rate (10.24 MHz).

With the downstream symbol clock locked to the CMTS master clock, the following equation MUST hold:

$$f_b' = f_m \times M/N$$

With the CMTS master clock locked to the downstream symbol clock, the following equation MUST hold:

$$f_m' = f_b \times N/M$$

M and N MUST be unsigned integer values, each capable of being represented in 16 bits. (These are specified in the channel TLV parameters of the UCD). When the downstream symbol clock and the CMTS master clock are not locked together (sync mode = 0), the values of M and N are not valid and are ignored by the CM.

The values of M and N MUST result in a value of f_b' or f_m', which is not more than ±1 ppm from its specified nominal value. Table F.15 lists the downstream modes of operation, their associated nominal symbol rates, f_b, example values for M and N, the resulting synchronized clock rates, and their offsets from their nominal values.

Table F.15 – Downstream symbol rates and example parameters for synchronization with the CMTS master clock

Downstream mode	Nominal specified symbol rate, f_b (MHz)	M/N	CMTS master clock rate, f_m' (MHz)	Downstream symbol rate, f_b' (MHz)	Offset from nominal
Annex A, 64QAM and 256QAM (8 MHz)	6.952	869/1280	10.24	6.952	0 ppm

F.6.3.9 Downstream symbol clock jitter for synchronous operation

The downstream symbol clock MUST meet the following double sideband phase noise requirements over the following frequency ranges:

< [–50 + 20 × log(f_{DS}/6.952)] dBc (i.e., < 0.07 ns RMS)	10 Hz to 100 Hz
< [–50 + 20 × log(f_{DS}/6.952)] dBc (i.e., < 0.07 ns RMS)	100 Hz to 1 kHz
< [–50 + 20 × log(f_{DS}/6.952)] dBc (i.e., < 0.07 ns RMS)	1 kHz to 10 kHz
< [–33 + 20 × log(f_{DS}/6.952)] dBc (i.e., < 0.5 ns RMS)	10 kHz to 100 kHz
< [–27 + 20 × log(f_{DS}/6.952)] dBc (i.e., < 1 ns RMS)	100 kHz to (f_{DS}/2)

where f_{DS} is the frequency of the measured clock in MHz. The value of f_{DS} MUST be an integral multiple or divisor of the downstream symbol clock. For example, an f_{DS} = 27.808-MHz clock may be measured if there is no explicit 6.952 MHz clock available.

The CMTS MUST provide a test mode in which:

- the downstream QAM symbol sequence is replaced with an alternating binary sequence (1, –1, 1, –1, 1, –1, ...) at nominal amplitude, on both I and Q;
- the CMTS generates the downstream symbol clock from the 10.24-MHz reference clock as in normal synchronous operation.

If an explicit downstream symbol clock that is capable of meeting the above phase noise requirements is available (e.g., a smooth clock without clock domain jitter), this test mode is not required.

F.6.3.10 Downstream symbol clock drift for synchronous operation

Refer to clause 6.3.10.

F.7 Downstream transmission convergence sublayer

F.7.1 Introduction

Refer to clause 7.1.

F.7.2 MPEG packet format

Refer to clause 7.2.

F.7.3 MPEG header for Euro-DOCSIS data-over-cable

The format of the MPEG transport stream header is defined in clause 2.4 of [ITU-T H.222.0]. The particular field values that distinguish data-over-cable MAC streams are defined in Table F.16. Field names are from [ITU-T H.222.0].

The MPEG header consists of 4 bytes that begin the 188-byte MPEG packet. The format of the header for use on an Euro-DOCSIS data-over-cable PID is restricted to that shown in Table F.16. The header format conforms to the MPEG standard, but its use is restricted in this Recommendation to NOT ALLOW inclusion of an adaptation_field in the MPEG packets.

Table F.16 – MPEG Header format for Euro-DOCSIS data-over-cable packets

Field	Length (bits)	Description
sync_byte	8	0x47; MPEG packet sync byte
transport_error_indicator	1	Indicates an error has occurred in the reception of the packet. This bit is reset to zero by the sender, and set to one whenever an error occurs in transmission of the packet.
payload_unit_start_indicator	1	A value of one indicates the presence of a pointer_field as the first byte of the payload (fifth byte of the packet)
transport_priority	1	Reserved; set to zero
PID	13	Euro-DOCSIS data-over-cable well-known PID (0x1FFE)
transport_scrambling_control	2	Reserved; set to '00'
adaptation_field_control	2	'01'; use of the adaptation_field is NOT ALLOWED on the Euro-DOCSIS PID
continuity_counter	4	Cyclic counter within this PID

F.7.4 MPEG payload for Euro-DOCSIS data-over-cable

The MPEG payload portion of the MPEG packet will carry the Euro-DOCSIS MAC frames. The first byte of the MPEG payload will be a 'pointer_field' if the payload_unit_start_indicator (PUSI) of the MPEG header is set.

F.7.4.1 stuff_byte

This Recommendation defines a stuff_byte pattern having a value (0xFF) that is used within the Euro-DOCSIS payload to fill any gaps between the Euro-DOCSIS MAC frames. This value is chosen as an unused value for the first byte of the Euro-DOCSIS MAC frame. The 'FC' byte of the MAC header will be defined to never contain this value (FC_TYPE = '11' indicates a MAC-specific frame, and FC_PARM = '11111' is not currently used and, according to this Recommendation, is defined as an illegal value for FC_PARM).

F.7.4.2 pointer_field

The pointer_field is present as the fifth byte of the MPEG packet (first byte following the MPEG header) whenever the PUSI is set to one in the MPEG header. The interpretation of the pointer_field is as follows.

The pointer_field contains the number of bytes in this packet that immediately follow the pointer_field that the CM decoder must skip past before looking for the beginning of a Euro-DOCSIS MAC frame. A pointer field MUST be present if it is possible to begin a data-over-cable MAC Frame in the packet, and MUST point to either:

1) the beginning of the first MAC frame to start in the packet; or
2) any stuff_byte preceding the MAC frame.

F.7.5 Interaction with the MAC sublayer

Refer to clause 7.5.

F.7.6 Interaction with the physical layer

The MPEG-2 packet stream MUST be encoded according to [b-EN 300 429].

F.7.7 MPEG header synchronization and recovery

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.

The format of MAC frames is described in detail in clause 8.

F.8 Media access control specification

F.8.1 Introduction

Refer to clause 8.1.

F.8.2 MAC frame formats

Refer to clause 8.2.

F.8.3 MAC management messages

F.8.3.1 MAC management message header

Refer to clause 8.3.1.

F.8.3.2 Time synchronization (SYNC)

Time synchronization (SYNC) MUST be transmitted by CMTS at a periodic interval to establish MAC sublayer timing. The message MUST use an FC field with FC_TYPE = MAC-specific header and FC_PARM = timing MAC header. This MUST be followed by a frame PDU in the format shown in Figure F.1.

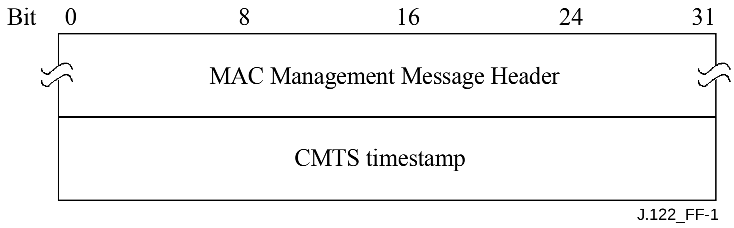

Figure F.1 – Format of packet PDU following the timing header

The parameters shall be as defined below.

CMTS timestamp: The count state of an incrementing 32-bit binary counter clocked with the CMTS 10.24-MHz master clock.

The CMTS timestamp represents the count state at the instant that the first byte (or a fixed time offset from the first byte) of the time synchronization MAC management message is transferred from the downstream transmission convergence sublayer to the downstream physical media dependent sublayer as described in clause 6.3.7. A CMTS MUST always put the SYNC-message at the start of an MPEG-packet. This is required for compatibility with certain CM implementations.

Annex G

DOCS 2.0 and 1.0/1.1 interoperability

(This annex forms an integral part of this Recommendation)

DOCS 2.0 is the third generation of the DOCS Recommendation. The terms DOCS 2.0, DOCS 1.1, and DOCS 1.0 refer to these three different Recommendations.

The DOCS 2.0 Recommendation primarily aims at enhancing the limited upstream physical layer performance of a DOCS 1.0- or 1.1-based cable access system. Two new MAC management message types have been defined, and several new parameter encodings have been defined in the existing MAC messages. A DOCS 2.0 CMTS is capable of supporting a higher upstream throughput for a given channel bandwidth as well as increased tolerance to noise experienced in the upstream.

As well as supporting DOCS 2.0-capable CMs, the DOCS 2.0 CMTS must be backwards compatible with DOCS 1.0 and DOCS 1.1 CMs. Furthermore, it is necessary for a DOCS 2.0 CM to function like a 1.0 CM when interoperating with a 1.0 CMTS and to function like a 1.1 CM when interoperating with a 1.1 CMTS.

This annex describes the interoperability issues and trade-offs involved when the operator wishes to support DOCS 1.0 and/or DOCS 1.1 CMs as well as DOCS 2.0 CMs on the same cable access channel.

G.1 General interoperability issues

This clause addresses the general DOCS 1.x/2.0 interoperability issues that do not depend on the modulation type used for the upstream channel.

G.1.1 Provisioning

The parameters of the TFTP configuration file for a DOCS 2.0 CM are (except for the addition of one optional TLV) identical to those for a DOCS 1.1 CM, and are a superset of those for a DOCS 1.0 CM. Configuration-file editors that support DOCS 1.1 may need to be modified to support the new TLV defined in DOCS 2.0.

A TFTP configuration file containing class of service TLVs is considered a "DOCS 1.0 style" configuration file. A TFTP configuration file containing service flow TLVs is considered a "DOCS 1.1/2.0 style" configuration file. A TFTP configuration file containing both class of service and service flow TLVs will be rejected by the CMTS (see clause 11.2.9).

If a DOCS 2.0 CM is provisioned with a DOCS 1.0-style TFTP configuration file, it will register as specified in clause G.1.2, although in the REG-REQ it MUST still specify "DOCS 2.0" in the DOCS version modem capability and MAY specify additional advanced (i.e., DOCS 1.1) modem capabilities that it supports. Thus, a DOCS 2.0 CM can be provisioned to work seamlessly on either a DOCS 1.0, a DOCS 1.1 or a DOCS 2.0 network. However, a DOCS 2.0 modem on a DOCS 1.x network would be clearly unable to support any DOCS 2.0-specific features.

If a DOCSIS 2.0 CM operating on an S-CDMA channel with the maximum scheduled codes feature enabled (see clause 8.3.3), and provisioned with a DOCSIS 1.0-style configuration file, SHOULD support fragmentation and indicate that support in the modem capabilities encoding in the REG-REQ message. If a DOCS 2.0 CM supports certain advanced capabilities when registered as a DOCS 1.0 CM (as indicated by the modem capabilities encoding), those features MUST function according to the requirements defined in the DOCS 2.0 Recommendations.

On the other hand, DOCS 1.0 CMs do not recognize (and ignore) many of the new TLVs in a DOCS 1.1/2.0 style configuration file, and will be unable to register successfully if provisioned with a DOCS 1.1/2.0 configuration file. To prevent any functionality mismatches, a DOCS 2.0 CMTS MUST reject any registration request with DOCS 1.1/2.0-specific configuration parameters that are not supported by the associated modem capabilities encoding in the REG-REQ (see clause C.1.3.1).

G.1.2 Registration

A DOCS 2.0 CMTS is designed to handle the registration TLVs from DOCS 1.0 CMs as well as the TLVs from DOCS 1.1 (TLV types 22 to 38) or DOCS 2.0 (TLV types 22 to 39) CMs. Furthermore, a DOCS 2.0 CM can handle any TLV in a configuration file usable by a DOCS 1.0 CM.

There is a slight difference in the registration-related messaging procedure when the DOCS 2.0 CMTS is responding to a DOCS 1.1 or 2.0 CM as opposed to a DOCS 1.0 CM (or a DOCS 1.1 CM using a 1.0-style configuration file). There is a further difference in the way a DOCS 2.0 CMTS handles registration from a DOCS 2.0 CM using a DOCS 1.0-style configuration file, depending on whether the upstream on which the registration is occurring has DOCS 2.0 features.

A DOCS 1.1 or 2.0 CM could be configured to use the service class name, which is statically defined at the CMTS, instead of explicitly asking for the service class parameters. When the DOCS 2.0 CMTS receives such a registration request, it encodes the actual parameters of that service class in the registration response and expects the registration acknowledge MAC message from the CM. If the detailed capabilities in the registration response message exceed those the CM is capable of supporting, the CM is required to indicate this to the CMTS in its registration acknowledge.

When a DOCS 1.0 CM (or a 1.1 CM using a DOCS 1.0-style configuration file) registers with the same CMTS, the absence of service class names eliminates the need for the DOCS 2.0 CMTS to explicitly specify the service class parameters in the registration response using DOCS 1.1 or 2.0 TLVs. The registration request from a DOCS 1.0 CM explicitly requests all non-default service class parameters in the registration request per its provisioning information. When a DOCS 2.0 CMTS receives a registration request containing DOCS 1.0 class of service encodings, it will respond with the DOCS 1.0-style registration response and, if the CM is a DOCS 1.x CM, not expect the CM to send the registration acknowledge MAC message. A DOCS 1.0 CM can be further identified by the absence of the "DOCS version" modem capabilities encoding in the registration request.

In the case where a DOCS 2.0 CM is using a DOCS 1.0-style configuration file, there is an additional consideration. This is because in the case where the upstream is a type 2 upstream (see clause 11.2.2), and therefore supports both TDMA and A-TDMA features, the registration acknowledge message is also used to synchronize switching from TDMA (DOCS 1.x) operation to A-TDMA (DOCS 2.0) operation. It is important that this switch be coordinated correctly between the CM and the CMTS in order for the CMTS to be able to correctly interpret bandwidth requests from the CM (see clause 11.2.9). Therefore, when a DOCS 2.0 CM registers using a DOCS 1.0-style configuration file on a type 2 or type 3 upstream, it transmits a registration acknowledgment with a confirmation code of OK/SUCCESS (since DOCS 1.0-style registration does not allow for the CM to reject the registration response). The CMTS knows to expect this because the modem capabilities field in the registration request indicated that the CM was a DOCS 2.0 CM. Table G.1 summarizes registration behaviour for all cases involving a DOCS 2.0 CM.

Table G.1 – Registration behaviour for a DOCS 2.0 CM

Configuration file	DOCS 1.0 CMTS	DOCS 2.0 CMTS with a type 1 upstream or a DOCS 1.1 CMTS	DOCS 2.0 CMTS with a type 2 or type 3 upstream
DOCS 1.1/2.0-style configuration file that does not disable DOCS 2.0 mode	N/A	CM sends DOCS 1.1/2.0-style REG-REQ. CMTS sends DOCS 1.1/2.0-style REG-RSP and CM responds with REG-ACK.	CM sends DOCS 1.1/2.0-style REG-REQ. CMTS sends DOCS 1.1/2.0-style REG-RSP and CM responds with REG-ACK.
DOCS 1.0-style configuration file that does not disable DOCS 2.0 mode	CM sends DOCS 1.0-style REG-REQ. CMTS sends DOCS 1.0-style REG-RSP. CM MUST NOT send REG-ACK.	CM sends DOCS 1.0-style REG-REQ. CMTS sends DOCS 1.0-style REG-RSP. CM MUST NOT send REG-ACK.	CM sends DOCS 1.0-style REG-REQ. CMTS sends DOCS 1.0-style REG-RSP. CM MUST send REG-ACK with SUCCESS confirmation code. CMTS MUST wait for REG-ACK.
DOCS 1.1/2.0-style configuration file that disables DOCS 2.0 mode	N/A	CM sends DOCS 1.1/2.0-style REG-REQ. CMTS sends DOCS 1.1/2.0-style REG-RSP and CM responds with REG-ACK.	CM sends DOCS 1.1/2.0-style REG-REQ. CMTS sends DOCS 1.1/2.0-style REG-RSP and CM responds with REG-ACK.
DOCS 1.0-style configuration file that disables DOCS 2.0 mode	CM sends DOCS 1.0-style REG-REQ. CMTS sends DOCS 1.0-style REG-RSP. CM MUST NOT send REG-ACK.	CM sends DOCS 1.0-style REG-REQ. CMTS sends DOCS 1.0-style REG-RSP. CM MUST NOT send REG-ACK.	CM sends DOCS 1.0-style REG-REQ. CMTS sends DOCS 1.0-style REG-RSP. CM MUST NOT send REG-ACK.

Another minor issue is that a DOCS 1.0 CM will request a bidirectional (with upstream/downstream parameters) service class from the CMTS using a class of service configuration setting.

Since a DOCS 2.0 CMTS typically operates with unidirectional service classes, it can easily translate a DOCS 1.0 class of service configuration setting into DOCS 1.1 or 2.0 service flow encodings for setting up unidirectional service classes in local QoS implementation. However, for DOCS 1.0 modems, the DOCS 2.0 CMTS MUST continue to maintain the QosProfile table (with bidirectional class parameters) for backward compatibility with the DOCS 1.0 MIB.

Thus, if properly provisioned, a DOCS 1.0, a DOCS 1.1 and a DOCS 2.0 CM can all successfully register with the same DOCS 2.0 CMTS, and a DOCS 2.0 CM can register with a 1.0 CMTS. Furthermore, a DOCS 2.0 CM can use a DOCS 1.0-style configuration file, register on a DOCS 2.0 CMTS and still use DOCS 2.0 enhanced physical-layer features with DOCS 1.0 class of service features.

G.1.3 Dynamic service establishment

There are 8 MAC messages that relate to dynamic service establishment. A DOCS 1.0 CM will never send them to any CMTS since they are unsupported. A DOCS 1.1 or 2.0 CM will never send them to a DOCS 1.0 CMTS because:

a) to register successfully it has to be provisioned as a DOCS 1.0 CM; and

b) when provisioned as a DOCS 1.0 CM, it acts identically.

When a DOCS 1.1 or 2.0 CM is connected to a DOCS 1.1 or 2.0 CMTS, these messages work as expected.

G.1.4 Fragmentation

Fragmentation is initiated by the CMTS. Thus, a DOCS 1.0 CMTS will never initiate fragmentation since it knows nothing about it. A DOCS 1.1 or 2.0 CMTS can only initiate fragmentation for DOCS 1.1 or 2.0 CMs. A DOCS 1.1 or 2.0 CMTS MUST NOT attempt to fragment transmissions from a DOCS 1.0 CM that has not indicated a modem capabilities encoding for fragmentation support with a value of 1.

G.1.5 Multicast support

It is mandatory for DOCS 1.0 CMs to support forwarding of multicast traffic. However, the Recommendation is silent on IGMP support. The only standard mechanism for controlling IP-multicast on DOCS 1.0 CMs is through SNMP and packet filters. Designers of DOCS 1.0 networks will have to deal with these limitations and expect no different from DOCS 1.0 CMs on a DOCS 2.0 network.

DOCSIS 2.0 CMs in 1.0 mode MUST still comply with the requirements for IGMP and the forwarding of multicast traffic as per clauses 5.1.2.3.2 and 5.3.1.

G.1.6 Changing upstream channels

A DOCS 2.0 CMTS is capable of specifying the level of re-ranging to be performed only when it issues a DCC-Request to the CM. This re-ranging technique parameter is specified by the DOCS 2.0 CMTS using a TLV in the DCC-Request MAC message.

DOCS 1.1 or 2.0 CMs can benefit by only re-ranging to the level specified by this TLV. This can help in reducing the re-initialization time following a DCC for the DOCS 1.1 or 2.0 CM carrying a voice call. A DOCS 2.0 CMTS is aware of the type of CM to which it is issuing the channel change request. It MUST refrain from sending a DCC-Request for DOCS 1.0 CMs, instead choosing to send a UCC-Request. If a DOCS 2.0 CMTS sends the UCC-Request, the DOCS 1.0 CMs will perform the default DOCS 1.0 re-ranging from start (initial ranging) with its existing non-zero primary SID.

G.2 Hybrid devices

Some DOCS 1.0 CM designs may be capable of supporting individual DOCS 1.1 features via a software upgrade. Similarly, some DOCS 1.0 CMTSs may be capable of supporting individual DOCS 1.1 features. To facilitate these "hybrid" devices, the majority of DOCS 1.1 features are individually enumerated in the modem capabilities.

DOCS 1.0 hybrid CMs MAY request DOCS 1.1 features via this mechanism. However, unless a CM is fully DOCS 1.1-compliant (i.e., not a hybrid), it MUST NOT send a "DOCS version" modem capability that indicates DOCS 1.1. Similarly, unless a CM is fully DOCS 2.0-compliant, it MUST NOT send a "DOCS version" modem capability which indicates DOCS 2.0.

If a hybrid CM intends to request such DOCS 1.1 capabilities from the CMTS during registration, it MUST send the ASCII coded string in option code 60 of its DHCP request, "docsis1.0:xxxxxxx", where xxxxxxx MUST be an ASCII representation of the hexadecimal encoding of the modem capabilities. Refer to clauses C.1.3.1 and D.1.1 for details. The DHCP server MAY use such information to determine what configuration file the CM is to use.

In order to control the hybrid operation of modems, if a DOCS 2.0 CMTS receives a DOCS 1.0-style registration request message from a CM, the CMTS MUST, by default, force the modem to operate in a "pure" DOCS 1.0 mode with respect to certain features by disabling those features via the modem capabilities encoding in the registration response. Specifically, the CMTS MUST support the six default values given in square brackets in Table G.2. The CMTS MAY provide switches, as indicated in Table G.2, for the operator to selectively allow certain hybrid features to be enabled. As an exception to these defaults, the DOCSIS 2.0 CMTS SHOULD allow the use of fragmentation for DOCSIS 2.0 CMs registering in DOCSIS 1.0 mode on an S-CDMA channel that has the maximum scheduled codes feature (see clause 8.3.3) enabled.

Table G.2 – Hybrid mode controls

	Concatenation Support	Fragmentation Support	Privacy Support
DOCS 1.0 CM	Allow/[deny]	Allow/[deny]	Allow BPI+/[force BPI]
DOCS 1.1 or 2.0 CM in DOCS 1.0 mode	Allow/[deny]	Allow/[deny]	Allow BPI+/[force-BPI]

Normally, a DOCS 1.0 CMTS sets all unknown modem capabilities to "off" in the registration response indicating that these features are unsupported and MUST NOT be used by the CM. A DOCS 1.0 hybrid CMTS MAY leave supported modem capabilities set to "on" in the registration response. However, unless a CMTS is fully DOCS 1.1- or 2.0-compliant (i.e., not a hybrid), it MUST still set all "DOCS version" modem capabilities to DOCS 1.0.

As always, any modem capability set to "off" in the registration response must be viewed as unsupported by the CMTS and MUST NOT be used by the CM.

G.3 DOCS 2.0 TDMA interoperability

G.3.1 Mixed-mode operation with TDMA

In mixed-mode operation with both DOCS 1.x and DOCS 2.0 TDMA, a single channel is defined with a single UCD that contains both type 4 and type 5 burst descriptors. DOCS 1.x and 2.0 modems use the type 4 burst descriptors; DOCS 2.0 modems MUST also use the type 5 burst descriptors. DOCS 2.0 modems will use IUCs 9 and 10.

The following rules of operation apply:

1) Prior to and during registration, a DOCS 2.0 TDMA-capable modem operating on a channel of type 1 or 2 (refer to clause 11.2.2) MUST calculate its request size based on DOCS 1.x IUC parameters, and the CMTS MUST make all grants using DOCS 1.x IUCs.
2) On a type 2 channel, a DOCS 2.0 TDMA CM MUST switch to DOCS 2.0 TDMA mode after transmission of the registration acknowledgement (REG-ACK) message. If the CM receives a registration response (REG-RSP) message after transmission of the REG-ACK message, the CM MUST switch back to DOCS 1.1 mode before it continues with the registration process (see Figure 11-12).
3) A CM in DOCS 2.0 TDMA mode MUST calculate its request size based on IUC types 9 and 10. The CMTS MUST make grants of IUC types 9 and 10 to that CM after it receives the registration acknowledgement message from the CM (see clause 11.2.9).
4) On a type 2 channel, the CM MUST ignore grants with IUCs that are in conflict with its operational mode (e.g., the CM receives a grant with IUC 5 when it is in DOCS 2.0 TDMA mode).
5) On a type 3 channel, the CMTS MUST use type 5 burst descriptors in order to prevent DOCS 1.x modems from attempting to use the channel. All data grants are in IUC types 9 and 10.
6) On a type 2 channel, only advanced PHY short (IUC 9) and advanced PHY long (IUC 10) bursts may be classified as burst descriptor type 5.

7) A DOCS 1.x modem that does not find appropriate type 4 burst descriptors for long or short data grant intervals MUST consider the UCD, and the associated upstream channel, unusable.

G.3.2 Interoperability and performance

This clause addresses the issue of performance impact on the upstream channel when DOCS 1.x CMs are provisioned to share the same upstream MAC channel as DOCS 2.0 TDMA CMs.

Since the initial maintenance, station maintenance, request and request/data IUCs are common to both DOCS 2.0 TDMA and DOCS 1.x CMs, the overall channel will experience reduced performance compared to a dedicated DOCS 2.0 TDMA upstream channel. This is due to broadcast/contention regions not being capable of taking advantage of improved physical layer parameters.

G.4 DOCS 2.0 S-CDMA interoperability

G.4.1 Mixed-mode operation with S-CDMA

In mixed-mode operation with both TDMA and S-CDMA, two logically separate upstream channels are allocated by the CMTS, one for TDMA modems, and another for DOCS 2.0 modems operating in S-CDMA mode. Each channel has its own upstream channel ID, and its own UCD. However, these two channels are both allocated the same RF centre frequency on the same cable plant segment. The CMTS controls allocation to these two channels in such a way that the channel is shared between the two groups of modems. This can be accomplished by reserving bandwidth through the scheduling of data grants to the NULL SID on all channels other than the channel that is to contain the potential transmit opportunity. Using this method, an upstream channel can support a mixture of differing physical layer DOCS modems, with each type capitalizing on their individual strengths. The channel appears as a single physical channel that provides transmission opportunities for both DOCS 1.x and 2.0 modems. The mixed-mode configuration of the channel will be transparent to the CMs.

The following rule of operation applies:

- The CMTS MUST use only type 5 burst descriptors on the S-CDMA channel in order to prevent DOCS 1.x modems from attempting to use the channel.

G.4.2 Interoperability and performance

This clause addresses the issue of performance impact on the S-CDMA upstream channel when the upstream centre frequency is shared with an upstream TDMA channel.

Due to the lack of ability to share the upstream transmit opportunities, the channels will not experience the statistical multiplexing benefits during contention regions across the CMs. Dedicated initial maintenance regions will be required on both logical MAC channels slightly reducing the overall performance available. Request and request/data regions will also not be capable of being shared although an intelligent CMTS scheduler will be able to reduce most performance impact.

Annex H

The DOCS MAC/PHY interface (DMPI)

(This annex forms an integral part of this Recommendation)

H.1 Scope

Integrated circuit (IC) chip sets with separate MAC and PHY chips used in the implementation of a CMTS SHOULD implement DMPI. DMPI does not apply to IC chip sets which integrate MAC and PHY components together into one chip.

Any usage of "MUST", "SHOULD", or "MAY" within the DMPI specification applies only if DMPI is implemented.

H.2 Conventions

H.2.1 Terminology

Throughout this annex, the terms "MAC" and "PHY" are used extensively. MAC is used to refer to the device which provides the interface between the PHY devices and the system. The term "PHY" refers to the device which performs the physical layer processing for a single RF channel. It is important to note that both of these terms refer to physical devices as opposed to layers in the IP protocol stack. For the purposes of this Recommendation, integrated circuit chips that handle multiple RF channels simultaneously are considered to contain multiple PHY devices.

H.2.2 Ordering of bits and bytes

The following rules control the order of transmission of bits and bytes over all the interfaces specified in this Recommendation. In all cases, fields of data blocks are transmitted in the order in which they appear in the data block format description.

- Multibyte quantities are transmitted most significant byte first (big-endian byte ordering). This byte ordering applies regardless of the width of the interface (byte, nibble, single bit).
- On nibble-wide interfaces, the most significant nibble (bits 7:4) is transmitted first.
- On bit-wide interfaces, the most significant bit of each field is transmitted first.

H.2.3 Signal naming conventions

Signal names that end with an "_N" are active low. Signals without this suffix are active high.

H.2.4 Active clock edge

All signals are driven and sampled on the rising edge of the clock except where otherwise noted.

H.2.5 Timing specifications

The timing specifications for DMPI use the following terminology (see Table H.1):

Table H.1 – Timing parameters

Parameter	Symbol	Description
Clock frequency	f	The frequency of the interface clock
Clock low pulse width	t_{lpw}	The low time of the interface clock
Clock high pulse width	t_{hpw}	The high time of the interface clock
Clock rise/fall time	t_{rf}	The transition time of the clock
Input setup time to clock	t_{su}	From when an interface signal is valid to the following rising clock edge
Input hold time from clock	t_h	From the rising clock edge to when an interface signal becomes invalid
Clock to signal valid delay	t_{cq}	From the rising edge of the interface clock to an interface signal becoming valid

Following are some usage notes for these timing parameters:

- Set-up and hold time specifications are given from the point of view of the DMPI interface and not from the point of view of a device on the DMPI interface. The clock to output, on the other hand, specifies the timing requirement of a DMPI device.
- The t_{su} parameter specifies the minimum guaranteed amount of set-up time provided by the DMPI interface measured at the receiving device. Therefore, inputs on DMPI devices should require no more than this amount of set-up time.
- The t_h parameter specifies the minimum guaranteed amount of hold time provided by the DMPI interface measured at the receiving device. Therefore, inputs on DMPI devices should require no more than this amount of hold time.
- The t_{cq} parameter specifies the minimum and maximum clock to output time at the driving device. The purpose of the minimum specification is to allow for clock skew between the driving and receiving DMPI device. For example, a 1-ns minimum specification and a 0-ns DMPI hold time requirement allows for at most 1-ns of clock skew between devices. The maximum specification is to allow for the settling time of signals from the driving device to the receiving device and clock skew between devices.

H.3 Overview

This annex describes the DOCS MAC/PHY interface (DMPI). DMPI is used to connect a DOCS MAC device to DOCS downstream and upstream PHY devices. While DMPI is a single interface, for the purposes of clarity, DMPI signals have been grouped into four separate groups. Each group serves a specific purpose and is independent of the others. For this reason, each group of signals is also referred to as an interface.

A downstream PHY MUST include a downstream data interface and an SPI bus interface. An upstream PHY must include an upstream data interface, an upstream control interface, and an SPI bus interface. PHY chips that integrate multiple PHYs into a single package MUST have one set of interfaces for each PHY that has been integrated, with the following exception:

An integrated PHY device MAY use a single select and a single SPI bus for all internal PHYs (using the SPI bus protocol described in clause H.8.4). An integrated upstream PHY device MAY have only one TS_CLK input and only one US_CLK input.

A MAC MUST include one downstream data interface for each downstream PHY it supports and one set of upstream interfaces (upstream data and upstream control) for each upstream PHY it supports. It MUST include at least one SPI bus interface.

DMPI has been defined with the following goals in mind:

- vendor independence;
- flexibility for future growth and vendor differentiation;
- minimization of PHY-specific logic in the MAC.

Figure H.1 shows an example application of DMPI. Note that this figure shows the connections required for a single DS PHY and a single US PHY. Obviously, other applications with multiple DS and US PHYs are possible.

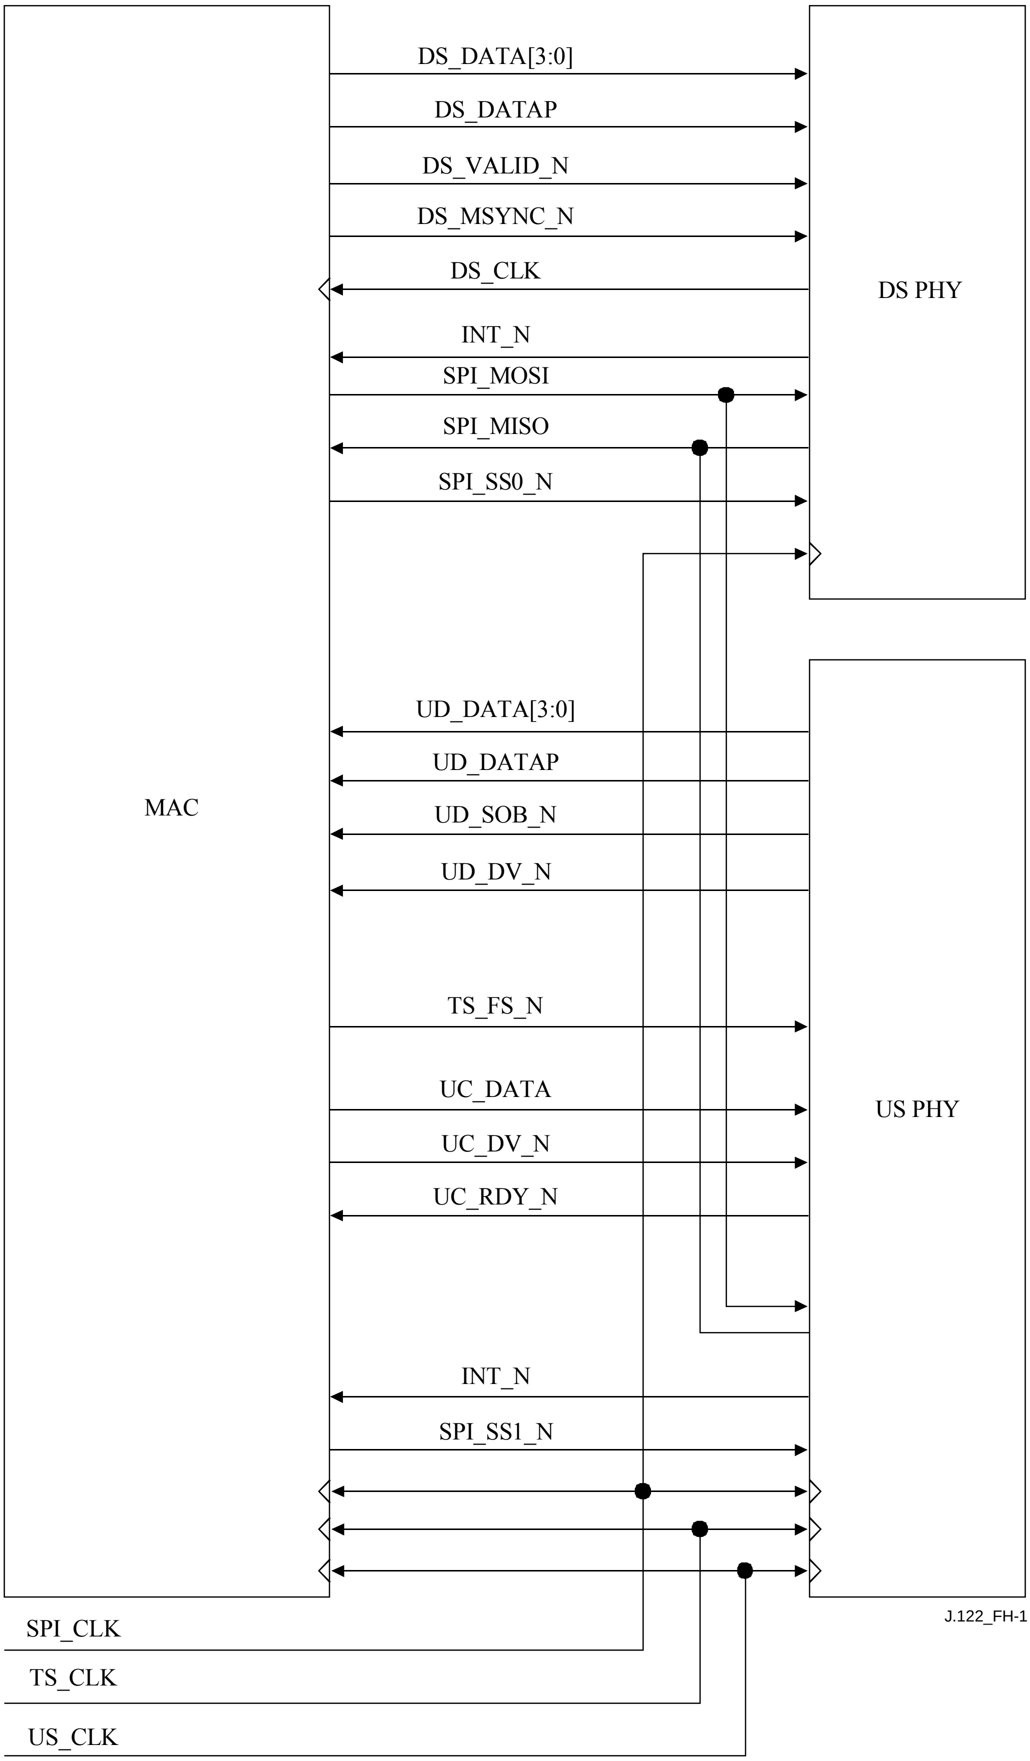

Figure H.1 – DMPI application

H.3.1 Downstream data

The downstream data interface carries data from the MAC to the PHY for transmission on the downstream. All signals on the interface are synchronous with respect to a clock driven by the PHY and received by the MAC. Four bits of data are transferred on each clock. The frequency of this clock is proportional to the downstream bit rate. Its precise frequency is a function of the downstream symbol rate, the modulation type (64QAM or 256QAM), and the physical layer framing in use (see Annex A or B of [ITU-T J.83]).

H.3.2 Upstream data

The upstream data interface carries data from the PHY to the MAC which has been received on the upstream. The interface is synchronous to a dedicated interface clock whose frequency is not directly related to the upstream bit rate.

Data is transferred over the interface using a mixture of TLVs and TVs (a TLV for which the length is implied by the type). Along with the DOCS burst data, certain status information about the burst is also transferred to the MAC. There is also a TLV which allows the PHY to indicate that it did not receive a burst when one was expected.

H.3.3 Upstream control

The upstream control interface is used for two purposes. The first is to initialize the PHY's timestamp counter, frame counter and mini-slot counter, and to check that the PHY's timestamp counter remains synchronized to the MAC's during operation. The second is to allow the MAC to pass information to the PHY regarding upcoming bursts.

This interface uses two clocks. The clock used for the counter synchronization is the 10.24 MHz CMTS master clock. A single signal, that is synchronized to this clock, is used to perform this counter synchronization. The other clock used for this interface is shared with the upstream data interface and has a frequency unrelated to the upstream modulation clock or the 10.24 MHz CMTS master clock. This clock, along with an associated set of signals, is used to transfer descriptions of future bursts.

H.3.4 SPI bus

The serial peripheral interconnect (SPI) bus is used to read and write registers in the PHYs. The system MAY use one or more SPI buses to provide register access to the PHYs. The number of SPI buses in the system is a function of the system's SPI bus performance requirements. Each SPI bus has a single master device which MAY be the MAC. Alternatively, an SPI bus master MAY be some other device in the system (e.g., a microprocessor). References to the SPI bus in this recommendation assume that the MAC is the master. The PHYs MUST only be slave devices. Each PHY MUST have one SPI bus interface. Multiple PHYs MAY share the same SPI bus.

The SPI bus definition includes an interrupt signal (INT_N). Each PHY MUST drive an interrupt. The interrupt signals MAY be received by an SPI bus master or they MAY be received by some other device in the system that provides the ability to monitor their state.

H.4 Signals

H.4.1 Downstream data

The signals used for the downstream data interface are defined in Table H.2.

Table H.2 – Downstream data interface signals

Signal	Description
DS_CLK	DS transmit clock Driven by the downstream PHY See clauses H.5.1 and H.5.2 for detailed requirements for this clock
DS_MSYNC_N	Downstream MPEG sync Driven by the MAC Marks first nibble of sync byte; active low
DS_VALID_N	Downstream data valid Driven by the MAC Indicates that valid data is present on DS_DATA
DS_DATA[3:0]	DS transmit data Driven by the MAC
DS_DATAP	Downstream parity Driven by the MAC Even parity for DS-DATA (the number of 1s across DS_DATA and DS_DATAP is even) DS_DATA and its corresponding DS_DATAP are driven on the same clock. Parity is not delayed a clock cycle as it is in some interfaces.

H.4.2 Upstream data

The signals used for the upstream data interface are defined in Table H.3.

Table H.3 – Upstream data interface signals

Signal	Description
US_CLK	Upstream data/control clock Driven by external clock source (input to MAC and PHY)
UD_SOB_N	Upstream data start of data block Driven by upstream PHY asserted when the first nibble or first byte of the data block is on UD_DATA
UD_DV_N	Upstream data valid Driven by upstream PHY indicates valid data on UD_DATA
UD_DATA[3:0]	Upstream data Driven by upstream PHY
UD_DATAP	Upstream data parity Driven by upstream PHY Even parity for UD_DATA (the number of 1s across UD_DATA and UD_DATAP is even) UD_DATA and its corresponding UD_DATAP are driven on the same clock. Parity is not delayed a clock cycle as it is in some other interfaces.

H.4.3 Upstream control

Table H.4 lists the signals that are used for the upstream control interface.

Table H.4 – Upstream control interface signals

Signal	Description
US_CLK	Upstream clock Driven by external clock source (input to MAC and PHY)
UC_DV_N	Upstream control data valid Driven by the MAC Indicates valid upstream control message data on UC_DATA
UC_DATA	Upstream control data Driven by the MAC
UC_RDY_N	Upstream control ready Driven by the PHY Indicates that the PHY is ready to receive an interval description message
TS_CLK	10.24-MHz master clock Driven by external clock source (input to MAC and PHY)
TS_FS_N	Timestamp frame sync Driven by the MAC

H.4.4 SPI bus

Table H.5 – SPI bus signals

Signal name	Description
SPI_CLK	SPI bus clock Driven by a source external to the MAC and PHY or driven by the MAC
SPI_MOSI	Master out/slave in Serial data from the MAC to the PHY
SPI_MISO	Slave out/master in Serial data from the PHY to the MAC MAY be driven by the PHY from the falling edge of SPI_CLK
SPI_SSx_N	Slave select Selects a slave for a transaction. One slave select signal is provided by the MAC for each PHY (x = 1 to N). Addressing of devices within a package is provided by the protocol layer described in clause H.8.4. MAY be sampled by the PHY on the falling edge of SPI_CLK
INT_N	Interrupt Driven by PHYs Open drain

H.4.5 Parity

The downstream data, upstream data, and upstream control interfaces use parity to maintain data integrity on the interface. Parity SHOULD be implemented.

The SPI bus does not have parity.

Parity is even and covers only the data lines of the interface. Specific rules for parity checking are detailed in the following clauses.

H.4.5.1 Downstream data

Parity must be checked by the downstream PHY and covers DS_DATA. Since the downstream transmit data is protected (DOCS frame HCS and CRC), detection of a parity error is not considered fatal and MUST NOT cause the processing of transmit data to halt. The PHY must generate an interrupt to the system when it detects a parity error so that the system can be made aware of its occurrence. Parity checking on this interface provides a way to distinguish between data errors on the interface and those in other parts of the data path.

H.4.5.2 Upstream data

Parity is checked by the MAC and covers UD_DATA. Since the upstream receive data is protected (DOCS frame HCS and CRC), detection of a parity error is not considered fatal and MUST NOT cause the processing of receive data to halt. The MAC must generate an interrupt to the system when it detects a parity error so that the system can be made aware of its occurrence. Parity checking on this interface provides a way to distinguish between data errors on the interface and those in other parts of the data path.

H.4.5.3 Upstream control

Parity is checked by the upstream PHY and covers the entire upstream control message. A parity error on this interface is considered a fatal error. The PHY MUST NOT process the upstream control message which was received with a parity error as well as any subsequently received message. The PHY MAY process any upstream control messages received prior to the occurrence of the parity error. This processing MAY include the passage of various types of upstream data blocks to the MAC.

H.4.6 Interrupts

Various places in this Recommendation make reference to the assertion of an interrupt by the PHY. The characteristics of this interrupt MUST be as follows:

- one active low interrupt line of level type;
- driven open drain;
- cause of interrupt line assertion determined by software read(s) of PHY register(s) that contains one bit for each interrupt source;
- no hardware prioritization of interrupt sources;
- each interrupt source separately cleared by software write(s) to PHY register(s);
- asserted until all interrupt source bits are cleared (interrupt line is a simple OR of all interrupt sources).

H.5 Protocol

H.5.1 Downstream data (ITU-T Rec. J.83 Annex A)

Figure H.2 shows the protocol for ITU-T Rec. J.83 Annex A operation.

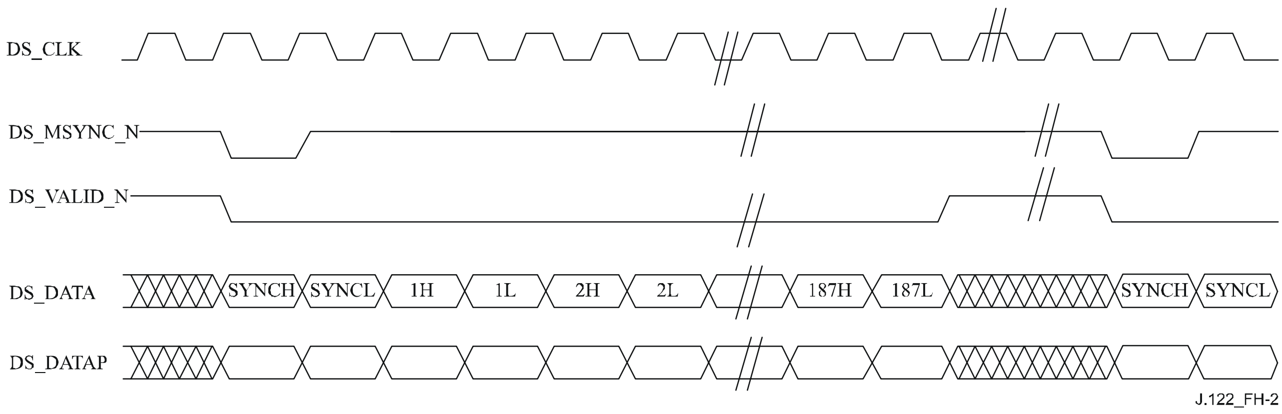

Figure H.2 – Downstream data J.83 Annex A protocol

The following behaviour of DS_CLK and DS_VALID_N is required:

- DS_CLK MUST NOT be gapped (it must have a constant frequency).
- DS_CLK frequency MUST be 1/4 of the downstream line rate. The DS line rate is the data rate including the ITU-T Rec. J.83 Annex A framing overhead.
- The MAC MUST assert DS_VALID_N for the entire 188-byte MPEG packet transfer and then MUST de-assert it for exactly 32 clocks following the transfer of the last nibble of the MPEG packet.

H.5.2 Downstream data (ITU-T Rec. J.83 Annex B)

Figure H.3 shows the protocol used to transfer data across this interface for ITU-T Rec. J.83 Annex B operation.

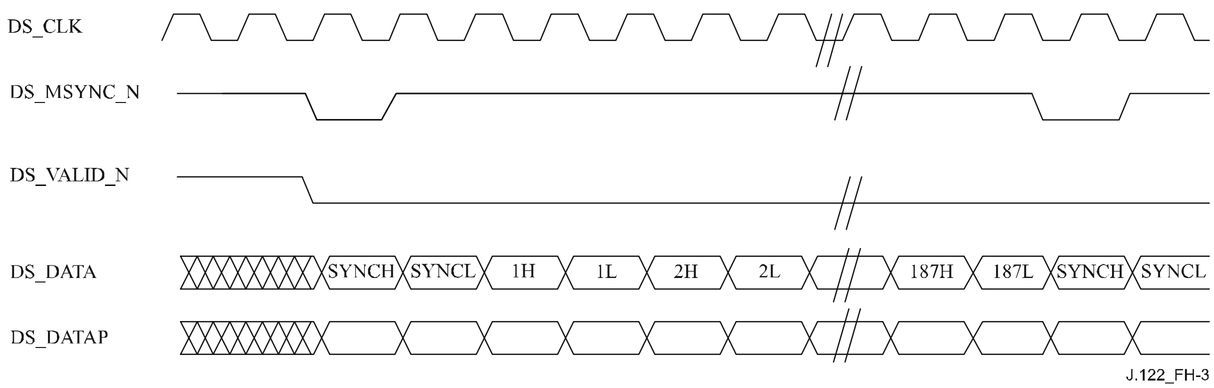

Figure H.3 – Downstream data signal protocol for J.83 Annex B operation

The following behaviour of DS_CLK and DS_VALID_N is required:

- DS_CLK MUST NOT be gapped (it must have a constant frequency).
- DS_CLK frequency MUST be 1/4 of the downstream payload rate. The downstream payload rate is the data rate excluding the ITU-T Rec. J.83 Annex B framing overhead.
- The MAC MUST keep DS_VALID_N always asserted.

H.5.3 Upstream data

Figure H.4 shows the signalling protocol for this interface.

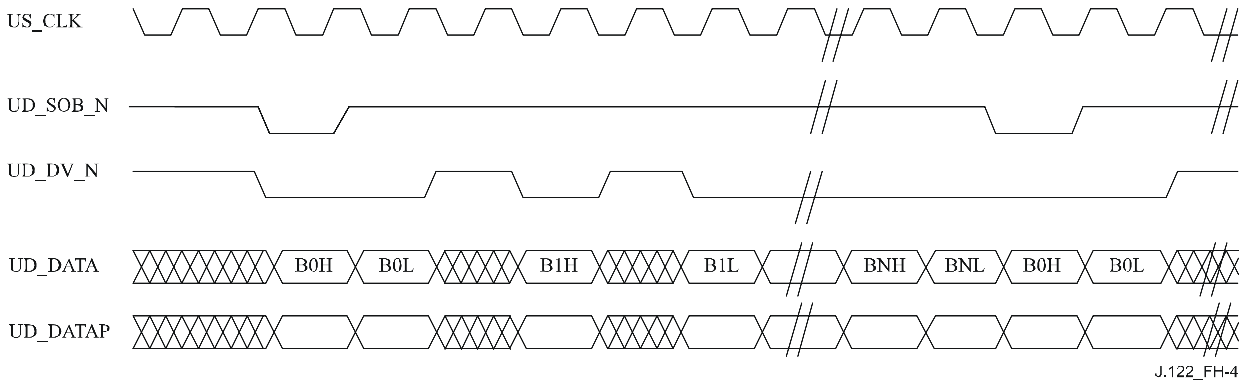

Figure H.4 – Upstream data protocol

It is a very simple protocol in which the upstream PHY indicates the presence of valid data on UD_DATA by asserting UD_DV_N. The MAC has no ability to control the flow of data and is required to sample UD_DATA on every rising clock edge on which UD_DV_N is asserted. The start of a data block is indicated by the PHY's assertion of UD_SOB_N. This signal MUST be asserted when the first nibble of the first byte of the data block is driven onto UD_DATA.

The MAC MUST keep track of length of each data block as it relates to the assertion of UD_SOB_N. If UD_SOB_N is asserted before the entire previous data block has been transferred, the MAC MUST drop the associated burst and generate an interrupt.

If the FIRST_STATUS byte indicates the absence of a PHY_STATUS data block but the PHY transfers one, the PHY_STATUS data block MUST be discarded by the MAC and an error MUST be signalled to the system.

H.5.4 Upstream control

H.5.4.1 Counter synchronization

The master timestamp counter MUST reside in the MAC. The master mini-slot counter and master frame counter MUST reside in the PHY. The PHY MUST capture a timestamp snapshot on every frame boundary. When the system needs a timestamp snapshot for a UCD, it MUST read this snapshot using a single SPI bus transaction. The PHY MUST ensure that the timestamp snapshot does not change during the SPI bus read transaction.

A common timestamp clock, TS_CLK, MUST be externally provided to the upstream PHYs and the MACs. The frequency of this timestamp clock MUST be 10.24 MHz ±5 ppm. The MAC MUST synchronize all PHYs to the timestamp value of the MAC. To accomplish this, the MAC MUST provide a frame sync pulse, TS_FS_N, to the PHYs that is synchronous to the positive edge of TS_CLK and has a pulse width equal to one period of TS_CLK.

The 32-bit timestamp counter consists of a group of upper bits and a group of lower bits. The MAC and PHY MUST provide at least the following choices of upper and lower bit boundaries shown in Table H.6.

Table H.6 – Timestamp counter initialization options

Upper bits	Lower bits	Frame sync interval
8	24	1638.4 ms
9	23	819.2 ms
10	22	409.6 ms
11	21	204.8 ms
12	20	102.4 ms

Figure H.5 shows an example of the proper assertion of the TS_FS_N signal. Note that the TIMESTAMP is shown for reference and is not part of the upstream control interface. In this example, upper bits = 8.

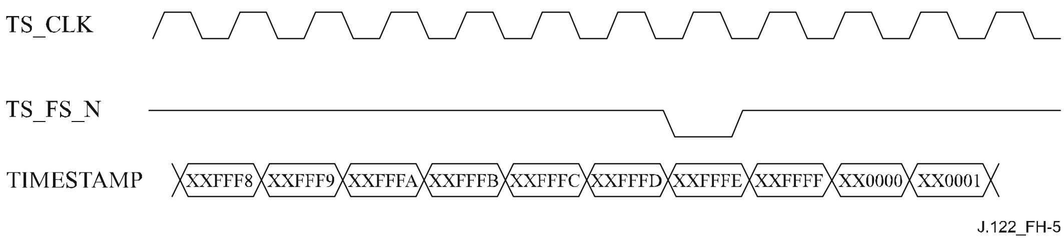

Figure H.5 – Counter synchronization

The MAC MUST assert TS_FS_N two 10.24-MHz clock periods prior to the lower bits of the MAC timestamp counter equalling all zeros. The MAC SHOULD provide some sort of maskable indication to the system when TS_FS_N occurs so that the system will have time to programme the registers of the PHYs prior to the next assertion of TS_FS_N. The period of TS_FS_N is a function of the timestamp bit time and the number of lower bits from Table H.6. The variation of the TS_FS_N period is to allow the system designer to trade off system response time versus the time available to initialize a PHY chip.

The PHY MUST provide all combinations of the following three initialization options when TS_FS_N is asserted:

- The upper bits of the timestamp counter are specified and the lower bits are set to zero.
- The full 8 bits of the frame counter are specified.
- The full 32 bits of the mini-slot counter are specified.

The specification of these counters is supplied across the SPI bus prior to the next frame sync pulse. Two TS_CLK clock cycles after TS_FS_N occurs, the PHY chip MUST initialize the specified counters. These counters are loaded at configuration time, and not on every assertion of TS_FS_N. A single PHY may be re-initialized without the need to re-initialize or otherwise interrupt the operation of other PHYs or the MAC.

During normal operation, the PHY MUST check that the lower bits of the PHY timestamp counter are exactly all zeros two 10.24-MHz clock cycles following every assertion of TS_FS_N. If the check is negative, the PHY MUST generate an interrupt and MUST provide status accessible over the SPI bus.

H.5.4.2 Upstream control messages

Figure H.6 shows a sample transaction.

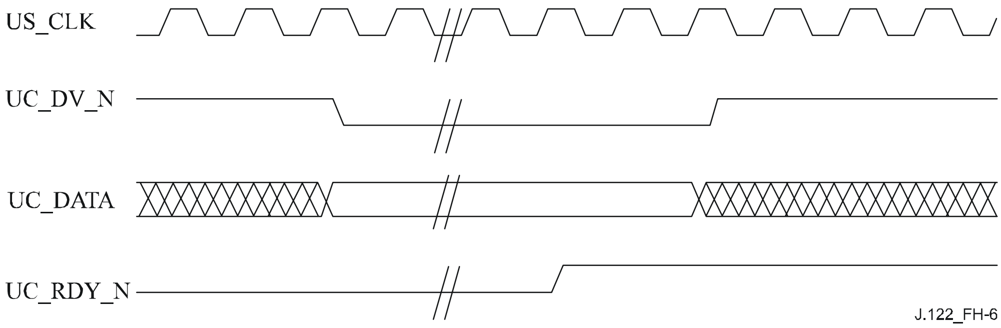

Figure H.6 – Upstream control message transfer

The upstream control interface is used to transfer time-critical configuration information (messages) to the PHY. The most common type of message is an interval description message. This message informs the PHY of the arrival time and characteristics of an upcoming burst. The protocol of this interface is very simple. Following is a description of how this interface works:

A transaction transfers a single upstream control message.

- UC_DV_N MUST remain asserted for the entire duration of the upstream control message transfer.
- The length of each upstream control message is inferred by its type.
- UC_DV_N MUST be de-asserted for a minimum of one US_CLK clock period to indicate the end of a transaction.
- UC_RDY_N MAY be used to stop and start the flow of interval description messages. UC_RDY_N does not affect the transfer of other message types. If the PHY is receiving an interval description and does not want to receive a subsequent interval description, the PHY MUST de-assert UC_RDY_N at least two clock cycles of US_CLK prior to the end of the current interval description. This de-assertion behaviour is shown in Figure H.6. The MAC MUST transfer a new interval description message within 10 US_CLK periods of the assertion of UC_RDY_N if a new interval description message is available.

H.5.5 SPI bus

Figure H.7 shows a SPI bus transaction.

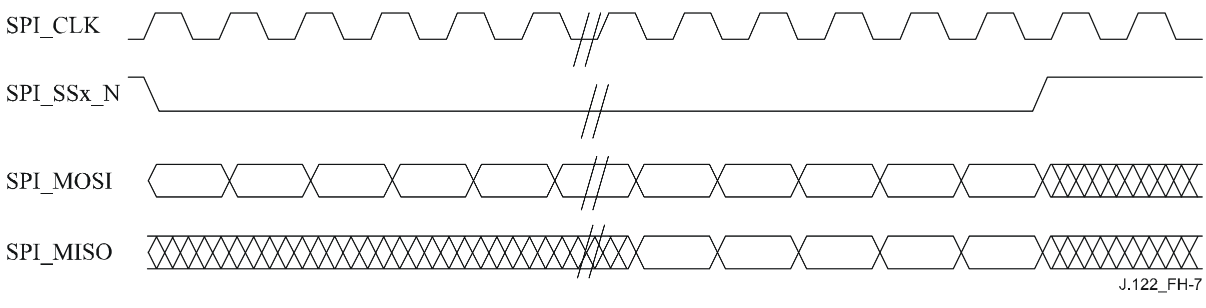

Figure H.7 – SPI bus transaction

A transaction proceeds as follows:

- The master asserts the select (SPI_SSx_N) of the desired slave device.
- The master drives SPI_MOSI with the appropriate command and data as described in clause H.8.4.

- For write commands, the first byte of data driven on SPI_MOSI is written to the register specified by the address in the command. The second byte of data (if it exists), is written to the next higher numbered address. Writes continue in this way until the master terminates the transaction by de-asserting SPI_SSx_N.
- For read commands, the slave drives the read data on SPI_MISO which is indicated by the address in the command. The first bit of this read data is driven one clock after the last bit of the command has been sampled. Read data from consecutively numbered addresses is driven until the master terminates the transaction by de-asserting SPI_SSx_N.

SPI_CLK MUST be driven (oscillate) for at least one clock period prior to the assertion of SPI_SSx_N, during the entire SPI bus transaction, and for one clock after the de-assertion of SPI_SSx_N. SPI_CLK MAY be driven high or low at all other times.

H.6 Electrical specifications

H.6.1 DC specifications

Devices which connect to DMPI must meet the requirements listed in Table H.7. Note that output high voltage and output high current specifications do not apply to the INT_N output as it is open drain.

Table H.7 – DC characteristics

Parameter	Symbol	Min	Max	Units	Comments
Input capacitance			10	pF	
Input low voltage	V_{il}		0.8	V	
Input high voltage	V_{ih}	2.0		V	
Output low voltage	V_{ol}		0.4	V	
Output high voltage	V_{oh}	2.4		V	
Output low current	l_{ol}	4		mA	
Output high current	l_{oh}	–4		mA	

H.7 Timing specifications

H.7.1 Downstream data

Table H.8 – Downstream data interface timing

Parameter	Symbol	Min.	Max.	Units
DS_CLK frequency	f		25	MHz
DS_CLK low pulse width	t_{lpw}	10		ns
DS_CLK high pulse width	t_{hpw}	10		ns
DS_CLK rise/fall time	t_{rf}		4	ns
DS_CLK jitter	t_j		97.66	ns
Input set-up time to DS_CLK	t_{su}	10		ns
Input hold time from DS_CLK	t_h	0		ns
DS_CLK to signal valid delay	t_{cq}	1	15	ns

H.7.2 Upstream data

Table H.9 – Upstream data interface timing

Parameter	Symbol	Min.	Max.	Units
US_CLK frequency	f	33	40.96	MHz
US_CLK low pulse width	t_{lpw}	6.5		ns
US_CLK high pulse width	t_{hpw}	6.5		ns
US_CLK rise/fall time	t_{rf}		1.5	ns
Input set-up time to US_CLK	t_{su}	6		ns
Input hold time from US_CLK	t_h	0		ns
US_CLK to signal valid delay	t_{cq}	1	9	ns

H.7.3 Upstream control

Table H.10 – Upstream control interface timing

Parameter	Symbol	Min.	Max.	Units
Input set-up time to US_CLK	t_{su}	6		ns
Input hold time from US_CLK	t_h	0		ns
US_CLK to signal valid delay	t_{cq}	1	9	ns
Input set-up time to TS_CLK	t_{su}	10		ns
Input hold time from TS_CLK	t_h	0		ns
TS_CLK to signal valid delay	t_{cq}	1	15	ns

H.7.4 SPI bus

Table H.11 – SPI bus timing

Parameter	Symbol	Min.	Max.	Units
SPI_CLK frequency	f		10.24	MHz
SPI_CLK low pulse width	t_{lpw}	43.9		ns
SPI_CLK high pulse width	t_{hpw}	43.9		ns
SPI_CLK rise/fall time	t_{rf}		4	ns
SPI_MOSI or SPI_MISO set-up time to SPI_CLK	t_{su}	15		ns
SPI_MOSI or SPI_MISO hold time from SPI_CLK	t_h	0		ns
SPI_SSx_N set-up time to SPI_CLK rising	t_{su}	50		ns
SPI_SSx_N setup time to SPI_CLK falling	t_{su}	25		ns
SPI_SSx_N hold time from SPI_CLK	t_h	0		ns
SPI_CLK to signal valid delay	t_{cq}	1	12	ns

H.8 Data format and usage

H.8.1 Downstream data

The data that passes from the MAC to the PHY is a stream of MPEG packets. The start of the SYNC byte is indicated by the assertion of the DS_MSYNC_N signal. Including the SYNC byte, each MPEG packet is 188 bytes in length.

The MAC MUST generate NULL MPEG packets when there are no DOCS frames to be transmitted.

H.8.2 Upstream data

H.8.2.1 Block format

Data is passed from the upstream PHY to the MAC using a combination of variable sized units called upstream data blocks. Each of these data blocks has the generic format described in Table H.12 (except for the CHANNEL data block type as indicated in clause H.8.2.8.5).

Table H.12 – Upstream data block format

Size (bytes)	Name	Description
1	Block type	Identifies the type of block
2	Block length	Length of block data field in bytes (N) Not present for CHANNEL block type
N	Block data	Block data

As can be seen from this table, each data block starts with a data block type. This type is used by the MAC to determine which type of data block data is being transferred. The data block length field contains the length in bytes of the data block data and is used by the MAC to find the end of the data block data field. In most cases, the data block type determines the format of the data block data field. The exception to this is the PHY_STATUS type where the format of the data block data field is PHY-specific.

Table H.13 gives a complete list of all block types.

Table H.13 – Upstream data block types

Type	Name	Description
0x00	Reserved	Reserved
0x01	FIRST_DATA	First data of burst contains 7 bytes of fixed-format status data and first data of burst
0x02	MIDDLE_DATA	Middle data of burst
0x03	LAST_DATA	Last data of burst contains 4 bytes of fixed format status data and last data of burst
0x04	PHY_STATUS	Status that should be passed to software The maximum length of this block is 128 bytes.
0x05	NO_BURST	Indicates that no burst was received during a transmit opportunity
0x06	CHANNEL	Used to indicate the channel to which the next data block belongs.
0x07-0xff	Reserved	Reserved

H.8.2.2 FIRST_DATA block

Table H.14 shows the format of the FIRST_DATA block.

The FIRST_DATA block is used by the PHY to transfer the beginning of a received burst. This block MUST contain the seven bytes of status information defined in the table. It MAY contain burst data as well. The block length of the FIRST_DATA block MUST NOT be less than seven. Note that N = 7 is allowed.

Table H.14 – FIRST_DATA data format

Size (bytes)	Name	Description
1	FIRST_STATUS	Bit 7:6, reserved, MUST be zero Bit 5, new UCD, 1 ≥ First burst received on new UCD Bit 4, PHY_STATUS data block present, 1 ≥ PHY_STATUS data block present Bit 3:0, IUC, taken from the upstream control interval description message
2	SID	Bit 15:14, reserved, MUST be zero Bit 13:0, SID, taken from the upstream control interval description message
4	START_MINISLOT	Derived from the upstream control interval description message parameters
N – 7	BURST_DATA	First data of burst

H.8.2.3 MIDDLE_DATA block

Table H.15 shows the format of the MIDDLE_DATA block. The MIDDLE_DATA block is used to transfer burst data.

Table H.15 – MIDDLE_DATA data format

Size (bytes)	Name	Description
N	BURST_DATA	Middle data of burst

H.8.2.4 LAST_DATA block

Table H.16 shows the format of the LAST_DATA block. The LAST_DATA block is used to transfer burst data. This block MUST contain the four bytes of status information defined in the table. It MAY also contain burst data. The block length of the LAST_DATA block MUST NOT be less than four. Note that N = 4 is allowed.

Table H.16 – LAST_DATA data format

Size (bytes)	Name	Description
N – 4	BURST_DATA	Last data of burst
1	LAST_STATUS	Bit 7:3, reserved, must be zero Bit 2, internal PHY error, 1 ≥ internal PHY error Bit 1, low energy; indicates that the burst power was below the desired threshold, 1 ≥ low energy Bit 0, high energy; indicates that the burst power was above the desired threshold, 1 ≥ high energy
1	GOOD_FEC	The number of good FEC blocks in the burst; must stop incrementing when count reaches 255; must be zero if FEC is disabled for associated interval
1	CORRECTED_FEC	The number of corrected FEC blocks in the burst; must stop incrementing when count reaches 255; must be zero if FEC is disabled for associated interval
1	UNCORRECTED_FEC	The number of uncorrected FEC blocks in the burst; must stop incrementing when count reaches 255; must be zero if FEC is disabled for associated interval

H.8.2.5 PHY_STATUS block

Table H.17 shows the format of the PHY_STATUS block. The PHY_STATUS block is used to transfer PHY unique status to the MAC. The contents of this block are vendor-unique and are unrestricted.

Table H.17 – PHY_STATUS data format

Size (bytes)	Name	Description
N	PHY_STATUS	PHY-specific status information such as channel characteristics (e.g., timing error, power error, frequency error, EQ coefficients)

H.8.2.6 NO_BURST block

Table H.18 shows the format of the NO_BURST block. This block is used by the PHY to indicate that a valid burst was not received when one was expected. Absence of a valid burst may be caused by either no transmitter, multiple transmitters or a noise-corrupted transmission. DMPI does not specify the criteria by which the PHY distinguishes between these cases.

Table H.18 – NO_BURST data format

Size (bytes)	Name	Description
2	SID_STATUS	Bit 15, collision, collision occurred Bit 14, no energy, no energy detected Bit 13:0, SID, taken from the upstream control interval description message
4	START_MINISLOT	Derived from the upstream control interval description message parameters

Table H.18 – NO_BURST data format

Size (bytes)	Name	Description
1	IUC	Bit 7:5, reserved, must be zero Bit 4, new UCD, 1=> first NO_BURST block received on new UCD Bit 3:0, IUC, taken from the upstream control interval description message
2	LENGTH	Taken from the upstream control interval description message Note that for contention intervals, this is the length of the interval and not the length of each individual transmit opportunity in the interval

H.8.2.7 CHANNEL block

Table H.19 shows the format of the CHANNEL block. The CHANNEL block is used by the PHY to indicate to which logical channel subsequent blocks belong.

Table H.19 – CHANNEL data format

Size (bytes)	Name	Description
1	CHANNEL	Bit 7:3, reserved, must be zero Bit 2:0, channel number

H.8.2.8 Block usage

H.8.2.8.1 Overview

At least one data block MUST be transferred for every transmit opportunity. If a burst is received during a transmit opportunity, the appropriate series of data blocks MUST be transferred to the MAC (FIRST_DATA, MIDDLE_DATA, LAST_DATA, PHY_STATUS). If no burst is received, a NO_BURST data block MUST be transferred unless the region was allocated to a SID that the system has reserved for no CM (e.g., the null SID as defined in clause A.2.1). Note that since contention regions have multiple transmit opportunities, more than one set of data blocks will likely be transferred to over the interface for each region (interval).

The minimum amount of payload in a data block (the length of the block data field) MUST be 16 bytes with the following exceptions:

- data blocks for bursts that are less than 16 bytes in length;
- any LAST_DATA data block.

The upstream PHY SHOULD minimize the number of data blocks required to transfer a burst so as to minimize the amount of overhead on DMPI. However, nothing specific other than what is mentioned above is required.

For non-contention intervals, the START_MINISLOT MUST be equal to the START_MINISLOT that was passed to the PHY in the corresponding interval description message (described in clause H.8.3). For contention intervals (IE types REQ and REQ/data), the PHY MUST calculate an accurate START_MINISLOT value and return it in the appropriate data block (FIRST_DATA or NO_BURST). In general terms, this means that PHY MUST calculate the START_MINISLOT for each data block by taking into account the number of mini-slots that have passed since the start of the interval. Specifically, the upstream PHY SHOULD use the IUC and SID in the upstream control interval description message to calculate a burst start offset from the original START_MINISLOT value received in this message. The offset is then added to this START_MINISLOT and returned to the MAC as the START_MINISLOT in the appropriate upstream data block.

H.8.2.8.2 Burst data transfer

The transfer of a burst MUST be accomplished by transferring the following data blocks in the following order:

- one FIRST_DATA block;
- zero to N MIDDLE_DATA blocks;
- one LAST_DATA block;
- zero or one PHY_STATUS block.

The only data block type that MAY be transferred after a FIRST_DATA data block and before a LAST_DATA data block is a MIDDLE_DATA data block. Any other data block transferred between these two data blocks MUST be discarded by the MAC.

In general, each data block will contain one FEC block of data. However, there is no specific requirement as to which data block types contain which parts of the burst data. The data MAY be distributed between the various data block types at the discretion of the PHY as long as the data block ordering shown above is maintained and the minimum block length requirements are respected. A data block type with a length of zero is also allowed. Every burst, regardless of size, MUST be transferred to the MAC using at least a FIRST_DATA block and a LAST_DATA data block. The PHY_STATUS data block is optional with its presence indicated in the FIRST_STATUS byte in the FIRST_DATA data block. The MIDDLE_DATA data block is optional.

Typically, there will be some arbitrary delay between the transfer of one data block and the transfer of the next. It is the PHY's responsibility to assure that these delays do not interfere with the PHY's ability to keep up with the incoming data rate.

Note that this series of data blocks is passed to the MAC any time a burst is received, regardless of the type of interval in which the burst was received (contention or non-contention).

H.8.2.8.3 No burst status transfer

It is sometimes useful for the system to know when no usable burst was received during a transmit opportunity. This can happen when there is no transmitter (no energy) in the opportunity, there is more than one transmitter (a collision), or noise corrupted a transmission. For a contention region, knowledge of unused opportunities or those with collisions helps software optimize its scheduling of contention regions (their duration and frequency). For non-contention regions, these same events could be an indication of a problem with a CM, or they could be a result of illegal or malicious use of the US bandwidth.

The NO_BURST data block contains two status bits. The one called "collision" indicates that a collision occurred during the transmit opportunity. The other, called "no energy", indicates that there was no energy detected during the transmit opportunity. If neither is set, it means that there was energy but that no preamble was found. Both of these bits must not be set at the same time.

H.8.2.8.4 UCD change indication

In order to allow the system to properly size grants for bandwidth requests that were received prior to a UCD change but are granted after such a UCD change, the MAC needs to be notified that a new UCD is in effect. This notification is achieved via "new UCD" status bits in the NO_BURST and FIRST_DATA data blocks. The PHY MUST set the new UCD bit of the first data block sent to the MAC after a UCD change (FIRST_DATA or NO_BURST, whichever is sent first). The new UCD bit of these data blocks MUST be zero at all other times.

H.8.2.8.5 Logical channel support

For upstream PHYs that support multiple logical channels, a data block type called CHANNEL is used to specify to that logical channel each data block belongs. This data block contains a single byte of payload that is the channel number (zero to seven inclusive). Since the data block is a fixed length and is potentially required for every other data block transferred, the length bytes are omitted from the normal data block format and only the data block type and block data are transferred. So, a CHANNEL block is always two bytes long (including the type byte).

It is important to note that the channel data block is only used to distinguish between data received on logical channels within the same RF channel. Since each PHY has its own DMPI interface, the RF channel to which data belongs is inferred by the PHY's connection to the MAC.

The CHANNEL data block is used as follows:

- The CHANNEL data block sets the "current" channel for transmitted data blocks. After reset, the MAC must set the current channel to zero.
- The current channel is always the channel number contained in the most recently transmitted CHANNEL data block. For this reason, transmission of a CHANNEL data blocks is only required when a change in the current channel is desired.

Since the MAC sets the current channel to zero prior to receipt of any CHANNEL data blocks, PHYs that support a single channel are not required to support this data block type. In cases where multiple CHANNEL data blocks are transferred in succession, the last one received prior to the transfer of one of the other data blocks will be considered valid and the others that preceded it will be ignored. NO_BURST data blocks may be preceded by a CHANNEL data block. If a series of NO_BURST data blocks for the same channel are transmitted to the MAC, only one CHANNEL data block is required (transferred prior to the first NO_BURST data block).

All data blocks associated with a single burst MUST be transferred contiguously over the upstream data interface. Specifically, this would mean that FIRST_DATA, MIDDLE_DATA, LAST_DATA, PHY_STATUS would all be transferred for a given burst of a given channel before any other data blocks were transferred for another channel. A CHANNEL data block MUST precede the first data block (NO_BURST or FIRST_DATA) that belongs to a channel that is different than the one which preceded it. The PHY MAY transfer a CHANNEL data block prior to the FIRST_DATA block of every burst. CHANNEL data blocks MUST NOT be transferred immediately before any of the other data block associated with a burst.

H.8.3 Upstream control

The upstream control interface carries two different messages. One of them is used to describe upcoming bursts. The other is used to indicate UCD changes.

The format of an upstream control message is shown in Table H.20.

Table H.20 – Upstream control message format

Size (bits)	Name	Description
3	TYPE	Message type
3	CHANNEL	Logical channel number
N	PAYLOAD	Payload of message
1	PARITY	Even parity for all bits in the message (the number of 1s across all bits in {TYPE, CHANNEL, PAYLOAD, PARITY} is even)

Table H.21 shows the message type encoding.

Table H.21 – Upstream message types

Type	Name	Description
0x0	INTERVAL_ DESCRIPTION	Describes an interval
0x1	UCD_CHANGE	Indicates a UCD change has occurred
0x2-0x7	Reserved	Reserved

H.8.3.1 Interval description message

Table H.22 describes the format of the interval description message PAYLOAD.

Table H.22 – Upstream control burst description PAYLOAD format

Size (bits)	Name	Description
14	SID	Expected SID from MAP IE
4	IUC	IUC from MAP IE
14	LENGTH	Length in mini-slots
32	START_ MINISLOT	Starting mini-slot of interval (alloc start time + offset of IE)
3	PSC	PHY_STATUS control

The MAC builds these interval description messages from the information present in the DOCS MAPs that have been generated for the logical channels that the PHY is servicing. The MAC MUST transfer only one interval description message to the PHY for an interval allocation that might describe an interval that has more than one transmit opportunity (e.g., a REQ, REQ/DATA). The MAC MAY generate interval description messages for interval allocations to the NULL SID, but MUST NOT generate interval description messages for interval allocations for the NULL SID if they overlap with interval allocations for non-NULL SIDs on other logical channels being serviced by the same PHY. Interval description messages for the NULL SID MAY be associated with any currently active logical channel. In contrast to MAP messages, the set of all interval description messages from the MAC taken together need not describe every mini-slot on the logical channels in question; the MAC MAY refrain from sending interval description messages to describe inactive time periods on any or all logical channels. So as to minimize the complexity and buffering requirements of the PHY, the MAC MUST sort the interval descriptions from all logical channels, put them into chronological order, and deliver them to the PHY in this order. Note that interval description messages MUST NOT be transferred for the NULL IE, data acknowledgement IEs, or data grants pending (since none of these is an interval allocation).

The system is allowed to schedule the initial maintenance regions of all logical channels of a physical channel to occur simultaneously. This type of overlap MUST be handled as follows:

- The MAC MUST transfer an interval description for only one of the logical channels.
- The interval description that is transferred MUST be the one with the earliest start time. If more than one interval description has the earliest start time, the MAC MAY choose any of these overlapping interval descriptions to pass to the PHY.
- The PHY MUST accept any of the logical channel numbers it supports for this interval description.

The system software is responsible for knowing that bursts received during initial ranging could be from CMs on any of the logical channels.

It is possible for there to be illegal overlap of intervals for the logical channels. An illegal overlap is defined to be an overlap of intervals other than initial maintenance. The PHY MAY detect these illegal overlaps. If the PHY performs this function, it MUST generate an interrupt to alert the system of such an event. It MUST capture the illegally overlapped interval description and hold it in SPI bus-accessible registers until software acknowledges its receipt.

The PSC field of the interval description message is used to control the contents of the PHY_STATUS block. The usage of this field is summarized below:

- If PSC = 000, the contents of the PHY_STATUS block is determined through PHY programmable registers.
- If PSC is any other value, the contents of the PHY_STATUS block are vendor-specific.

The MAC and PHY MUST support PSC = 000.

The MAC and PHY MAY support other values.

H.8.3.2 UCD change message

Table H.23 describes the format of the UCD change message PAYLOAD.

Table H.23 – UCD change PAYLOAD format

Size (bits)	Name	Description
8	CCC	Configuration change count from the MAP

The MAC MUST send this message before sending the first interval description message after a UCD change. This message MUST NOT be sent at any other time.

H.8.4 SPI bus

In order to perform an SPI bus transaction, the master MUST drive SPI_MOSI with a bitstream of the following format:

Table H.24 – SPI bus transaction format

Size (bits)	Name	Description
4	DEVICE_ID	Device ID
3	RSVD	Reserved
1	WRITE	1 = write, 0 = read
16	REGISTER_ADD	Register address
N × 8	WRITE_DATA	Write data; ignored for reads

The DEVICE_ID is used to address PHY devices that are integrated into the same physical package and share a single SPI select. DEVICE_ID MUST be zero for accesses to single PHY devices.

Annex I

(Set aside)

NOTE – This annex is left blank intentionally to avoid any possible confusion with Appendix I.

Annex J

Japan specification additions

(This annex forms an integral part of this Recommendation)

This annex applies to the third technology option referred to in clause 1.1, "scope".

This annex describes the physical layer MAC layer specifications required for what are generally called Japan-DOCS cable-modems. This is an optional annex and in no way affects certification of North American and European DOCS 1.x and DOCS 2.0 modems.

The numbering of the clauses has been made so that the suffix after the J refers to the part of the specification which has changed. As a consequence, some clauses are missing in this annex because no change is required.

J.1 Scope and purpose

See clause 1.

J.2 References

See clause 2.

J.3 Glossary

See clause 3.

J.4 Functional assumptions

See clause 4.

J.4.1 Broadband access network

A coaxial-based broadband access network is assumed. This may take the form of either an all-coax or hybrid-fibre/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 analogue transmission. The key functional characteristics assumed in this annex are the following:

- two-way transmission;
- a maximum optical/electrical spacing between the CMTS and the most distant CM of 160 km, although typical maximum separation may be 16 to 24 km;
- a maximum differential optical/electrical spacing between the CMTS and the closest and most distant modems of 160 km, although this would typically be limited to 24 km.

J.4.2 Equipment assumptions

J.4.2.1 Frequency plan

In the downstream direction, the cable system is assumed to have a passband with a lower edge 90 MHz and an upper edge that is implementation-dependent but is typically in the range of 350 MHz to 770 MHz. Within that passband, NTSC analogue television signals in 6-MHz channels are assumed to be present on the standard Japan frequency plans, as well as other narrow-band and wideband digital signals.

In the upstream direction, the cable system MAY have a subsplit (10 MHz to 55 MHz) passband. NTSC analogue television signals in 6 MHz channels MAY be present, as well as other signals.

J.4.2.2 Compatibility with other services

Refer to clause 4.2.2.

J.4.2.3 Fault isolation impact on other users

Refer to clause 4.2.3.

J.4.2.4 Cable system terminal devices

The CM MUST meet and SHOULD exceed all applicable regulations for cable system termination devices and cable ready consumer equipment as defined in voluntary control council for interference by information technology equipment (VCCI). None of these specific requirements may be used to relax any of the specifications contained elsewhere within VCCI. The CMTS SHOULD meet and exceed all applicable regulations for Class-B ITE as well.

J.4.3 RF channel assumptions

Refer to clause 4.3.

J.4.3.1 Transmission downstream

The RF channel transmission characteristics of the cable network in the downstream direction are described in Table J.4-1. 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 J.4-1 assume average power in a bandwidth in which the impairment levels are measured in a standard manner for cable TV system. For analogue signal levels, the numbers in Table J.4-1 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 Recommendation.

Table J.4-1 – Assumed downstream RF channel transmission characteristics (see Note 1)

Parameter	Value
Frequency range	Cable system normal downstream operating range is from 90 MHz to as high as 770 MHz.
RF channel spacing (design bandwidth)	6 MHz
Transit delay from headend to most distant customer	≤0.800 ms (typically much less)
Carrier-to-noise ratio in a 6-MHz band	Not less than 26 dBrms (@ 5.274 MHz) for 64QAM Not less than 33 dBrms (@ 5.274 MHz) for 256QAM (Note 2)
Carrier-to-composite triple beat distortion ratio	Not less than 40 dBrms for 64QAM Not less than 51 dBrms for 256QAM (Note 2)
Carrier-to-any other discrete interference (ingress)	Not less than 26 dBrms for 64QAM Not less than 33 dBrms for 256QAM (Note 2)
Amplitude ripple	3 dB within the design bandwidth
Micro reflections bound for dominant echo	Figure J.4-1
Maximum analogue video carrier level at the CM input	85 dBµV peak
Maximum number of carriers	111 (770 MHz system)
NOTE 1 – Transmission is from the headend combiner to the CM input at the customer location. NOTE 2 – Measured relative to a QAM signal level (rms) that is –10 dB for 64QAM, –4 dB for 256QAM to the nominal video level (peak) in the plant.	

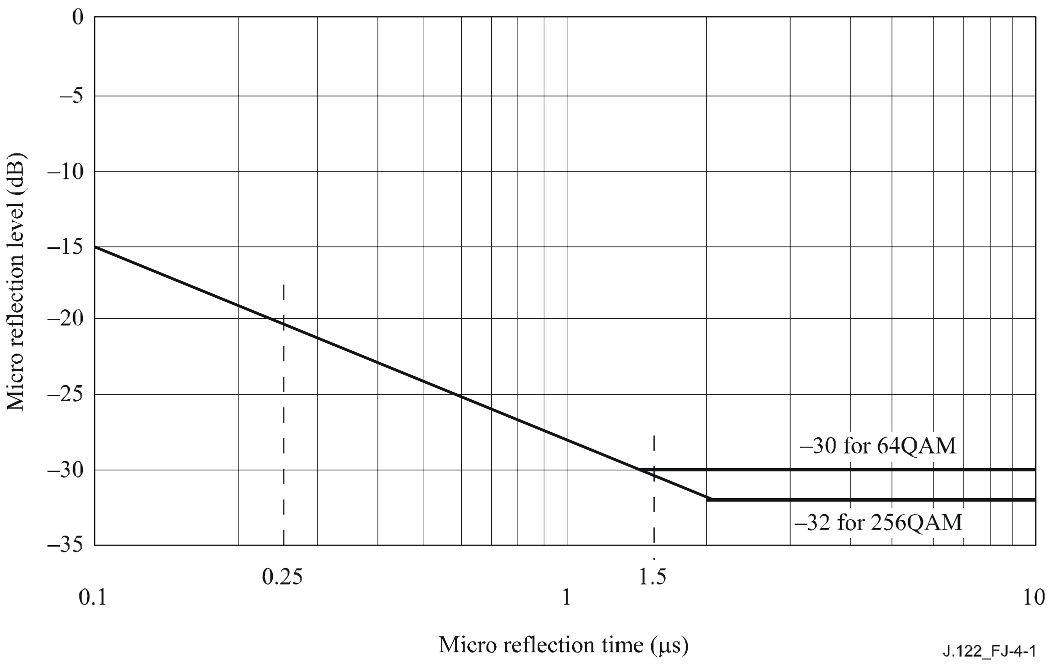

Figure J.4-1 – Micro reflections bound for dominant echo

J.4.3.2 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 J.4-2. All conditions are present concurrently.

Table J.4-2 – Assumed upstream RF channel transmission characteristics (see Note 1)

Parameter	Value
Frequency range	10 MHz to 55 MHz edge to edge
Transit delay from the most distant CM to the nearest CM or CMTS	≤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 (Note 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 (Notes 3 and 4)
Amplitude ripple 10 MHz to 55 MHz	0.5 dB/MHz
Group delay ripple 10 MHz to 55 MHz	200 ns/MHz

Table J.4-2 – Assumed upstream RF channel transmission characteristics (see Note 1)

Parameter	Value
Micro reflections – single echo	–10 dB @ ≤ 0.5 µs –20 dB @ ≤ 1.0 µs –30 dB @ > 1.0 µs
Seasonal and diurnal reverse gain (loss) variation	Not greater than 14 dB min. to max.
NOTE 1 – Transmission is from the CM output at the customer location to the headend. NOTE 2 – 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. NOTE 3 – Amplitude and frequency characteristics sufficiently strong to partially or wholly mask the data carrier. NOTE 4 – Impulse noise levels more prevalent at lower frequencies (<15 MHz).	

J.4.3.2.1 Availability

Refer to clause 4.3.2.1.

J.4.4 Transmission levels

The nominal power level of the downstream CMTS signal(s) within a 6 MHz channel is targeted to be in the range –10 dBc to –6 dBc relative to the analogue video carrier level and will normally not exceed the analogue video carrier level. The 256QAM downstream carrier level SHOULD be carefully chosen for two reasons. One is to avoid any interference to the adjacent analogue video carrier, the other is to maintain required carrier-to-noise ratio. Normally the 256QAM downstream signal MAY NOT be allocated to any channels that are adjacent to the analogue video carrier.

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.

J.4.5 Frequency inversion

Refer to clause 4.5.

J.5 Communication protocols

Refer to clause 5.

J.6 Physical media dependent sublayer specification

J.6.1 Scope

This clause applies to the third technology option referred to in clause 1.1, "scope". In those cases where the requirements for three technology options are identical, a reference is provided to the main text.

Whenever any reference in this clause to spurious emissions conflicts with any legal requirement for the area of operation, the latter shall take precedence.

J.6.2 Upstream

J.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 a single RF channel unless reconfigured to change channels. 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 Recommendation, the following naming conventions are used. For TDMA, the term "modulation rate" refers to the RF channel symbol rate (144 to 4608 ksymb/s). For S-CDMA, the term "chip rate", which is the modulation rate (1152 to 4608 kHz) of a single bit of the S-CDMA spreading code, may be used interchangeably with "modulation rate". 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 entity 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 are used to pad data blocks; this indicates data with zero-valued binary bits, which result in non-zero transmitted RF energy. In other cases, a numerical zero is used; this denotes, for example, symbols that 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 levels.

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 centre of the last symbol of one burst and the centre of the first symbol of the preamble of an immediately following burst are separated by at least the duration of five symbols. The guard time 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 clause 6.2.19. Maximum timing error and guard time may vary with CMTSs from different vendors.

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. 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 that 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 centred 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) narrow-band 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 clauses 6.2.19, 6.2.5.1, and A.2).

J.6.2.2 Signal processing requirements

Refer to clause 6.2.2.

J.6.2.3 Modulation formats

The modulation formats listed here specify requirements for J.122-compliant equipment. Cable operators are free to configure the modulation format to best address their system characteristics and application requirements.

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

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

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

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

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

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

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

J.6.2.4 R-S encode

Refer to clause 6.2.4.

J.6.2.5 R-S frame structure

Refer to clause 6.2.5.

J.6.2.6 TDMA byte interleaver

Refer to clause 6.2.6.

J.6.2.7 Scrambler (randomizer)

Refer to clause 6.2.7.

J.6.2.8 TCM encoder

Refer to clause 6.2.8.

J.6.2.9 Preamble prepend

Refer to clause 6.2.9.

J.6.2.10 Modulation rates

In TDMA mode, the CM upstream modulator MUST provide all modulations at 144, 288, 576, 1152, 2304 and 4608 ksymb/s. In S-CDMA mode, the CM upstream modulator MUST provide all modulations at 1152, 2304 and 4608 ksymb/s.

In TDMA mode, the CMTS upstream demodulator MUST be able to support demodulation at 144, 288, 576, 1152, 2304 and 4608 ksymb/s. In S-CDMA mode, the CMTS upstream demodulator MUST be able to support demodulation at 1152, 2304 and 4608 ksymb/s.

This variety of modulation rates, and flexibility in setting upstream carrier frequencies, permits operators to position carriers in gaps in the pattern of narrow-band ingress, as discussed in Annex G.

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.

J.6.2.11 S-CDMA framer and interleaver

J.6.2.11.1 S-CDMA framing considerations

Refer to clause 6.2.11.1.

J.6.2.11.2 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 (these messages are described in clause 8.3.2, "time synchronization (SYNC)", and clause 8.3.3, "upstream channel descriptor (UCD)".)

Mini-slots are mapped onto frames starting at the first active code (usually code number 0), are numbered sequentially through the remainder of the frame (code number 127), and then wrap to the next sequential frame. 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 9.216-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 MUST identify frame boundaries relative to the timestamp counter on a periodic basis. This is called the timestamp snapshot and must be sent in the UCD for each upstream S-CDMA channel.

The CMTS MUST maintain a frame counter and a mini-slot counter, and MUST sample these values along with the timestamp, on a frame boundary, as shown in Figure J.6-5. The CMTS MUST obtain a new sample prior to sending each UCD message.

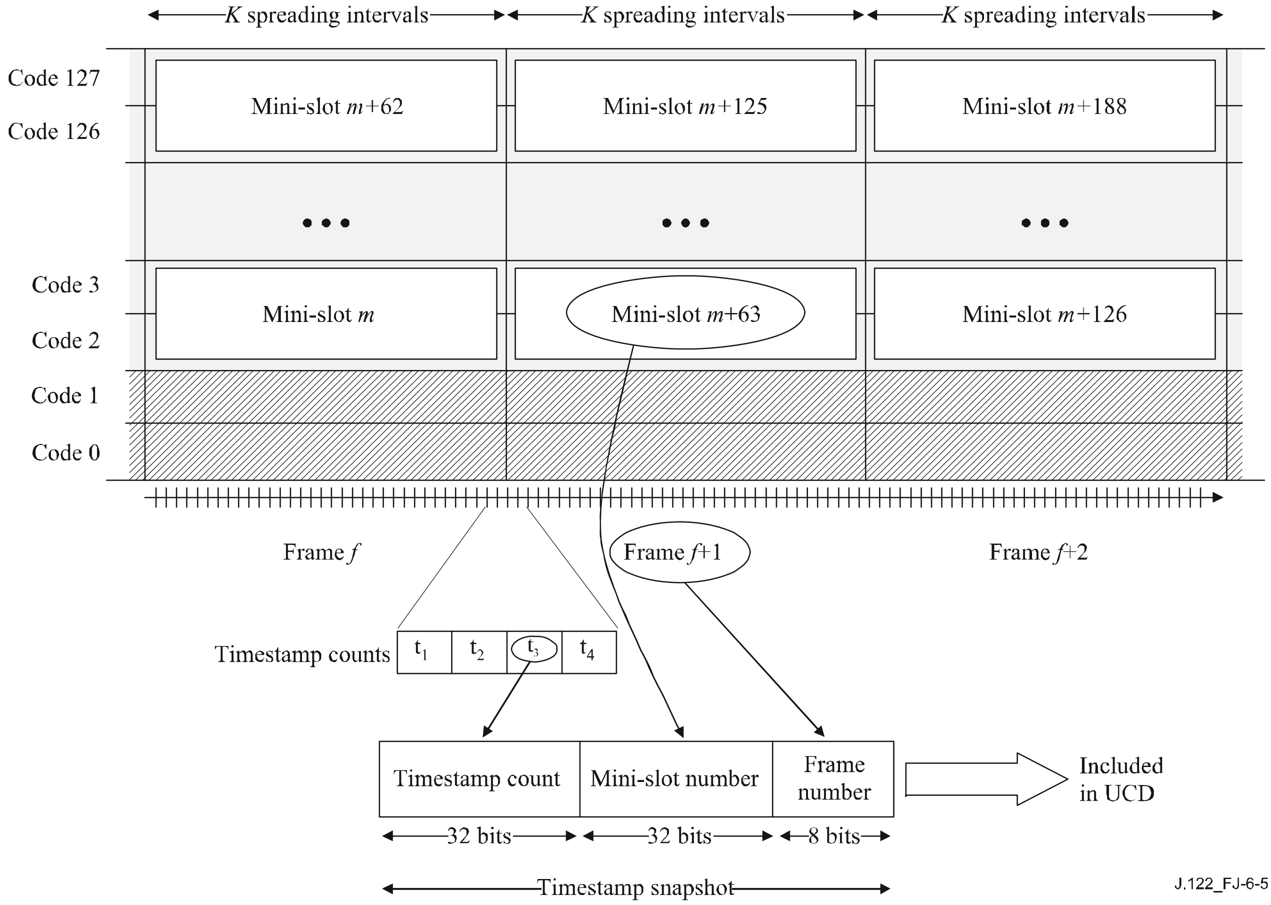

Figure J.6-5 – Timestamp snapshot

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 clause 6.2.14.1, "code hopping".

The frame structure above relates to the entire upstream and not necessarily to the transmission from a single CM. The codes are resources that 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 clause 6.2.12.

J.6.2.11.2.1 Mini-slot numbering parameters in UCD

Refer to clause 6.2.11.2.1.

J.6.2.11.2.2 Mini-slot numbering examples

Refer to clause 6.2.11.2.2.

J.6.2.11.3 Transmission time

Refer to clause 6.2.11.3.

J.6.2.11.4 Latency considerations

S-CDMA frame timing is derived directly from (is phase-locked to) the 9.216-MHz CMTS master clock. Based on the allowable signalling rates and the fact that there are 128 signalling periods in a spreading interval, the S-CDMA frame time MUST always be a multiple of 27.7 µs.

Selecting the number of spreading intervals per frame and the signalling rate therefore exactly define the S-CDMA frame duration. As a specific example, a burst profile defined with 10 spreading intervals per frame with a signalling rate of 2.304 Mbaud would result in a frame duration of 555.5 µs.

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 clause 6.2.17.

J.6.2.11.5 Spreader-off bursts for maintenance on S-CDMA channel

Refer to clause 6.2.11.5.

J.6.2.12 S-CDMA framer

Refer to clause 6.2.12.

J.6.2.13 Symbol mapping

Refer to clause 6.2.13.

J.6.2.14 S-CDMA spreader

Refer to clause 6.2.14.

J.6.2.15 Transmit pre-equalizer

Refer to clause 6.2.15.

J.6.2.16 Spectral shaping

The upstream transmitter MUST approximate a Nyquist square-root raised-cosine pulse-shaping filter with roll-off factor alpha = 0.25. The –30 dB transmitted bandwidth MUST NOT exceed the channel width values in Table J.6-1. The channel width values are given analytically by:

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

The occupied spectrum MUST NOT exceed the channel widths shown in Table J.6-1.

Table J.6-1 – Maximum channel width

Modulation rate (kHz)	Channel width (kHz)	Recommended channel spacing (kHz)
144	180	187.5
288	360	375
576	720	750
1152	1440	1500
2304	2880	3000
4608	5760	6000

J.6.2.16.1 Upstream frequency agility and range

The upstream PMD sublayer MUST support operation over the frequency range of 10-55 MHz edge to edge.

Offset frequency resolution MUST be supported having a range of ±32 kHz (increment 1 Hz; implement within ±10 Hz).

J.6.2.16.2 Spectrum format

Refer to clause 6.2.16.2.

J.6.2.17 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 specified below.

The CM MAP processing delay, $D_{p,}$, is given by the equations:

$$D_p = 200 + \frac{M}{4.608}\,\mu\text{s},$$

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

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 DOCS 1.x mode, $M = 0$. Note that in the above equations, the values for B_r, and $I_r \times N_r$, 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 1117 µs, assuming a modulation rate of 4.608 MHz.

NOTE 1 – The CM MAP processing delay does not include downstream FEC de-interleaving delay.

NOTE 2 – 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.

J.6.2.18 Transmit power requirements

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

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

The protocol by which power adjustments are performed is defined in clause 11.2.4, "ranging and automatic adjustments". Such adjustments by the CM MUST be within the ranges of tolerances described below. A CM MUST confirm that the transmit power limits are met after a RNG-RSP is received or after a UCD change.

Transmit power is defined as the average RF power in the occupied bandwidth (channel width) transmitted in the data symbols of a burst, assuming equally likely QAM symbols, measured at the F-connector of the CM.

Maximum and minimum transmit power level requirements refer to the CM's target transmit power level, defined as the CM's estimate of its actual transmit power. The actual transmitted power MUST be within ±2 dB of the target power. The target transmit power MUST be variable over the range specified in Table J.6-4.

Transmit power, as reported by the CM in the MIB, is referenced to the 64QAM constellation. When transmitting with other constellations, a slightly different transmit power will result, depending on the constellation gain in Table J.6-2 (see clause 6.2.13). As an example, if the reported power is 90 dBµV, 64QAM will be transmitted with a target power of 90 dBµV, while QPSK will be transmitted with 88.82 dBµV.

Table J.6-2 – Constellation gains and power limits

Constellation	Constellation gain G_{const} relative to 64QAM (dB)	P_{min} (dBµV)	P_{max} (dBµV) TDMA	P_{max} (dBµV) S-CDMA	P_{min} G_{const} (dBµV)	P_{max} G_{const} (dBµV) TDMA	P_{max} G_{const} (dBµV) S-CDMA
QPSK	–1.18	68	118	113	69.18	119.18	114.18
8QAM	–0.21	68	115	113	68.21	115.21	113.21
16QAM	–0.21	68	115	113	68.21	115.21	113.21
32QAM	0.00	68	114	113	68.00	114.00	113.00
64QAM	0.00	68	114	113	68.00	114.00	113.00
128QAM	0.05	68	N/A	113	67.95	N/A	112.95

The actual transmitted power within a burst MUST be constant to within 0.1 dB peak to peak. 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.

J.6.2.18.1 TDMA transmit power calculations

In TDMA mode, the CM determines its target transmit power P_t as follows. Define:

P_r =	Reported power level (dBµV) of CM in MIB (refers to 64QAM constellation)
ΔP =	Power level adjustment (dB); for example, as commanded in ranging response message
G_{const} =	Constellation gain (dB) relative to 64QAM constellation (see Table J.6-2)
P_{min} =	Minimum target transmit power permitted for the CM per clause 6.2.21.1 (see Table J.6-2)
P_{max} =	Maximum target transmit power permitted for the CM per clause 6.2.21.1 (see Table J.6-2)
P_{hi} =	min($P_{max} – G_{const}$) over all burst profiles used by the CM (see Table J.6-2)
P_{low} =	max($P_{max} – G_{const}$) over all burst profiles used by the CM (see Table J.6-2)
P_t =	Target transmit power level (dBµV) of CM (actual transmitted power as estimated by CM)

The CM updates its reported power by the following steps:

1)	P_r =	$P_r + \Delta P$	//Add power level adjustment to reported power level
2)	P_r =	$\min[P_r, P_{hi}]$	//Clip at max power limit
3)	P_r =	$\max[P_r, P_{low}]$	//Clip at min power limit

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

Usually, the reported power level is a relatively constant quantity, while the transmitted power level varies dynamically as different burst profiles, with different constellation gains, are transmitted. A CM's target transmit power MUST never be below P_{min}, or above P_{max}. This implies that, in some cases, the extreme transmit power levels (e.g., 118 dBµV for QPSK and 68 dBµV) may not be permitted if burst profiles with multiple constellations are active. Also, if only QPSK is used, the reported power may be greater than 118 dBµV, although the target transmit power will not exceed 118 dBµV.

For example, if only QPSK and 64QAM burst profiles are active, P_{hi} = 114 dBµV and P_{low} = 69.2 dBµV. The maximum permitted QPSK transmitted power is 114 – 1.2 = 112.8 dBµV, the minimum QPSK power is 69.2 –1.2 = 68 dBµV, the maximum 64QAM power is 114 dBµV, and the minimum 64QAM power is 69.2 dBµV.

J.6.2.18.2 S-CDMA transmit power calculations

In S-CDMA mode, the CM determines its target transmit power P_t as follows. Define:

P_{hi} =	$\min[P_{max} – G_{const}]$ over all burst profiles used by the CM (see Table J.6-2)
P_{low} =	$\max[P_{min} – G_{const}] + 10 \log(number\_active\_codes/number\_of\_codes\_per\_mini\text{-}slot)$ where the maximum is over all burst profiles used by the CM (see Table J.6-2)

The CM updates its reported power by the following steps:

1)	P_r =	$P_{min} + \Delta P$	//Add power level adjustment to reported power level
2)	P_r =	$\min[P_r, P_{hi}]$	//Clip at max power limit
3)	P_r =	$\max [P_r, P_{low}]$	//Clip at min power limit

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

$$P_{t,i} = P_r + G_{const,i} – 10 \log(number\_active\_codes)$$

i.e., the reported power 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 P_t in a frame is the sum of the individual transmit powers $P_{t,i}$ of each code, where the sum is performed using absolute power quantifies (non-dB domain).

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

The transmitted power level 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 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 (e.g., 68 and 113 dBμV) may not be permitted. Also if, for example, only QPSK is used, the reported power may be greater than 113 dBμV, although the target transmit power will not exceed 113 dBμV.

If, for example, QPSK and 64QAM burst profiles are active, P_{hi} = 113 dBμV and P_{low} = 69.2 dBμV. The maximum permitted QPSK transmitted power is 113 – 1.2 = 111.8 dBμV when all active codes are transmitted, the minimum QPSK power is 69.2 – 1.2 = 68 dBμV when one mini-slot is transmitted, the maximum 64QAM power is 113 dBμV when all active codes are transmitted, and the minimum 64QAM power is 69.2 dBμV with one mini-slot transmitted. The minimum QPSK power permitted while transmitting, for example, 2 mini-slots is 71 dBμV, and the minimum 64QAM power permitted while transmitting 2 mini-slots is 72.2 dBμV.

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.

J.6.2.18.3 Transmit power step size

Refer to clause 6.2.18.3.

J.6.2.19 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:

i) the modulation rate (six rates from 144 ksymb/s to 4.608 Msymb/s in octave steps);
ii) the centre frequency (Hz);
iii) the 1536-bit preamble superstring; and
iv) the S-CDMA channel parameters.

The channel parameters are further described in Table 8-18; these characteristics are shared by all users on a given channel. The burst profile attributes are listed in Table J.6-3 and are further described in Table 8-19; these parameters are the shared attributes corresponding to a burst type. 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 J.6-4.

The CM MUST generate each burst at the appropriate time as conveyed in the mini-slot grants provided by the CMTS MAPs (see clause 8.3.4).

The CM MUST support all burst profiles commanded by the CMTS via the burst descriptors in the UCD (see clause 8.3.3), and subsequently assigned for transmission in a MAP (see clause 8.3.4).

Table J.6-3 – Burst profile attributes

Burst profile attributes	Configuration settings
Modulation	QPSK, 8QAM, 16QAM, 32QAM, 64QAM, 128QAM (TCM only)
Differential encoding	On/off
TCM encoding	On/off
Preamble length	0 to 1536 bits (see clause J.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 FEC on)
Scrambler seed	15 bits
Maximum burst length (mini-slots)[a)]	0 to 255
Guard time	5 to 255 modulation intervals 1 for S-CDMA channels
Last codeword length	Fixed, shortened
Scrambler on/off	On/off
Byte interleaver depth (I_r)[b)]	0 to floor (2048/N_r)[c)]
Byte interleaver block size (B_r)[d)]	2 × N_r to 2048
S-CDMA spreader[e)]	On/off
S-CDMA codes per subframe[e)]	1 to 128
S-CDMA interleaver step[e)]	1 to (spreading intervals per frame – 1)

a) 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.

b) If depth = l, no interleaving; if depth = 0, dynamic mode.

c) N_r is the R-S codeword size: k + 2T as defined in clause 6.2.6.1.

d) Used only in dynamic mode.

e) Used only for S-CDMA channels.

Table J.6-4 – User-unique burst parameters

User-unique parameter	Configuration settings
Power level	TDMA: +68 to +114 dBμV (32QAM, 64QAM) +68 to +115 dBμV (8QAM, 16QAM) +68 to +118 dBμV (QPSK) S-CDMA: +68 to +113 dBμV (all modulations) 1-dB steps
Offset frequency[a)]	Range = ±32 kHz; increment = 1 Hz; implement within ±10 Hz
Ranging offset	Integer part: 0 to (2^{16}–1), increments of 6.9$\underline{4}$ μs/64 Fractional part: unsigned 8-bit fractional extension, units of 6.9$\underline{4}$ μs/(64 × 256) = 0.42385525 ns
Burst length (mini-slots) if variable on this channel (changes burst-to-burst)	1 to 255 mini-slots
Transmit equalizer coefficients	Up to 64 coefficients; 4 bytes per coefficient: 2 real, 2 imaginary

a) Values in this table apply for this given channel and modulation rate.

NOTE – Underlined digit represents an infinite series of the same digit: 6.9$\underline{4}$ = 6.9444…

The CM MUST implement the offset frequency to within ±10 Hz.

J.6.2.19.1 Ranging offset

Ranging offset is the delay correction applied by the CM to the CMTS upstream frame time derived at the CM. It is an advancement equal to roughly the round-trip delay of the CM from the CMTS, and is needed to synchronize upstream transmissions in the TDMA and S-CDMA schemes. The CMTS MUST provide feedback correction for this offset to the CM, based on reception of one or more successfully received bursts (i.e., satisfactory result from each technique employed: error correction and/or CRC), with resolution of 1/16'384 of the frame tick increment (6.9$\underline{4}$ μs/(64 × 256) = 0.42385525 ns). The CMTS sends adjustments to the CM, 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 correction 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 ±0.25 μs plus ±1/2 symbol owing to resolution. As an example, for the maximum symbol rate of 4608 ksymb/s, the corresponding symbol period would be 217 ns, the corresponding maximum resolution for the timing correction MUST be 217 ns, and the corresponding minimum accuracy MUST be ±359 ns. The accuracy of CM burst timing of ±0.25 μs plus ±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.

For S-CDMA channels, the CM MUST implement the ranging offset correction to within ±0.01 of the nominal chip period. As an example, for the maximum chip rate of 4608 kHz, the corresponding maximum resolution for the timing correction would be 217 ns × (±0.01) or roughly ±2 ns.

J.6.2.19.2 TDMA reconfiguration times

Refer to clause 6.2.19.2.

J.6.2.19.3 S-CDMA reconfiguration times

Refer to clause 6.2.19.3.

J.6.2.20 Burst timing convention

Refer to clause 6.2.20.

J.6.2.21 Fidelity requirements

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

J.6.2.21.1 Spurious emissions

The spurious emissions MUST NOT exceed the values shown in Table J.6-5.

Table J.6-5 – Spurious emissions

Frequency	Active Period	Inactive Period
10 to 55 MHz, inband	Less than –40 dBc[b) c)]	Less than the greater of –72 dBc 2,4 or +1 dBμV[d)]
10 to 55 MHz, outband including adjacent band, carrier-related band and other noise powers within 10 to 55 MHz[a)]	Less than –45 dBc[b) c)]	
55 to 70 MHz	Less than the greater of –45 dBc[b) c)] or +35 dBμV[e)]	Less than +35 dBμV[e)]
70 to 90 MHz	Less than +35 dBμV[e)]	Less than +25 dBμV[e)]
90 to 770 MHz	Less than +25 dBμV[e)]	

a) The out-of-band spurious emissions are measured starting from the neighbouring –3 dB response frequencies of the adjacent channels at symbol rate R. To maintain 6-MHz upstream channelization, this frequency offset is given as the transmitted channel centre frequency ± Δf, where $\Delta f = [(1 + 2\alpha)/2 + 240'000/4'608'000] \times R$ Hz = 0.8 R Hz for roll-off factor $\alpha = 0.25$.

b) The modulated carrier power is measured within the occupied bandwidth of the data signal, which is equal to 1.25 times the upstream transmission symbol rate.

c) The total spurious emissions power is measured within the Nyquist bandwidth of the transmitted signal, which is equal to the upstream transmission symbol rate.

d) The total spurious emissions power is measured within the Nyquist bandwidth of the lowest transmitted symbol rate, which is equal to 144 kHz. For any other symbol rate B Hz, this value is adjusted by 10 log (B/144,000) dB.

e) The total spurious emissions power is measured within the NTSC bandwidth, which is equal to 4.0 MHz.

J.6.2.21.1.1 Adjacent channel spurious emissions

Clause 6.2.21.1.1 is not applicable to Annex J.

J.6.2.21.1.2 Spurious emissions in 5 to 42 MHz

Clause 6.2.21.1.2 is not applicable to Annex J.

J.6.2.21.2 Spurious emissions during burst on/off transients

Clause 6.2.21.2 is not applicable to Annex J.

J.6.2.21.3 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.

J.6.2.21.3.1 Definitions

Symbol MER: MER_{symb} is defined as follows for TDMA or S-CDMA symbols. The transmitted RF waveform (after appropriate downconversion) 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_{symb}. 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_{symb}. MER_{symb} is defined at the F connector of the CM, except that when an echo channel is inserted, MER_{symb} is defined at the output of the echo channel. MER_{symb} is computed by the formula:

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

where:

E_{av} is the average constellation energy for equally likely symbols (see clause J.6.2.13 and Figure 6-18)

N is the number of symbols averaged

e_j is the error vector from the jth received symbol to the ideal transmitted QAM symbol on the grid of Figure 6-18

For S-CDMA, MER_{symb} 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 funnelling 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 $\overline{P_k}$ defined in clause J.6.2.14, "S-CDMA spreader".

MER_{chip} is defined as follows. The transmitted RF waveform (after appropriate downconversion) 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:

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

where:

p_j is the jth ideal transmitted composite chip

r_j is the jth received composite chip

N is the number of composite chips observed

J.6.2.21.3.2 Requirements

Unless otherwise stated, the MER MUST meet or exceed the following limits over the full transmit power range of Table J.6-4 for each modulation, each modulation rate, and over the full carrier frequency range, and for S-CDMA, over any valid number of active and allocated codes. The 10-55 MHz carrier frequency range refers more precisely to [10 MHz + modulation rate × 1.25/2] to [55 MHz – 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 2.304 MHz and below:

MER_{symb} ≥ 30 dB over 20 to 41 MHz carrier frequency.

MER_{symb} ≥ 27 dB over 15 to 20 MHz and 41 to 47 MHz carrier frequency.

MER_{symb} ≥ 23 dB over 10 to 15 MHz and 47 to 55 MHz carrier frequency.

Case 1b: for modulation rate 4.608 MHz:

MER_{symb} ≥ 27 dB over 20 to 41 MHz carrier frequency.

MER_{symb} ≥ 24 dB over 15 to 20 MHz and 41 to 47 MHz carrier frequency.

MER_{symb} ≥ 20 dB over 10 to 20 MHz and 47 to 55 MHz carrier frequency.

Case 2: Flat channel, transmit equalization ON

Case 2a: For TDMA/QPSK, MER_{symb} ≥ 30 dB.

Case 2b: For S-CDMA and all TDMA modulations except QPSK, MER_{symb} ≥ 35 dB.

Case 2c: For S-CDMA, MER_{chip} ≥ 33 dB.

Case 3: Echo channel, transmit equalization ON

Case 3a: In the presence of a single echo selected from the channel micro-reflections defined in Figure J.4-1, 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 Figure J.4-1 (at most one of each specified magnitude and delay), the measured MER_{symb} MUST be ≥ 29 dB.

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.

J.6.2.21.4 Filter distortion

Refer to clause 6.2.21.4.

J.6.2.21.5 Carrier phase noise

Refer to clause 6.2.21.5.

J.6.2.21.6 Channel frequency accuracy

Refer to clause 6.2.21.6.

J.6.2.21.7 Modulation rate accuracy

Refer to clause 6.2.21.7.

J.6.2.21.8 Modulation timing jitter

Refer to clause 6.2.21.8.

J.6.2.22 Upstream demodulator input power characteristics

The maximum total input power to the upstream demodulator MUST NOT exceed 95 dBμV in the 10-55 MHz frequency range of operation.

The intended received power in each carrier MUST be within the values shown in Table J.6-6.

Table J.6-6 – Maximum range of commanded nominal receive power in each carrier

Modulation rate (kHz)	Maximum range (dBμV)
144	+44 to +72
288	+47 to +75
576	+50 to +78
1152	+53 to +81
2304	+56 to +84
4608	+59 to +87

The demodulator MUST operate within its defined performance specifications with received bursts within ±6 dB of the nominal commanded received power.

J.6.2.23 Upstream electrical output from the CM

The CM MUST output an RF modulated signal with the characteristics delineated in Table J.6-7.

Table J.6-7 – Electrical output from the CM

Parameter	Value
Frequency	10 to 55 MHz edge to edge
Level range (one channel)	TDMA: +68 to +114 dBμV (32QAM, 64QAM) +68 to +115 dBμV (8QAM, 16QAM) +68 to +118 dBμV (QPSK) S-CDMA: +68 to +113 dBμV (all modulations of S-CDMA)
Modulation type	QPSK, 8QAM, 16QAM, 32QAM, 64QAM and 128QAM
Modulation rate (nominal)	TDMA: 144, 288, 576, 1152, 2304 and 4608 kHz S-CDMA: 1152, 2304 and 4608 kHz
Bandwidth	TDMA: 187.5, 375, 750, 1500, 3000 and 6000 kHz S-CDMA: 1500, 3000 and 6000 kHz
Output impedance	75 ohms
Output return loss	>6 dB (10-55 MHz)
Connector	F connector per [IEC 60169-24] (common with the input)

J.6.3 Downstream

J.6.3.1 Downstream protocol

The downstream PMD sublayer MUST conform to Annex C of [ITU-T J.83].

J.6.3.2 Interleaving

The interleave method for 64QAM MUST be compliant with Annex C of [ITU-T J.83]. The interleave method for 256QAM is equal to Annex C of [ITU-T J.83] except for interleave depth values. The interleave depth I = 12 MUST be supported by the CM and CMTS. The CM and CMTS MAY support interleaver depths of I = 34 or I = 204. Table J.6-8 shows the interleaver characteristics at 5.274 Msymb/s.

Table J.6-8 – Interleaver characteristics (@ 5.274 Msymb/s)

I (number of taps)	J (increment)	Burst protection 64QAM/256QAM	Latency 64QAM/256QAM
12	17	24 µs/18 µs	0.57 ms/0.43 ms
34	6	– /51 µs	– /1.28 ms
204	1	– /300 µs	– /7.85 ms

J.6.3.3 Downstream frequency plan

The downstream channel MUST support a 90 MHz to 770 MHz frequency range edge to edge.

J.6.3.4 CMTS electrical output

The transmission signal level at CMTS output connector MUST be adjustable over the range of +100 to +120 dBµVrms.

J.6.3.4.1 Modulation method

The modulation method MUST be 64QAM and 256QAM for downstream.

J.6.3.4.2 Signal constellation diagram

The signal constellation diagram and phase shift rule for 64QAM MUST be compliant with Annex C of [ITU-T J.83]. 256QAM signal constellation diagram MUST be compliant with Figure J.6-6, when the interleave depth I = 12 is selected. In case that the interleaver depth I = 34 or 204 is selected, the signal constellation diagram and phase shift rule MUST be as shown in Figure J.6-7.

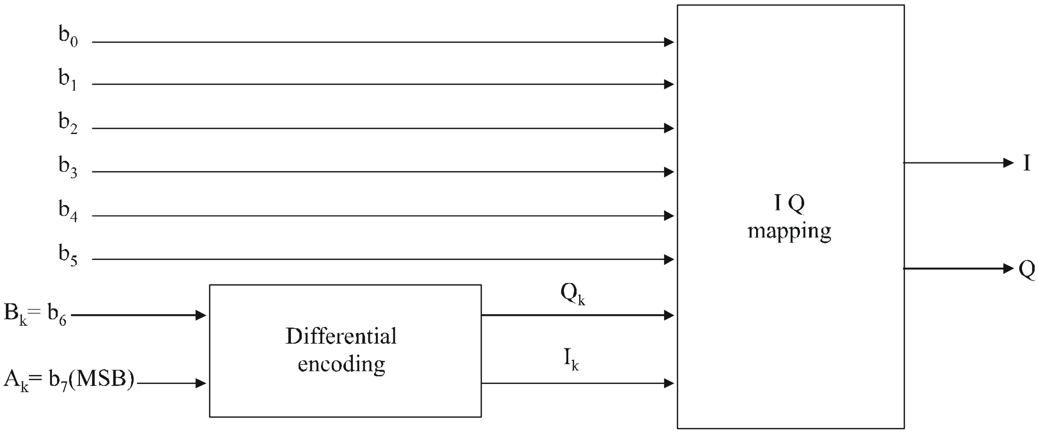

$$I_k = \overline{(A_k \oplus B_k)}.(A_k \oplus I_{k-1}) + (A_k \oplus B_k).(A_k \oplus Q_{k-1})$$

$$Q_k = \overline{(A_k \oplus B_k)}.(B_k \oplus Q_{k-1}) + (A_k \oplus B_k).(B_k \oplus I_{k-1})$$

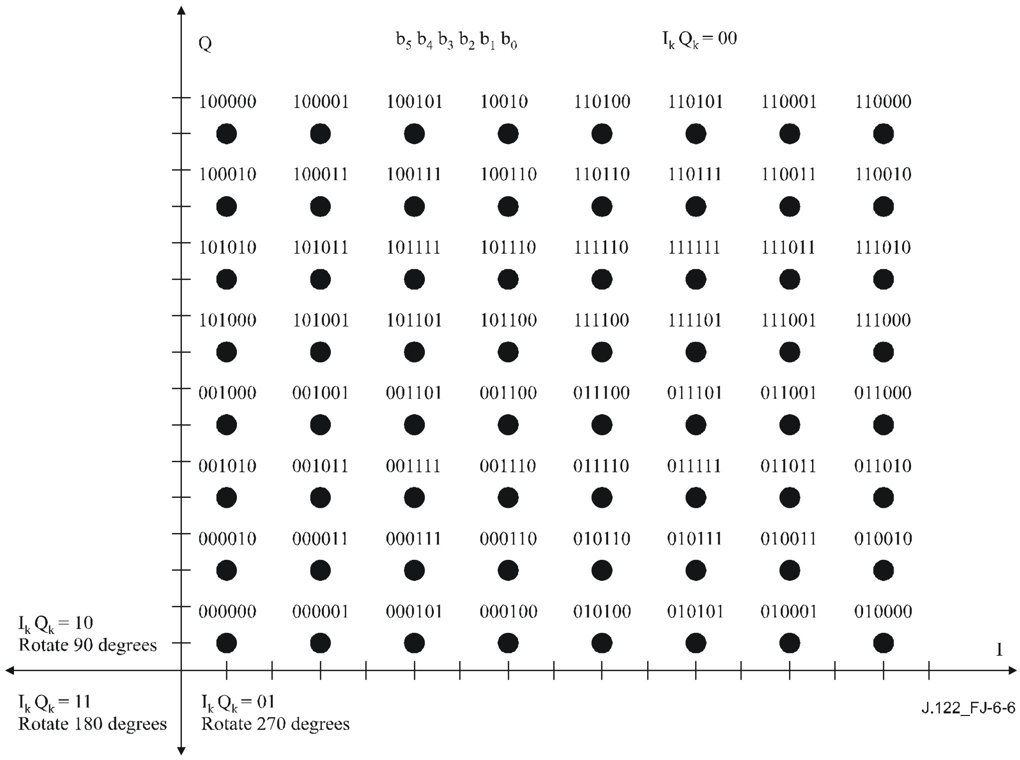

Figure J.6-6 – 256QAM signal constellation diagram and phase shift rule (I = 12)

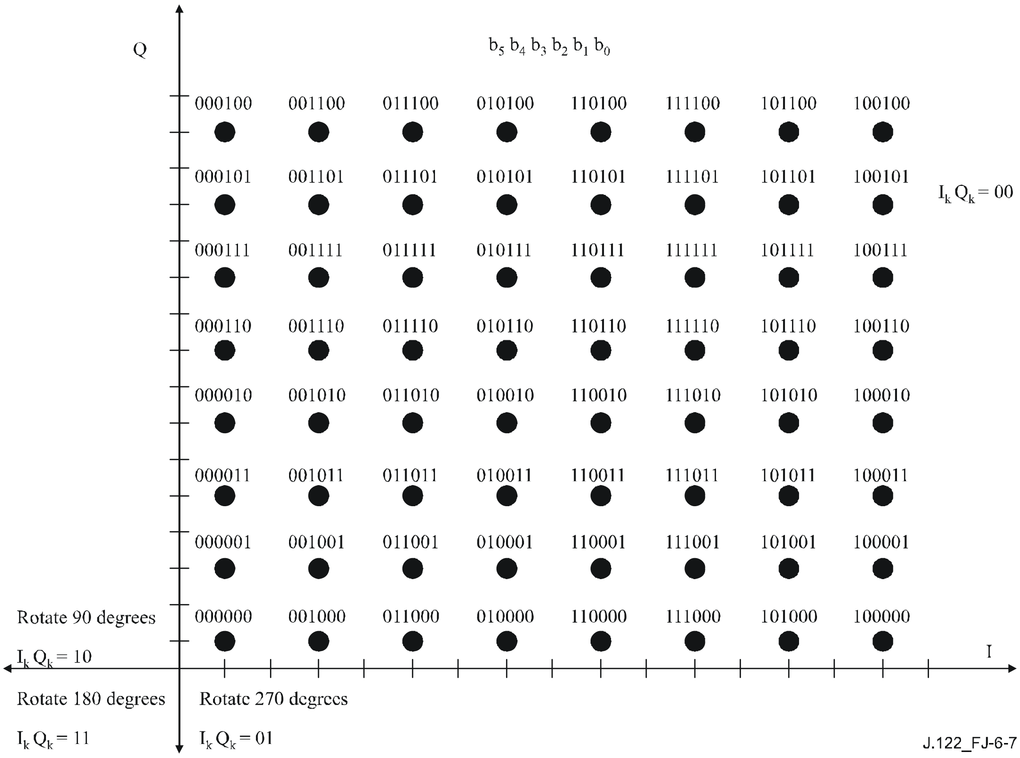

Figure J.6-7 – 256QAM signal constellation diagram and phase shift rule (I = 34 or 204)

J.6.3.4.3 Symbol rate, bandwidth and roll off

The symbol rate MUST be 5.274 Msymb/s. Bandwidth MUST be 6 MHz. The roll-off factor MUST be alpha = 0.13. Other parameters related to symbol rate, bandwidth and roll-off SHOULD be compliant with Annex C of [ITU-T J.83].

J.6.3.4.4 Frame structure

The frame structure SHOULD be compliant with Annex C of [ITU-T J.83].

J.6.3.4.5 Error correction

Error correction functionality SHOULD be considered for the noise environment in the cable television network. Code length and information byte length MUST be in accordance with Annex C of [ITU-T J.83].

The original Reed-Solomon code is defined as follows:

- Primitive polynomial: $P(x) = x^8 + x^4 + x^3 + x^2 + 1$
- Generator polynomial: $G(x) = (x + \alpha^0)(x + \alpha^1) - - - - - - (x + \alpha^{2t-1})$

where t is the error-correcting capability of a Reed-Solomon code and α is 02H and one of roots of equation P(x) = 0.

J.6.3.4.6 Randomization

A randomization function MUST be provided. The generator polynomial MUST be compliant with Annex C of [ITU-T J.83].

J.6.3.4.7 Transmission signal level

The transmission signal level at the CMTS output connector MUST be adjustable over the range of +100 to +120 dBμVrms.

J.6.3.4.8 Transmission spurious level

Transmission spurious level at the CMTS output connector MUST be less than –55 dBc over the range of 90 MHz to 770 MHz.

J.6.3.4.9 Channel frequency accuracy

Channel frequency accuracy MUST be within ± 20 kHz over a temperature range of 0 to 40° C.

J.6.3.4.10 Symbol rate accuracy

Symbol rate accuracy MUST be within ± 20 ppm over a temperature range of 0 to 40° C.

J.6.3.4.11 Impedance, return loss and connector

Impedance, return loss and connector at the CMTS output and CMTS input MUST meet the requirements shown in Table J.6-9.

Table J.6-9 – Impedance, return loss and connector type

	Impedance	Return loss	Connector type
CMTS Output	75 Ω	More than 14 dB 90-770 MHz	F-type, female
CMTS Input	75 Ω	More than 6 dB 10-55 MHz	F- type, female

J.6.3.5 Downstream electrical input to CM

J.6.3.5.1 Receiving signal level

CM MUST be able to operate at the level within a range of +45 to +75 dBμVrms for 64QAM, +51 to +81 dBμVrms for 256QAM at CM Input connector.

J.6.3.5.2 Impedance, return loss and connector

Impedance, return loss and connector at CM in/output MUST meet the requirements shown in Table J.6-10.

Table J.6-10 – Impedance, return loss and connector type

	Impedance	Return loss	Connector type
CM in/output	75 Ω	More than 6 dB 10-55 and 90-770 MHz	F-type, female

J.6.3.6 CM BER performance

The bit-error-rate performance of a CM MUST be as described in this clause. The requirements apply to the I = 12, J = 17 mode of interleaving.

J.6.3.6.1 64QAM

J.6.3.6.1.1 64QAM CM BER performance

Bit error rate MUST be less than 10^{-8} at CNR (Nyquist bandwidth) of 26 dBrms for 64QAM with error correction.

J.6.3.6.1.2 64QAM image rejection performance

Performance as described in clause J.6.3.6.1.1 MUST be met with an analogue or digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

J.6.3.6.1.3 64QAM adjacent channel performance

Performance as described in clause J.6.3.6.1.1 MUST be met with a digital signal at 0 dBc in the adjacent channels.

Performance as described in clause J.6.3.6.1.1 MUST be met with an analogue signal at +10 dBc in the adjacent channels.

Performance as described in clause J.6.3.6.1.1, with an additional 0.2-dB allowance, MUST be met with a digital signal at +10 dBc in the adjacent channels.

J.6.3.6.2 256QAM

J.6.3.6.2.1 256QAM CM BER performance

Bit error rate MUST be less than 10^{-8} at CNR (Nyquist bandwidth) of 33 dBrms for 256QAM with error correction.

J.6.3.6.2.2 256QAM image rejection performance

Performance as described in clause J.6.3.6.2.1 MUST be met with an analogue or digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

J.6.3.6.2.3 256QAM adjacent channel performance

Refer to clause J.4.4.

J.6.3.7 CMTS timestamp jitter

The CMTS timestamp jitter MUST be less than 500 ns peak-to-peak at the output of the downstream transmission convergence sublayer. This jitter is relative to an ideal downstream transmission convergence sublayer that transfers the MPEG packet data to the downstream physical media dependent sublayer with a perfectly continuous and smooth clock at the MPEG packet data rate. Downstream physical media dependent sublayer processing MUST NOT be considered in timestamp generation and transfer to the downstream physical media dependent sublayer.

Thus, any two timestamps Nl and N2 (N2 > N1) which were transferred to the downstream physical media dependent sublayer at times Tl and T2, respectively, must satisfy the following relationship:

$$| (N2 - Nl)/f_{CMTS} - (T2 - Tl) | < 500 \times 10^{-9}$$

In the equation, the value of (N2 – N1) is assumed to account for the effect of roll-over of the timebase counter, and Tl and T2 represent time in seconds. f_{CMTS} is the actual frequency of the CMTS master timebase and may include a fixed frequency offset from the nominal frequency of 9.216 MHz. This frequency offset is bounded by a requirement further below in this clause.

The jitter includes inaccuracy in timestamp value and the jitter in all clocks. The 500 ns allocated for jitter at the downstream transmission convergence sublayer output MUST be reduced by any jitter that is introduced by the downstream physical media dependent sublayer.

The CM is expected to meet the burst timing accuracy requirements in clause 6.2.19 when the timestamps contain this worst-case jitter.

NOTE – Jitter is the error (i.e., measured) relative to the CMTS master clock (the CMTS master clock is the 9.216-MHz clock used for generating the timestamps).

J.6.3.7.1 CMTS master clock jitter for asynchronous operation

The CMTS 9.216-MHz master clock MUST have frequency accuracy of ≤ ±5 ppm, drift rate: $\leq 10^{-8}$ per second, and edge jitter of: ≤10 ns peak-to-peak (±5 ns) over a temperature range of 0 to 40 degrees C up to ten years from date of manufacture[30]. (The drift rate and jitter requirements on the CMTS master clock implies that the duration of two adjacent segments of 9 216 000 cycles will be within 30 ns, due to 10-ns jitter on each segment's duration, and 10 ns due to frequency drift. Durations of other counter lengths also may be deduced: adjacent 9 216 000 segments, ≤21 ns; 9 216 000 length segments separated by one 92 160 000 cycle segment, ≤30 ns; adjacent 921 600 000 segments, ≤120 ns. The CMTS master clock MUST meet such test limits in 99% or more measurements.)

J.6.3.7.2 CMTS master clock jitter for synchronous operation

In addition to the requirements in clause 6.3.7.1, the 9.216-MHz CMTS master clock MUST meet the following double sideband phase noise requirements over the specified frequency ranges:

$< [-50 + 20 \times \log(f_{MC}/9.216)]$ dBc (i.e., < 0.05 ns RMS)	10 Hz to 100 Hz
$< [-58 + 20 \times \log(f_{MC}/9.216)]$ dBc (i.e., < 0.02 ns RMS)	100 Hz to 1 kHz
$< [-50 + 20 \times \log(f_{MC}/9.216)]$ dBc (i.e., < 0.05 ns RMS)	1 to 10 kHz
$< [-50 + 20 \times \log(f_{MC}/9.216)]$ dBc (i.e., < 0.05 ns RMS)	10 kHz to $f_{MC}/2$

where f_{MC} is the frequency of the measured master clock in MHz. The value of f_{MC} MUST be either an integral multiple or divisor of 9.216 MHz. For example, if a 18.432-MHz oscillator is used as the master clock frequency source, and there is no explicit 9.216-MHz clock to test, the 18.432-MHz clock may be used with f_{MC} equal to 18.432 in the above expressions.

J.6.3.7.3 CMTS master clock frequency drift for synchronous operation

Refer to clause 6.3.7.3.

J.6.3.8 CMTS clock generation

The CMTS has the following three options related to the synchronization of the CMTS master clock and the downstream symbol clock:

1) not locked;
2) downstream symbol clock locked to cmts master clock;
3) CMTS master clock locked to downstream symbol clock.

For S-CDMA operation, the master clock and the downstream symbol clock MUST be locked using either option 2 or 3.

Let f_b' represent the rate of the downstream symbol clock which is locked to the CMTS master clock, and let f_m' represent the rate of the CMTS master clock locked to the downstream symbol clock. Let f_b' represent the nominal specified downstream symbol rate, and let f_m' represent the nominal CMTS master clock rate (9.216 MHZ).

With the downstream symbol clock locked to the CMTS master clock, the following equation MUST hold:

$$f_b' = f_m \times M/N$$

[30] This Recommendation MAY also be met by synchronizing the CMTS master clock oscillator to an external frequency reference source. If this approach is used, the internal CMTS master clock MUST have frequency accuracy of ±20 ppm over a temperature range of 0 to 40 degrees C up to ten years from date of manufacture when no frequency reference source is connected. The drift rate and edge jitter MUST be as specified above.

With the CMTS master clock locked to the downstream symbol clock, the following equation MUST hold:

$$f_m' = f_b \times N/M$$

M and N MUST be unsigned integer values each representable in 16 bits (these are specified in the channel TLV parameters of the UCD). When the downstream symbol clock and the CMTS master clock are not locked together (sync mode = 0), the values of M and N are not valid and are ignored by the CM.

The values of M and N MUST result in a value of f_b' or f_m', which is not more than ±l ppm from its specified nominal value. Table J.6-11 lists the downstream modes of operation, their associated nominal symbol rates, f_b, example values for M and N, the resulting synchronized clock rates, and their offsets from their nominal values.

Table J.6-11 – Downstream symbol rates and example parameters for synchronization with the CMTS master clock

Downstream mode	Nominal specified symbol rate, f_b (MHz)	M/N	CMTS master clock rate, f_m'(MHz)	Downstream symbol rate, f_b'(MHz)	Offset from nominal
64QAM and 256QAM	5.274	293/512	9.216	5.274	0 ppm

J.6.3.9 Downstream symbol clock jitter for synchronous operation

The downstream symbol clock MUST meet the following double sideband phase noise requirements over the following frequency ranges:

$< [-53 + 20 \times \log(f_{DS}/5.274)]$ dBc (i.e., < 0.07 ns rms)	10 to 100 Hz
$< [-53 + 20 \times \log(f_{DS}/5.274)]$ dBc (i.e., < 0.07 ns rms)	100 Hz to 1 kHz
$< [-53 + 20 \times \log(f_{DS}/5.274)]$ dBc (i.e., < 0.07 ns rms)	1 to 10 kHz
$< [-36 + 20 \times \log(f_{DS}/5.274)]$ dBc (i.e., < 0.5 ns rms)	10 to 100 kHz
$< [-30 + 20 \times \log(f_{DS}/5.274)]$ dBc (i.e., < 1 ns rms)	100 kHz to $(f_{DS}/2)$

where f_{DS} is the frequency of the measured clock in MHz. The value of f_{DS} MUST be an integral multiple or divisor of the downstream symbol clock. For example, an f_{DS} = 21.096-MHz clock may be measured if there is no explicit 5.274-MHz clock available.

The CMTS MUST provide a test mode in which:

- the downstream QAM symbol sequence is replaced with an alternating binary sequence (1, –1, 1, –1, 1, –1, ...) at nominal amplitude, on both I and Q;
- the CMTS generates the downstream symbol clock from the 9.216-MHz reference clock as in normal synchronous operation.

If an explicit downstream symbol clock that is capable of meeting the above phase noise requirements is available (e.g., a smooth clock without clock domain jitter), this test mode is not required.

J.6.3.10 CMTS downstream symbol clock drift for synchronous operation

Refer to clause 6.3.10.

J.7 Downstream transmission convergence sublayer

J.7.1 Introduction

In order to improve demodulation robustness, facilitate common receiving hardware for both video and data, and provide an opportunity for the possible future multiplexing of video and data over the PMD sublayer bitstream defined in clause 6, a sublayer is interposed between the downstream PMD sublayer and the data-over-cable MAC sublayer.

The downstream bitstream is defined as a continuous series of 188-byte MPEG [ITU-T H.222.0] packets. These packets consist of a 4-byte header followed by 184 bytes of payload. The header identifies the payload as belonging to the data-over-cable MAC. Other values of the header may indicate other payloads. the mixture of MAC payloads and those of other services is optional and is controlled by the CMTS.

Figure J.7-1 illustrates the interleaving of data-over-cable (DOC) MAC bytes with other digital information (digital video in the example shown).

Header = DOC	DOC MAC payload
Header = video	Digital video payload
Header = video	Digital video payload
Header = DOC	DOC MAC payload
Header = video	Digital video payload
Header = DOC	DOC MAC payload
Header = video	Digital video payload
Header = video	Digital video payload
Header = video	Digital video payload

Figure J.7-1 – Example of interleaving MPEG packets in downstream

J.7.2 MPEG packet format

Refer to clause 7.2.

J.7.3 MPEG header for DOCS data-over-cable

Refer to clause 7.3.

J.7.4 MPEG payload for DOCS data-over-cable

Refer to clause 7.4.

J.7.5 Interaction with the MAC sublayer

Refer to clause 7.5.

J.7.6 Interaction with the physical layer

The MPEG-2 packet stream MUST be encoded according to Annex C of [ITU-T J.83].

J.7.7 MPEG header synchronization and recovery

Refer to clause 7.7.

J.8 Media access control specification

Refer to clause 8, except for the following changes:

- Time tick 6.25 µs MUST change to 6.94 µs.
- Master clock 10.24 MHz MUST change to 9.216 MHz.
- Base rate 160 KHz MUST change to 144 KHz.

J.8.3.20.1.2.3 Downstream symbol rate

This TLV specifies the symbol rate that is used on the new downstream channel.

Subtype	Length	Value
2.3	1	0: 5.056941 Msymb/s 1: 5.360537 Msymb/s 2: 6.952 Msymb/s 3: 5.274 Msymb/s 4-255: reserved

The CMTS SHOULD include this sub-TLV. The CM SHOULD observe this sub-TLV.

J.9 Media access control protocol operation

J.9.3.4 Timing units and relationships

The SYNC message conveys a time reference with a resolution of 6.94/64 microseconds (9.216 MHz) to allow the CM to track the CMTS clock with a small phase offset. Since this timing reference is decoupled from particular upstream channel characteristics, a single SYNC time reference may be used for all upstream channels associated with the downstream channel.

The bandwidth allocation MAP uses time units of "mini-slots". A mini-slot represents the time needed for transmission of a fixed number of symbols. For some modulations (e.g., QPSK), an integer number of bytes can be transmitted in a mini-slot. For these channels, the mini-slot is expected to represent 16 byte-times, although other values could be chosen.

A "mini-slot" is the unit of granularity for upstream transmission opportunities; there is no implication that any PDU can actually be transmitted in a single mini-slot.

J.9.3.4.1.1 Mini-slot capacity

On TDMA channels, the size of the mini-slot, expressed as a multiple of the SYNC time reference, is carried in the upstream channel descriptor. The example in Table J.9-1 relates mini-slots to the SYNC time ticks (assuming QPSK modulation):

Table J.9-1 – Example relating mini-slots to time ticks

Parameter	Example value
Time tick	6.94 microseconds
Bytes per mini-slot	16 (nominal, when using QPSK modulation)
Symbols/byte	4 (assuming QPSK)
Symbols/second	2 304 000
Mini-slots/second	36'000
Microseconds/mini-slot	27.7
Ticks/mini-slot	4

Note that the symbols/byte is a characteristic of an individual burst transmission, not of the channel. A mini-slot in this instance could represent a minimum of 16 or a maximum of 48 bytes, depending on the modulation choice.

In a channel allocated exclusively to DOCS 2.0 TDMA modems, the mini-slot size field of the UCD MAY take on the value 0, in which case the mini-slot size is 1 timebase tick. If a channel is to be accessible to both DOCS 1.x and 2.0 TDMA cable modems, the UCD MUST follow the DOCS 1.x requirements for timing units and relationships.

J.9.3.4.2.1 Mini-slot capacity

On S-CDMA channels, the size of the mini-slot is dependent on the modulation rate, the codes per mini-slot, and the spreading intervals per frame, which are all carried in the upstream channel descriptor. The timing units and relationships for S-CDMA are covered in detail in clause 6.2.11, "S-CDMA framer and interleaver". An example of the timing relationships (assuming 64QAM modulation) is shown in Table J.9-2:

Table J.9-2 – Example of mini-slot capacity in S-CDMA mode

Parameter	Example value
Spreading intervals per frame	10
Active code length	128
Codes per mini-slot	4
Mini-slots per frame	32
Symbols per mini-slot	40
Bytes per mini-slot	30 (nominal, when using 64QAM modulation)
Bits/symbol	6 (assuming 64QAM)
Symbols/second	4'608'000
Mini-slots/second	115'200
Microseconds/mini-slot	8.6805

J.9.5 Data link encryption support

Refer to clause 9.5

J.10 Quality of service and fragmentation

Refer to clause 10.

J.11 Cable modem – CMTS interaction

J.11.2.10 Baseline privacy initialization

Following registration, if the CM is provisioned to run baseline privacy, the CM MUST initialize baseline privacy operations, as described in [DOCS8]. A CM is provisioned to run baseline privacy if the privacy enable TLV (see clause C.1.1.16) in the DOCS 1.1-style configuration file is explicitly/implicitly set to enable or the baseline privacy configuration setting (see clause C.3.2) is contained in the DOCS 1.0-style configuration file as specified in clauses A.1.1 and C.2 of the BPI+ spec [DOCS8]. Note that the secure software download is required regardless of whether the CM is provisioned to run baseline privacy or not, as specified in Appendix D of the BPI+ spec [DOCS8].

In order to ease operations and reduce cost, the manufacturer CA certificate of the CM MAY be self-signed.

J.C Common radio frequency interface encodings

J.C.2.2.6.11 Unsolicited grant time reference

For unsolicited grant service and unsolicited grant service with activity detection, the value of this parameter specifies a reference time t_0 from which can be derived the desired transmission times

$t_i = t_0 + i \times$ interval, where interval is the nominal grant interval (refer to clause C.2.2.6.7). This parameter is applicable only for messages transmitted from the CMTS to the CM, and only when a UGS or UGS-AD service flow is being made active. In such cases, this is a mandatory parameter.

Type	Length	Value	Valid range
24.24	4	CMTS timestamp	0-4'294'967'295

The timestamp specified in this parameter represents a count state of the CMTS 9.216-MHz master clock. Since a UGS or UGS-AD service flow is always activated before transmission of this parameter to the modem, the reference time t_0 is to be interpreted by the modem as the ideal time of the next grant only if t_0 follows the current time. If t_0 precedes the current time, the modem can calculate the offset from the current time to the ideal time of the next grant according to:

$$\text{interval} - (((\text{current time} - t_0)/9.216) \text{ modulus interval})$$

where interval is in units of microseconds, and current time and t_0 are in 9.216-MHz units.

J.E The data-over-cable spanning tree protocol

J.E.4.1 Path cost

In [b-ISO/IEC 10038], the following formula is used:

Path_Cost = 1000/Attached_LAN_speed_in_Mb/s

For CMs, this formula is adapted as:

Path_Cost = 1000/(Upstream_modulation_rate × bits_per_Symbol_for_long_data_grant)

That is, the modulation type (QPSK or 16QAM) for the long data grant IUC is multiplied by the raw modulation rate to determine the nominal path cost. Table J.E-1 provides the derived values.

Table J.E-1 – CM path cost

Modulation rate	Default path cost	
kHz	**QPSK**	**16QAM**
144	3472	1736
288	1736	868
596	868	434
1152	434	217
2304	217	108

For CMTSs, this formula is:

Path_Cost = 1000/(Downstream_symbol_rate × bits_per_symbol)

Appendix I

MAC service definition

(This appendix does not form an integral part of this Recommendation)

This appendix is informative. In case of conflict between this appendix and any normative section of this Recommendation, the normative section takes precedence.

I.1 MAC service overview

The DOCS MAC provides a protocol service interface to upper-layer services. Examples of upper-layer services include a DOCS bridge, embedded applications (e.g., IPCablecom/VoIP), a host interface (e.g., NIC adapter with NDIS driver), and layer three routers (e.g., IP router).

The MAC service interface defines the functional layering between the upper-layer service and the MAC. As such it defines the functionality of the MAC which is provided by the underlying MAC protocols. This interface is a protocol interface, not a specific implementation interface.

The following data services are provided by the MAC service interface:

- A MAC service exists for classifying and transmitting packets to MAC service flows.
- A MAC service exists for receiving packets from MAC service flows. Packets may be received with suppressed headers.
- A MAC service exists for transmitting and receiving packets with suppressed headers. The headers of transmitted packets are suppressed based upon matching classifier rules. The headers of received suppressed packets are regenerated based upon a packet header index negotiated between the CM and CMTS.
- A MAC service exists for synchronization of grant timing between the MAC and the upper-layer service. This clock synchronization is required for applications such as embedded IPCablecom VoIP clients in which the packetization period needs to be synchronized with the arrival of scheduled grants from the CMTS.
- A MAC service exists for synchronization of the upper-layer clock with the CMTS controlled master clock.

It should be noted that a firewall and policy based filtering service may be inserted between the MAC layer and the upper-layer service, but such a service is not modelled in this MAC service definition.

The following control services are provided by the MAC service interface:

- A MAC service exists for the upper layer to learn of the existence of provisioned service flows and QoS traffic parameter settings at registration time.
- A MAC service exists for the upper layer to create service flows. Using this service, the upper layer initiates the admitted/activated QoS parameter sets, classifier rules and packet suppression headers for the service flow.
- A MAC service exists for the upper layer to delete service flows.
- A MAC service exists for the upper layer to change service flows. Using this service, the upper layer modifies the admitted/activated QoS parameter sets, classifier rules and packet suppression headers.
- A MAC service exists for controlling the classification of and transmission of PDUs with suppressed headers. At most, a single suppressed header is defined for a single classification rule. The upper-layer service is responsible for defining both the definition of suppressed headers (including wild-card do-not-suppress fields) and the unique

classification rule that discriminates each header. In addition to the classification rule, the MAC service can perform a full match of all remaining header bytes to prevent generation of false headers if so configured by the upper-layer service.

- A MAC service exists for controlling two-phase control of QoS traffic resources. Two-phase activation is controlled by the upper-layer service to provide both admitted QoS parameters and active QoS parameters within the appropriate service request. Upon receipt of an affirmative indication, the upper-layer service knows that the admitted QoS parameter set has been reserved by the CMTS, and that the activated QoS parameter set has been activated by the CMTS. Barring catastrophic failure (such as resizing of the bandwidth of the upstream PHY), admitted resources will be guaranteed to be available for activation, and active resources will be guaranteed to be available for use in packet transmission.

A control function for locating an unused service flow and binding it, or a specific identified service flow, to a specific upper-layer service may also exist. The details of such a function are not specified and are implementation-dependent.

Other control functions may exist at the MAC service interface, such as functions for querying the status of active service flows and packet classification tables, or functions from the MAC service to the upper-layer service to enable the upper-layer service to authorize service flows requested by the peer MAC layer service, but those functions are not modelled in this MAC service definition.

Other MAC services that are not service flow-related also exist, such as functions for controlling the MAC service MAC address and SAID multicast filtering functions, but those functions are not modelled in this MAC service definition.

I.1.1 MAC service parameters

The MAC service utilizes the following parameters. For a full description of the parameters, consult the theory of operation and other relevant clauses within the body of the RFI Recommendation.

- Service flow QoS traffic parameters

  MAC activate-service-flow and change-service-flow primitives allow common, upstream and downstream QoS traffic parameters to be provided. When such parameters are provided, they override whatever values were configured for those parameters at provisioning time or at the time the service flow was created by the upper-layer service.

- Active/admitted QoS traffic parameters

  If two-phase service flow activation is being used, then two complete sets of QoS traffic parameters are controlled. The admitted QoS parameters state the requirements for reservation of resources to be authorized by the CMTS. The activated QoS parameters state the requirements for activation of resources to be authorized by the CMTS. Admitted QoS parameters may be activated at a future time by the upper-layer service. Activated QoS parameters may be used immediately by the upper-layer service.

- Service flow classification filter rules

  Zero or more classification filter rules may be provided for each service flow that is controlled by the upper-layer service. Classifiers are identified with a classifier identifier.

- Service flow PHS suppressed headers

  Zero or more PHS suppressed header strings with their associated verification control and mask variables may be defined for each service flow. When such headers are defined, they are associated 1-to-1 with specific classification rules. In order to regenerate packets with suppressed headers, a payload header suppression index is negotiated between the CM and CMTS.

I.2 MAC data service interface

MAC services are defined for transmission and reception of data to and from service flows. Typically, an upper-layer service will utilize service flows for mapping of various classes of traffic to different service flows. Mappings to service flows may be defined for low priority traffic, high priority traffic and multiple special traffic classes such as constant bit-rate traffic that is scheduled by periodic grants from the CMTS at the MAC layer.

The following specific data service interfaces are provided by the MAC service to the upper-layer service. These represent an abstraction of the service provided and do not imply a particular implementation:

- MAC_DATA.request;
- MAC_DATA.indicate;
- MAC_GRANT_SYNCHRONIZE.indicate;
- MAC_CMTS_MASTER_CLOCK_SYNCHRONIZE.indicate;
- MAC_DATA.request.

Issued by the upper-layer service to request classification and transmission of an IEEE 802.3 or DIX-formatted PDU to the RF.

Parameters

- PDU – IEEE 802.3 or DIX-encoded PDU including all layer-2 header fields and optional FCS. PDU is the only mandatory parameter.
- Padding – Used when the PDU is less than 60 bytes and it is desired to maintain ISO/IEC 8802-3 transparency.
- ServiceFlowID – If included, the MAC service circumvents the packet classification function and maps the packet to the specific service flow indicated by the ServiceFlowID value.
- ServiceClassName, RulePriority – If included, this tuple identifies the service class name of an active service flow to which the packet is to be mapped so long as a classifier does not exist at a rule priority higher than the rule priority supplied.

Expanded service description

Transmit a PDU from upper-layer service to MAC/PHY. The only mandatory parameter is PDU. PDU contains all layer-2 headers, layer-3 headers, data and (optional) layer-2 checksum.

If PDU is the only parameter, the packet is subjected to the MAC packet classification filtering function in order to determine how the packet is mapped to a specific service flow. The results of the packet classification operation determine on which service flow the packet is to be transmitted and whether or not the packet should be transmitted with suppressed headers.

If the parameter ServiceFlowID is supplied, the packet can be directed to the specifically-identified service flow.

If the parameter tuple ServiceClassName, RulePriority is supplied, the packet is directed to the first active service flow that matches the service class name so long as a classifier does not exist at a rule priority higher than the rule priority supplied. This service is used by upper-layer policy enforcers to allow zero or more dynamic rules to be matched for selected traffic (e.g., voice) while all other traffic is forced to a service flow within the named service flow class. If no active service flow with the service class name exists, then the service performs normal packet classification.

In all cases, if no classifier match is found, or if none of the combinations of parameters maps to a specific service flow, the packet will be directed to the primary service flow.

The following pseudo-code describes the intended operation of the MAC_DATA.request service interface.

I.2.1 MAC_DATA.request

```
   PDU
   [ServiceFlowID]
   [ServiceClassName, RulePriority]
```

FIND_FIRST_SERVICE_FLOW_ID (ServiceClassName) returns ServiceFlowID of first service flow whose ServiceClassName equals the parameter of the procedure or NULL if no matching service flow is found.

SEARCH_CLASSIFIER_TABLE (PriorityRange) searches all rules within the specified priority range and returns either the ServiceFlowID associated with the rule or NULL if no classifier rule is found.

```
TxServiceFlowID = NULL
IF (ServiceFlowID DEFINED)
                TxServiceFlowID = MAC_DATA.ServiceFlowID
ELSEIF (ServiceClassName DEFINED and RulePriority DEFINED)
               TxServiceFlowID = FIND_FIRST_SERVICE_FLOW_ID (ServiceClassName)
               SearchID = SEARCH_CLASSIFIER_TABLE (All Priority Levels)
               IF (SearchID not NULL and ClassifierRule.Priority >=
MAC_DATA.RulePriority)
                                 TxServiceFlowID = SearchID
ELSE [PDU only]
              TxServiceFlow = SEARCH_CLASSIFIER_TABLE (All Priority Levels)
IF (TxServiceFlowID = NULL)
        TRANSMIT_PDU (PrimaryServiceFlowID)
ELSE
         TRANSMIT_PDU (TxServiceFlowID)
```

I.2.2 MAC_DATA.indicate

Issued by the MAC to indicate reception of an IEEE 802.3 or DIX PDU for the upper-layer service from the RF.

Parameters

- PDU – IEEE 802.3 or DIX-encoded PDU including all layer-2 header fields and FCS.

I.2.3 MAC_GRANT_SYNCHRONIZE.indicate

Issued by the MAC service to the upper-layer service to indicate the timing of grant arrivals from the CTMS. It is not stated how the upper layer derives the latency, if any, between the reception of the indication and the actual arrival of grants (within the bounds of permitted grant jitter) from the CMTS. It should be noted that in UGS applications it is expected that the MAC layer service will increase the grant rate or decrease the grant rate based upon the number of grants per interval QoS traffic parameter. It should also be noted that as the number of grants per interval is increased or decreased, the timing of grant arrivals will change also. It should also be noted that when synchronization is achieved with the CMTS downstream master clock, this indication may only be required once per active service flow. No implication is given as to how this function is implemented.

Parameters

- ServiceFlowID – Unique identifier value for the specific active service flow receiving grants.

I.2.4 MAC_CMTS_MASTER_CLOCK_SYNCHRONIZE.indicate

Issued by the MAC service to the upper-layer service to indicate the timing of the CMTS master clock. No implication is given as to how often or how many times this indication is delivered by the MAC service to the upper-layer service. No implication is given as to how this function is implemented.

Parameters

- No parameters specified.

I.3 MAC control service interface

A collection of MAC services are defined for control of MAC service flows and classifiers. It should be noted that an upper-layer service may use these services to provide an upper-layer traffic construct such as "connections" or "subflows" or "micro-flows". However, except for the ability to modify individual classifiers, no explicit semantics is defined for such upper layer models. Thus, control of MAC service flow QoS parameters is specified in the aggregate.

The following specific control service interface functions are provided by the MAC service to the upper-layer service. These represent an abstraction of the service provided and do not imply a particular implementation:

– MAC_REGISTRATION_RESPONSE.indicate;

– MAC_CREATE_SERVICE_FLOW.request/response/indicate;

– MAC_DELETE_SERVICE_FLOW.request/response/indicate;

– MAC_CHANGE_SERVICE_FLOW.request/response/indicate.

I.3.1 MAC_REGISTRATION_RESPONSE.indicate

Issued by the DOSCIS MAC to the upper-layer service to indicate the complete set of service flows and service flow QoS traffic parameters that have been provisioned and authorized by the registration phase of the MAC. Subsequent changes to service flow activation state or addition and deletion of service flows are communicated to the upper-layer service with indications from the other MAC control services.

Parameters

- Registration TLVs – Any and all TLVs that are needed for service flow and service flow parameter definition including provisioned QoS parameters. See the normative body of this Recommendation for more details.

I.3.2 MAC_CREATE_SERVICE_FLOW.request

Issued by the upper-layer service to the MAC to request the creation of a new service flow within the MAC service. This primitive is not issued for service flows that are configured and registered, but rather for dynamically-created service flows. This primitive may also define classifiers for the service flow and supply admitted and activated QoS parameters. This function invokes DSA signalling.

Parameters

- ServiceFlowID – Unique ID value for the specific service flow being created.
- ServiceClassName – service flow class name for the service flow being created.
- Admitted QoS parameters – Zero or more upstream, downstream and common traffic parameters for the service flow.
- Activated QoS parameters – Zero or more upstream, downstream and common traffic parameters for the service flow.

- Service flow payload header suppression rules – Zero or more PHS rules for each service flow that is controlled by the upper-layer service.
- Service flow classification filter rules – Zero or more classification filter rules for each service flow that is controlled by the upper-layer service. Classifiers are identified with a classifier identifier.

I.3.3 MAC_CREATE_SERVICE_FLOW.response

Issued by the MAC service to the upper-layer service to indicate the success or failure of the request to create a service flow.

Parameters

- ServiceFlowID – Unique identifier value for the specific service flow being created.
- ResponseCode – Success or failure code.

I.3.4 MAC_CREATE_SERVICE_FLOW.indicate

Issued by the MAC service to notify the upper-layer service of the creation of a new service flow within the MAC service. This primitive is not issued for service flows that have been administratively pre-configured, but rather for dynamically defined service flows. In this Recommendation, this notification is advisory only.

Parameters

- ServiceFlowID – Unique ID value for the specific service flow being created.
- ServiceClassName – service flow class name for the service flow being created.
- Admitted QoS parameters – Zero or more upstream, downstream and common traffic parameters for the service flow.
- Activated QoS parameters – Zero or more upstream, downstream and common traffic parameters for the service flow.
- Service flow payload header suppression rules – Zero or more PHS rules for each service flow that is controlled by the upper-layer service.
- Service flow classification filter rules – Zero or more classification filter rules for each service flow that is controlled by the upper-layer service. Classifiers are identified with a classifier identifier.

I.3.5 MAC_DELETE_SERVICE_FLOW.request

Issued by the upper-layer service to the MAC to request the deletion of a service flow and all QoS parameters including all associated classifiers and PHS rules. This function invokes DSD signalling.

Parameters

- ServiceFlowID(s) – Unique identifier value(s) for the deleted service flow(s).

I.3.6 MAC_DELETE_SERVICE_FLOW.response

Issued by the MAC service to the upper-layer service to indicate the success or failure of the request to delete a service flow.

Parameters

- ResponseCode – Success or failure code.

I.3.7 MAC_DELETE_SERVICE_FLOW.indicate

Issued by the MAC service to notify the upper-layer service of deletion of a service flow within the MAC service.

Parameters

- ServiceFlowID(s) – Unique identifier value(s) for the deleted service flow(s).

I.3.8 MAC_CHANGE_SERVICE_FLOW.request

Issued by the upper-layer service to the MAC to request modifications to a specific created and acquired service flow. This function is able to define both the complete set of classifiers and incremental changes to classifiers (add/remove). This function defines the complete set of admitted and active QoS parameters for a service flow. This function invokes DSC MAC-layer signalling.

Parameters

- ServiceFlowID – Unique identifier value for the specific service flow being modified.
- Zero or more packet classification rules with add/remove semantics and LLC, IP and 802.1PQ parameters.
- Admitted QoS parameters – Zero or more upstream, downstream and common traffic parameters for the service flow.
- Activated QoS parameters – Zero or more upstream, downstream and common traffic parameters for the service flow.
- Service flow payload header suppression rules – Zero or more PHS rules for each service flow that is controlled by the upper-layer service.

I.3.9 MAC_CHANGE_SERVICE_FLOW.response

Issued by the MAC service to the upper-layer service to indicate the success or failure of the request to change a service flow.

Parameters

- ServiceFlowID – Unique identifier value for the specific service flow being released.
- ResponseCode – Success or failure code.

I.3.10 MAC_CHANGE_SERVICE_FLOW.indicate

Issued by the DOSCIS MAC service to notify upper-layer service of a request to change a service flow. In this Recommendation, the notification is advisory only and no confirmation is required before the service flow is changed. Change-service-flow indications are generated based upon DSC signalling. DSC signalling can be originated based upon change-service-flow events between the peer upper-layer service and its MAC service, or based upon network resource failures such as a resizing of the total available bandwidth at the PHY layer. How the upper-layer service reacts to forced reductions in admitted or reserved QoS traffic parameters is not specified.

Parameters

- ServiceFlowID – Unique identifier for the service flow being activated.
- Packet classification rules with LLC, IP and 802.1PQ parameters, and with zero or more PHS_CLASSIFIER_IDENTIFIERs.
- Admitted QoS parameters – Zero or more upstream, downstream and common traffic parameters for the service flow.
- Activated QoS parameters – Zero or more upstream, downstream and common traffic parameters for the service flow.
- Service flow payload header suppression rules – Zero or more PHS rules for each service flow that is controlled by the upper-layer service.

I.4 MAC service usage scenarios

Upper-layer entities utilize the services provided by the MAC in order to control service flows and in order to send and receive data packets. The partition of function between the upper-layer service and the MAC service is demonstrated by the following scenarios.

I.4.1 Transmission of PDUs from upper-layer service to MAC DATA service

- Upper-layer service transmits PDUs via the MAC_DATA service.
- MAC_DATA service classifies transmitted PDUs using the classification table, and transmits the PDUs on the appropriate service flow. The classification function may also cause the packet header to be suppressed according to a header suppression template stored with the classification rule. It is possible for the upper-layer service to circumvent this classification function.
- MAC_DATA service enforces all service flow-based QoS traffic shaping parameters.
- MAC_DATA service transmits PDUs on DOCS RF as scheduled by the MAC layer.

I.4.2 Reception of PDUs to upper-layer service from MAC DATA service

PDUs are received from the DOCS RF.

If a PDU is sent with a suppressed header, the header is regenerated before the packet is subjected to further processing.

In the CMTS, the MAC_DATA service classifies the PDU's ingress from the RF using the classification table and then polices the QoS traffic shaping and validates addressing as performed by the CM. In the CM, no per-packet service flow classification is required for traffic ingress from the RF.

Upper-layer service receives PDUs from the MAC_DATA.indicate service.

I.4.3 Sample sequence of MAC control and MAC data services

A possible CM-oriented sequence of MAC service functions for creating, acquiring, modifying and then using a specific service flow is as follows:

- MAC_REGISTER_RESPONSE.indicate

  Learn of any provisioned service flows and their provisioned QoS traffic parameters.
- MAC_CREATE_SERVICE_FLOW.request/response

  Create new service flow. This service interface is utilized if the service flow was learned as not provisioned by the MAC_REGISTER_RESPONSE service interface. Creation of a service flow invokes DSA signalling.
- MAC_CHANGE_SERVICE_FLOW.request/response

  Define admitted and activated QoS parameter sets, classifiers and packet suppression headers. Change of a service flow invokes DSC signalling.
- MAC_DATA.request

  Send PDUs to MAC service for classification and transmission.
- MAC_DATA.indication

  Receive PDUs from MAC service.
- MAC_DELETE_SERVICE_FLOW.request/response

  Delete service flow. Would likely be invoked only for dynamically-created service flows, not provisioned service flows. Deletion of a service flow uses DSD signalling.

Appendix II

Example preamble sequence

(This appendix does not form an integral part of this Recommendation)

II.1 Introduction

A programmable preamble superstring, up to 1536 bits long, is part of the channel-wide profile or attributes, common to the all-burst profiles on the channel (see Table 8-18), but with each burst profile able to specify the start location within this sequence of bits and the length of the preamble (see Table 8-19). The first bit of the preamble pattern is designated by the preamble value offset as described in Table 8-19. The first bit of the preamble pattern is the first bit into the symbol mapper (Figure 6-2, "TDMA upstream transmission processing" and Figure 6-3, "S-CDMA upstream transmission processing"), and is the first symbol of the burst (see clause 6.2.13, "symbol mapping"). As an example, per Table 8-19, 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. An example 1536-bit-long preamble superstring is given in clause II.2.

II.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 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 III

Multiple upstream channels

(This appendix does not form an integral part of this Recommendation)

This appendix is informative. In case of conflict between this appendix and any normative part of this Recommendation, the normative part takes precedence.

Clause 7.2, "MPEG packet format", describes support for multiple upstream and multiple downstream channels within a DOCS domain. The permutations that a CM may see on the cable segment it is attached to include:

- single downstream and single upstream per cable segment;
- single downstream and multiple upstreams per cable segment;
- multiple downstreams and single upstream per cable segment;
- multiple downstreams and multiple upstreams per cable segment.

A typical application that will require one upstream and one downstream per CM is web browsing. Web browsing tends to have asymmetrical bandwidth requirements that match closely to the asymmetrical bandwidth of DOCS.

A typical application that will require access to one of multiple upstreams per CM is IP telephony. IP telephony tends to have symmetrical bandwidth requirements. If there is a large concentration of CMs in a geographical area all served by the same fibre node, more than one upstream may be required in order to provide sufficient bandwidth and prevent call blocking.

A typical application that will require access to one of multiple downstreams per CM is IP streaming video. IP streaming video tends to have extremely large downstream bandwidth requirements. If there is a large concentration of CMs in a geographical area, all served by the same fibre node, more than one downstream may be required in order to provide sufficient bandwidth and to deliver multiple IP video streams to multiple CMs.

A typical application that will require multiple downstreams and multiple upstreams is when the above applications are combined, and it is more economical to have multiple channels than it is to physically subdivide the HFC network.

The role of the CM in these scenarios would be to be able to move between multiple upstreams and between multiple downstreams. The role of the CMTS would be to manage the traffic load to all attached CMs, and balance the traffic between the multiple upstreams and downstreams by dynamically moving the CMs based upon their resource needs and the resources available.

This appendix looks at the implementation considerations for these cases. Specifically, the first and last application are profiled. These examples are meant to illustrate one topology and one implementation of that topology.

III.1 Single downstream and single upstream per cable segment

This clause presents an example of a single downstream channel and four upstream channels. In Figure III.1, the four upstream channels are on separate fibres serving four geographical communities of modems. The CMTS has access to the one downstream and all four upstreams, while each CM has access to the one downstream and only one upstream.

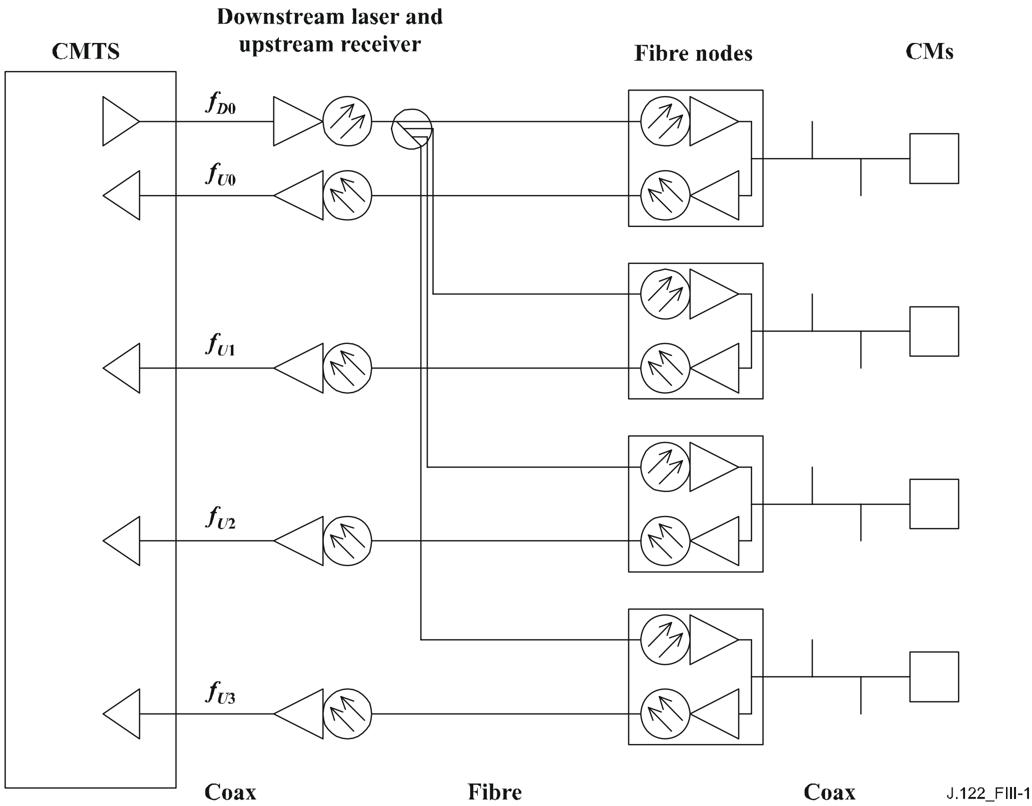

Figure III.1 – Single downstream and single upstream channels per CM

In this topology, the CMTS transmits upstream channel descriptors (UCDs) and MAPs for each of the four upstream channels related to the shared downstream channel.

Unfortunately, each CM cannot determine which fibre branch it is attached to because there is no way to convey the geographical information on the shared downstream channel. At initialization, the CM randomly picks a UCD and its corresponding MAP. The CM then chooses an initial maintenance opportunity on that channel and transmits a ranging request.

The CMTS will receive the ranging request and will redirect the CM to the appropriate upstream channel identifier by specifying the upstream channel ID in the ranging response. The CM MUST then use the channel ID of the ranging response, not the channel ID on which the ranging request was initiated. This is necessary only on the first ranging response received by the CM. The CM SHOULD continue the ranging process normally and proceed to wait for station maintenance IEs.

From then on, the CM will be using the MAP that is appropriate to the fibre branch to which it is connected. If the CM ever has to redo initial ranging, it may start with its previous known UCD instead of choosing one at random.

A number of constraints are imposed by this topology:

- All initial maintenance opportunities across all fibre nodes must be aligned. If there are multiple logical upstreams sharing the same spectrum on a fibre, then the initial maintenance opportunities for each of the logical upstreams MUST align with the initial maintenance opportunity of at least one logical upstream with the same centre frequency on each fibre node. When the CM chooses a UCD to use and then subsequently uses the MAP for that channel, the CMTS must be prepared to receive a ranging request at that initial maintenance opportunity. Note that only the initialization intervals must be aligned. Once the CM is successfully ranged on an upstream channel, its activities need only be aligned with other users on the same upstream channel. In Figure III.1, ordinary data transmission and requests for bandwidth may occur independently across the four upstream channels.

- All of the upstream channels on different nodes should operate at the same frequency or frequencies unless it is known that no other upstream service will be impacted due to a CM transmission of a ranging request on a "wrong" frequency during an initial maintenance opportunity. If the CM chooses an upstream channel descriptor arbitrarily, it could transmit on the wrong frequency if the selected UCD applied to an upstream channel on a different fibre node. This could cause initial ranging to take longer. However, this might be an acceptable system trade-off in order to keep spectrum management independent between cable segments.
- All of the upstream channels may operate at different modulation rates. However, there is a trade-off involved between the time it takes to acquire ranging parameters and flexibility of upstream channel modulation rate. If upstream modulation rates are not the same, the CMTS would be unable to demodulate the ranging request if it was transmitted at the wrong modulation rate for the particular upstream receiver of the channel. The result would be that the CM would retry as specified in the RFI Recommendation and then would eventually try other upstream channels associated with the currently used downstream. Increasing the probability of attempting ranging on multiple channels increases CM initialization time but using different modulation rates on different fibre nodes allows flexibility in setting the degree of burst noise mitigation.
- All initial maintenance opportunities on different channels may use different burst characteristics so that the CMTS can demodulate the ranging request. Again, this is a trade-off between time to acquire ranging and exercising flexibility in setting physical-layer parameters among different upstream channels. If upstream burst parameters for initial maintenance are not the same, the CMTS would be unable to demodulate the ranging request if it was transmitted with the wrong burst parameters for the particular channel. The result would be that the CM would retry the ranging request as specified in the RFI Recommendation and then would eventually try other upstream channels associated with the currently used downstream. Increasing the probability of attempting ranging on multiple channels increases CM initialization time but using different burst parameters for initial maintenance on different fibre nodes allows the ability to set parameters appropriate for plant conditions on a specific node.

III.2 Multiple downstreams and multiple upstreams per cable segment

This clause presents a more complex set of examples of CMs that are served by several downstream channels and several upstream channels and where those upstream and downstream channels are part of one MAC domain. The interaction of initial ranging, normal operation and dynamic channel change are profiled, as well as the impact of the multiple downstreams using synchronized or unsynchronized timestamps.

Synchronized timestamps refer to both downstream paths transmitting a time stamp that is derived from a common clock frequency and have common time bases. The timestamps on each downstream do not have to be transmitted at the same time in order to be considered synchronized.

III.2.1 Topologies

Suppose two downstream channels are used in conjunction with four upstream channels as shown in Figure III.2. In all three topologies, there are two geographical communities of modems, both served by the same two downstream channels. The difference in the topologies is found in their upstream connectivity.

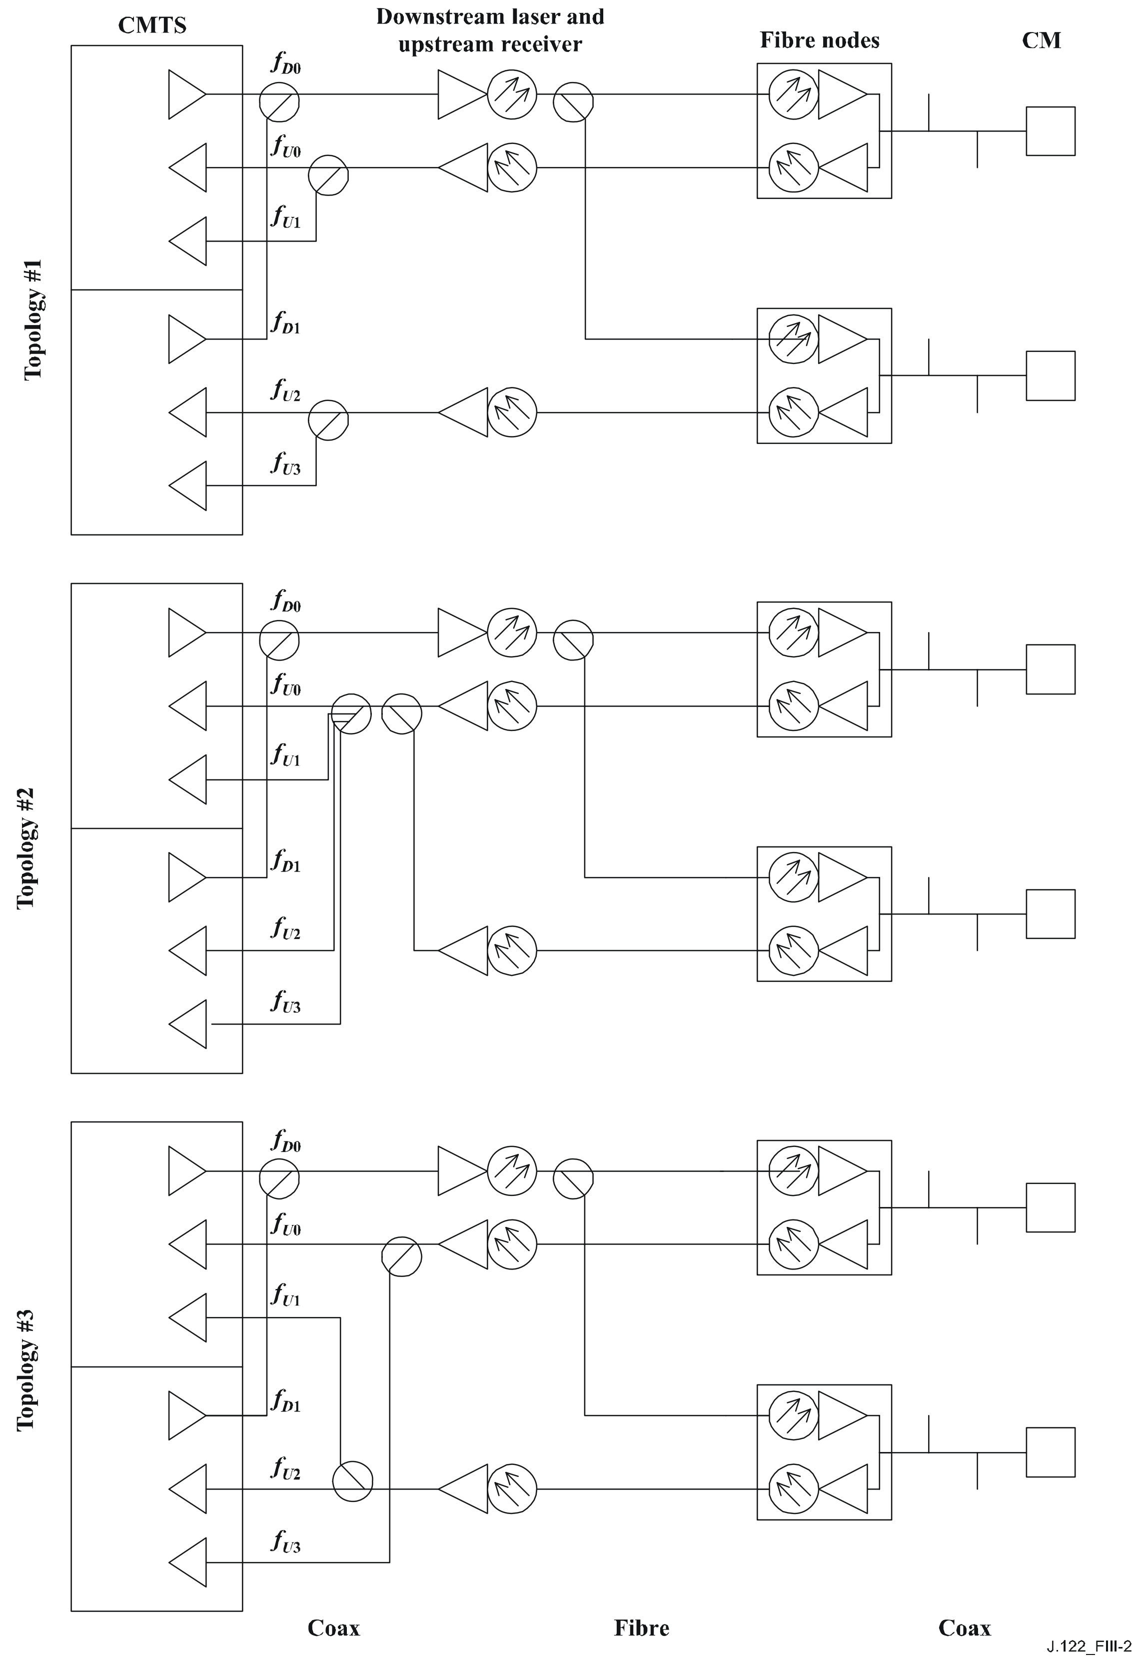

Figure III.2 – Multiple downstream and multiple upstream channels per CM

Topology No. 1 has the return path from each fibre node connected to a dedicated set of upstream receivers. A CM will see both downstream channels, but only one upstream channel that is associated with one of the two downstream channels.

Topology No. 2 has the return path from each fibre node combined and then split across all upstream receivers. A CM will see both downstream channels and all four upstream channels in use with both downstream channels.

Topology No. 3 has the return path from each fibre node split and then sent to multiple upstream receivers, each associated with a different downstream channel. A CM will see both downstream channels, and one upstream channel associated with each of the two downstream channels.

Topology No. 1 is the typical topology in use. Movement between downstreams can only occur if the timestamps on both downstreams are synchronized. Topology No. 2 and Topology No. 3 are to compensate for downstreams that have unsynchronized timestamps, and allow movement between downstream channels as long as the upstream channels are changed at the same time.

The CMs are capable of single frequency receive and single frequency transmit.

III.2.2 Normal operation

Table III.1 lists MAC messages that contain channel IDs.

Table III.1 – MAC messages with channel IDs

MAC message	Downstream channel ID	Upstream channel ID
UCD	Yes	Yes
MAP	No	Yes
RNG-REQ	Yes	No
RNG-RSP	No	Yes
DCC-REQ	Yes	Yes

With unsynchronized timestamps:

- Since upstream synchronization relies on downstream timestamps, each upstream channel must be associated with the time stamp of one of the downstream channels.
- The downstream channels should only transmit MAP messages and UCD messages that pertain to their associated upstream channels.

With synchronized timestamps:

- Since upstream synchronization can be obtained from either downstream channel, all upstreams can be associated with any downstream channel.
- All MAPs and UCDs for all upstream channels should be sent on all downstream channels. The UCD messages contains a downstream channel ID so that the CMTS can determine with the RNG-REQ message which downstream channel the CM is on. Thus the UCD messages on each downstream will contain different downstream channel IDs even though they might contain the same upstream channel ID.

III.2.3 Initial ranging

When a CM performs initial ranging, the topology is unknown and the timestamp consistency between downstreams is unknown. Therefore, the CM chooses either downstream channel and any one of the UCDs sent on that downstream channel.

In both cases:

- The upstream channel frequencies within a physical upstream or combined physical upstreams must be different.
- The constraints specified in clause III.1 apply.

III.2.4 Dynamic channel change

With unsynchronized timestamps:

- When a DCC-REQ is given, it must contain new upstream and new downstream frequency pairs that are both associated with the same timestamp.
- When the CM resynchronizes to the new downstream, it must allow for timestamp resynchronization without re-ranging unless instructed to do so with the DCC-REQ command.
- Topology No. 1 will support channel changes between local upstream channels present within a cable segment, but will not support changes between downstream channels. Topology No. 2 and No. 3 will support upstream and downstream channel changes on all channels within the fibre node as long as the new upstream and downstream channel pair are associated with the same timestamp.

With synchronized timestamps:

- Downstream channel changes and upstream channel changes are independent of each other.

Topologies No. 1, No. 2, and No. 3 will support changes between all upstream and all downstream channels present within the cable segment.

Appendix IV

DOCS transmission and contention resolution

(This appendix does not form an integral part of this Recommendation)

IV.1 Introduction

This appendix clarifies how the DOCS transmission and contention-resolution algorithms work. It contains a few minor simplifications and assumptions, but should be useful to help clarify this area of the Recommendation.

The simplifications include:

- The text does not explicitly discuss packet arrivals while deferring or waiting for pending grants, nor the sizing of piggyback requests.
- The CM always sends a piggyback request for the next frame in the last fragment and not inside one of the headers of the original frame.
- Much of this applies to concatenation, but no attempt is made to address all the subtleties of that situation.

The assumptions include, among others:

- The assumption is made that a request always fits in any request/data region.
- When a piggyback request is sent with a contention data packet, the state machine only checks for the grant to the request and assumes the data ack for the contention data packet was supplied by the CMTS.

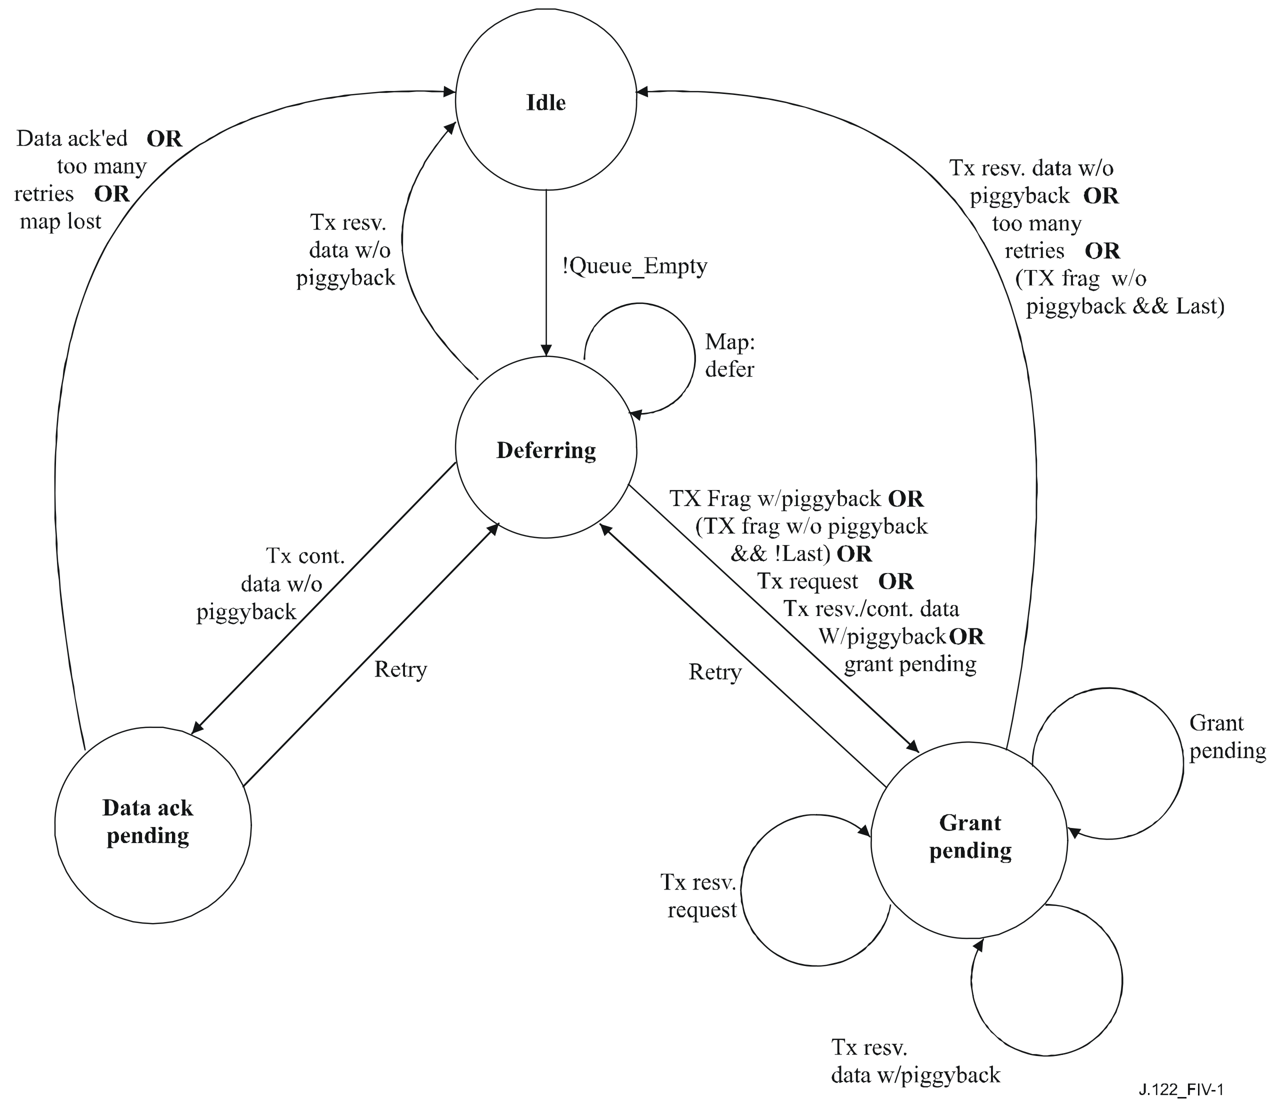

Figure IV.1 – Transmission and deference state transition diagram

IV.2 Variable definitions

Start	Data backoff start field from map "currently in effect"
End	Data backoff end field from map "currently in effect"
Window	Current backoff window
Random[n]	Random number generator that selects a number between 0 and n – 1
Defer	Number of transmit opportunities to defer before transmitting
Retries	Number of transmissions attempted without resolution
Tx_time	Saved time of when request or request/data was transmitted
Ack_time	Ack time field from current map
Piggyback	Flag set whenever a piggyback REQ is added to a transmit packet
Queue_Empty	Flag set whenever the data queue for this SID is empty
Lost_Map	Flag set whenever a MAP is lost while in state data ack pending
my_SID	Service ID of the queue that has a packet to transmit

pkt size	Data packet size including MAC and physical-layer overhead (including piggyback if used)
frag_size	Size of the fragment
Tx_Mode	{Full_Pkt; First_Frag; Middle_Frag; Last_Frag}
min_frag	Size of the minimum fragment

IV.3 State examples

IV.3.1 Idle – Waiting for a packet to transmit

```
Window = 0;
Retries = 0;

Wait for !Queue_Empty;  /* Packet available to transmit */
CalcDefer();
go to Deferring
```

IV.3.2 Data ack pending – Waiting for data ack only

```
Wait for next Map;

if (Data Acknowledge SID == my_SID)   /* Success! CMTS received data packet */
     go to state Idle;
else if (Ack_time > Tx_time)     /* COLLISION!!! or Pkt Lost or Map Lost */
{
     if (Lost_Map)
          go to state Idle;  /* Assume pkt was ack'ed to avoid sending
duplicates */
     else
          Retry();
}

stay in state Data Ack Pending;
```

IV.3.3 Grant pending – Waiting for a grant

```
Wait for next Map;

while (Grant SID == my_SID)
     UtilizeGrant();

if (Ack_time > Tx_time)  /* COLLISION!!!!! or Request denied/lost or Map Lost */
     Retry();
stay in state Grant Pending
```

IV.3.4 Deferring – Determine proper transmission timing and transmit

```
if (Grant SID == my_SID)              /* Unsolicited Grant */
{
     UtilizeGrant();
}
else if (unicast Request SID == my_SID)    /* Unsolicited Unicast Request */
{
     transmit Request in reservation;
     Tx_time = time;

     go to state Grant Pending;
}
else
{
     for (each Request or Request/Data Transmit Opportunity)
     {
```

```
        if (Defer != 0)
             Defer = Defer – 1; /* Keep deferring until Defer = 0 */
        else
        {
             if (Request/Data tx_op) and (Request/Data size >= pkt size)
                       /* Send data in contention */
             {
                  transmit data pkt in contention;
                  Tx_time = time;

                  if (Piggyback)
                       go to state Grant Pending;
                  else
                       go to state Data Ack Pending;
             }
             else           /* Send Request in contention */
             {
                  transmit Request in contention;
                  Tx_time = time;

                  go to state Grant Pending;
             }
        }
    }
}

Wait for next Map;
STAY IN STATE DEFERRING
```

IV.4 Function examples

IV.4.1 CalcDefer() – Determine defer amount

```
   if (Window < Start)
     Window = Start;

   if (Window > End)
     Window = End;

   Defer = Random[2^Window];
```

IV.4.2 UtilizeGrant() – Determine best use of a grant

```
if (Grant size >= pkt size)           /* CM can send full pkt */
{
     transmit packet in reservation;
     Tx_time = time;
     Tx_mode = Full_pkt

     if (Piggyback)
          go to state Grant Pending
     else
          go to state Idle;
}
else if (Grant size < min_frag && Grant Size > Request size)
     /* Can't send fragment, but can send a Request */
{
     transmit Request in reservation;
     Tx_time = time;
     go to state Grant Pending;
}
else if (Grant size == 0)             /* Grant Pending */
     go to state Grant Pending;
else
{
```

```
     while (pkt_size > 0 && Grant SID == my_SID)
     {
          if (Tx_mode == Full_Pkt)
               Tx_mode = First_frag;
          else
               Tx_mode = Middle_frag;

          pkt_size = pkt_size – frag_size;
          if (pkt_size == 0)
               Tx_mode = Last_frag;

          if (another Grant SID == my_SID) /* multiple grant mode */
               piggyback_size = 0
          else
               piggyback_size = pkt_size    /* piggyback mode */

          if (piggyback_size > 0)
               transmit fragment with piggyback request for remainder of packet
in reservation
          else
               transmit fragment in reservation;
     }
     go to state Grant Pending;
}
```

IV.4.3 Retry()

```
Retries = Retries + 1;
if (Retries > 16)
{
     discard pkt, indicate exception condition
     go to state Idle;
}

Window = Window + 1;

CalcDefer();

go to state Deferring;
```

Appendix V

IGMP example

(This appendix does not form an integral part of this Recommendation)

Clause 5.3.1 defines the requirements for CMTS and CM support of IGMP signalling. This appendix provides an example CM passive-mode state machine for maintaining membership of a single multicast group.

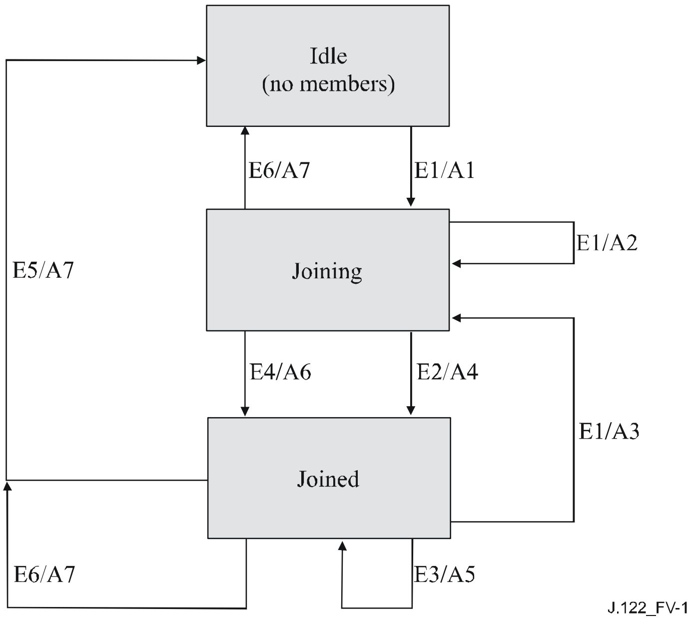

Figure V.1 – IGMP support – CM passive mode

V.1 Events

E1: MR received on CPE I/f

E2: M1 timer expired

E3: MQ received on RF I/f

E4: MR received on RF I/f

E5: M2 timer expired

E6: Auth failure[31]

[31] SA-MAP response returns an error code of 7, "not authorized for requested downstream traffic flow".

V.2 Actions

A1: MQI = 125 s; QRI = 10 s; start M1 timer with random value between 0 and 3 s; start M2 timer = 2 × MQI + QRI; start TEK machine, if necessary[32]; add multicast addr to multicast filter

A2: Discard MR packet

A3: Reset M2 timer = 2 × MQI + QRI; start M1 timer with random value between 0 and 3 s

A4: Transmit MR on RF I/f; set I = current time

A5: Recompute MQI = MAX(125, current time – I); set I = current time, forward MQ on CPE I/f

A6: Cancel M1 timer

A7: Delete multicast addr from multicast filter

[32] If the multicast traffic is encrypted, a TEK machine needs to be started to decrypt the encrypted multicast packets. To determine whether the multicast is encrypted, the CM makes a SA-MAP request to the CMTS to get the associated SAID of the multicast group address. If the SA-MAP response returns an SAID, then a TEK machine is started. No TEK machine is necessary if the SA-MAP response indicates that the multicast traffic is not encrypted. The SA-MAP response may also indicate that the CM is not authorized to receive this multicast traffic. In such a case, the CM terminates the multicast state machine and stops forwarding the multicast traffic.

Appendix VI

Unsolicited grant services

(This appendix does not form an integral part of this Recommendation)

This appendix discusses the intended use of the unsolicited grant service (UGS) and unsolicited grant service with activity detection (UGS-AD) and includes specific examples.

VI.1 Unsolicited grant service (UGS)

VI.1.1 Introduction

Unsolicited grant service is an upstream flow scheduling service type that is used for mapping constant bit rate (CBR) traffic onto service flows. Since the upstream is scheduled bandwidth, a CBR service can be established by the CMTS scheduling a steady stream of grants. These are referred to as unsolicited because the bandwidth is predetermined, and there are no ongoing requests being made.

The classic example of a CBR application of interest is voice over Internet Protocol (VoIP) packets. Other applications are likely to exist as well.

Upstream flow scheduling services are associated with service flows, each of which is associated with a single service ID (SID). Each service flow may have multiple classifiers. Each classifier may be associated with a unique CBR media stream. Classifiers may be added and removed from a service flow. Thus, the semantics of UGS must accommodate single or multiple CBR media streams per SID.

For the discussion within this appendix, a subflow will be defined as the output of a classifier. Since a VoIP session is identified with a classifier, a subflow in this context refers to a VoIP session.

VI.1.2 Configuration parameters

- Nominal grant interval.
- Unsolicited grant size.
- Tolerated grant jitter.
- Grants per interval.

Explanations of these parameters and their default values are provided in Annex C.

VI.1.3 Operation

When a service flow is provisioned for UGS, the nominal grant interval is chosen to equal the packet interval of the CBR application. For example, VoIP applications with 10-ms packet sizes will require a nominal grant interval of 10 ms. The size of the grant is chosen to satisfy the bandwidth requirements of the CBR application and relates directly to the length of the packet.

When multiple subflows are assigned to a UGS service, multiple grants per interval are issued. There is no explicit mapping of subflows to grants. The multiple grants per interval form a pool of grants in which any subflow can use any grant.

It is assumed in this operational example the default UGS case of no concatenation and no fragmentation.

VI.1.4 Jitter

Figure VI.1 shows the relationship between grant interval and tolerated grant jitter, and shows an example of jitter on subflows.

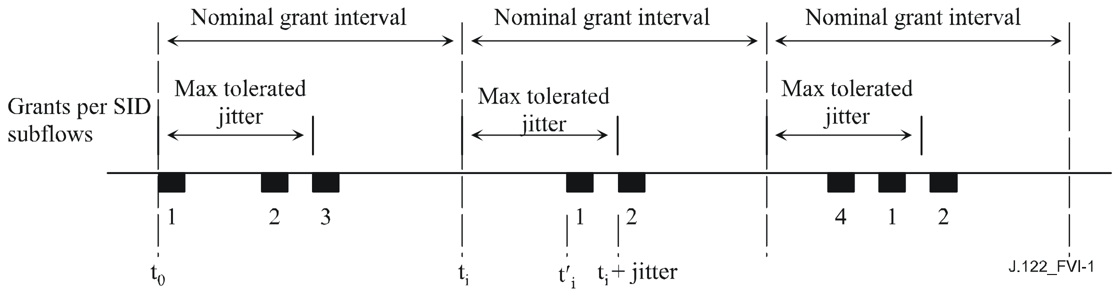

Figure VI.1 – Example jitter with multiple grants per SID

For only one grant per interval, the tolerated grant jitter is the maximum difference between the actual grant time (t_i') and the nominal grant time (t_i). For multiple grants per interval, the tolerated grant jitter is the maximum difference between the actual time of the last grant in the group of grants and the nominal grant time (t_i). If the arrival of any grant is at t_i', then $t_i \leq t_i' \leq t_i$ +jitter.

Figure VI.1 demonstrates how a subflow will be jittered even though the individual grants may not move from their relative position. During the first interval, three VoIP sessions are established, and they happen to fall on the three grants. In the second interval, VoIP session 3 has been torn down. Since the CMTS does not know which subflow is associated with which grant, it decides to remove the first grant. The remaining two calls shift to the other two grants. In the third interval, a new VoIP session 4 and a new grant have been added. The new call happens to fall on the new grant. The net effect is that the subflows may move around within their jitter interval.

The advantage of a small jitter interval is that the VoIP receive jitter buffer may be kept small. The disadvantage is that this places a scheduling constraint on the CMTS.

The boundary of a nominal grant interval is arbitrary and is not communicated between the CMTS and the CM.

NOTE – More dramatic events, like the loss of a downstream MAP or the frequency hopping of an upstream, may cause subflows to jitter outside of this jitter window.

VI.1.5 Synchronization issues

There are two synchronization problems that occur when carrying CBR traffic such as VoIP sessions across a network. The first is a frequency mismatch between the source clock and the destination clock. This is managed by the VoIP application, and is beyond the scope of this Recommendation. The second is the frequency mismatch between the CBR source/sinks, and the bearer channel that carries them.

Specifically, if the clock that generates the VoIP packets towards the upstream is not synchronized with the clock at the CMTS that is providing the UGS service, the VoIP packets may begin to accumulate in the CM. This could also occur if a MAP was lost, causing packets to accumulate.

When the CM detects this condition, it asserts the queue indicator in the service flow EH element. The CMTS will respond by issuing an occasional extra grant so as to not exceed 1% of the provisioned bandwidth (this corresponds to a maximum of one extra grant every one hundred grants). The CMTS will continue to supply this extra bandwidth until the CM deasserts this bit.

A similar problem occurs in the downstream. The far-end transmitting source may not be frequency-synchronized to the clock that drives the CMTS. Thus, the CMTS SHOULD police at a rate slightly higher than the exact provisioned rate to allow for this mismatch and to prevent delay buildup or packet drops at the CMTS.

VI.2 Unsolicited grant service with activity detection (UGS-AD)

VI.2.1 Introduction

Unsolicited grant service with activity detection (UGS-AD) is an upstream flow scheduling service type. This clause describes one application of UGS-AD, which is the support for voice activity detection (VAD). VAD is also known as silence suppression and is a voice technique in which the transmitting CODEC sends voice samples only when there is significant voice energy present. The receiving CODEC will compensate for the silence intervals by inserting silence or comfort noise equal to the perceived background noise of the conversation.

The advantage of VAD is the reduction of network bandwidth required for a conversation. It is estimated that 60% of a voice conversation is silence. With that silence removed, that would allow a network to handle substantially more traffic.

Subflows in this context will be described as active and inactive. Both of these states are within the MAC layer QOS state known as active.

VI.2.2 MAC configuration parameters

The configuration parameters include all of the normal UGS parameters, plus:

- nominal polling interval;
- tolerated poll jitter.

Explanation of these parameters and their default values are provided in Annex C.

VI.2.3 Operation

When there is no activity, the CMTS sends polled requests to the CM. When there is activity, the CMTS sends unsolicited grants to the CM. The CM indicates the number of grants per interval that it currently requires in the active grant field of the UGSH in each packet of each unsolicited grant. The CM may request up to the maximum active grants per interval. The CM constantly sends this state information so that no explicit acknowledgment is required from the CMTS.

It is left to the implementation of the CM to determine activity levels. Implementation options include:

- Having the MAC layer service provide an activity timer per classifier. The MAC layer service would mark a subflow inactive if packets stopped arriving for a certain time, and mark a subflow active the moment a new packet arrived. The number of grants requested would equal the number of active subflows.
- Having a higher layer service entity, such as an embedded media client, which indicates activity to the MAC layer service.

When the CM is receiving polled requests and it detects activity, the CM requests enough bandwidth for one grant per interval. If activity is for more than one subflow, the CM will indicate this in the active grant field of the UGSH beginning with the first packet it sends.

When the CM is receiving unsolicited grants, then detects new activity, and asks for one more grant, there will be a delay in time before it receives the new grant. During that delay, packets may build up at the CM. When the new unsolicited grant is added, the CMTS will burst extra grants to clear out the packet buildup.

When the CM is receiving unsolicited grants, then detects inactivity on a subflow and asks for one less grant, there will be a delay in time before the reduction in grants occurs. If there has been any build up of packets in the upstream transmit queue, the extra grants will reduce or empty the queue. This is fine, and keeps system latency low. The relationship of which subflow is getting which specific grant will also change. This effect appears as low-frequency jitter that the far end must manage.

When the CM is receiving unsolicited grants and detects no activity on any of its subflows, it will send one packet with the active grants field of the UGSH set to zero grants, and then cease transmission. The CMTS will switch from UGS mode to real time polling mode. When activity is again detected, the CM sends a request in one of these polls to resume delivery of unsolicited grants. The CMTS ignores the size of the request and resumes allocating grant size grants to the CM.

It is not necessary for the CMTS to separately monitor packet activity since the CM does this already. In the worst case, if the CMTS misses the last packet that indicated zero grants, the CMTS and CM would be back in sync at the beginning of the next talk spurt. Because of this scenario, when the CM goes from inactive to active, the CM must be able to restart transmission with either polled requests or unsolicited grants.

VI.2.4 Example

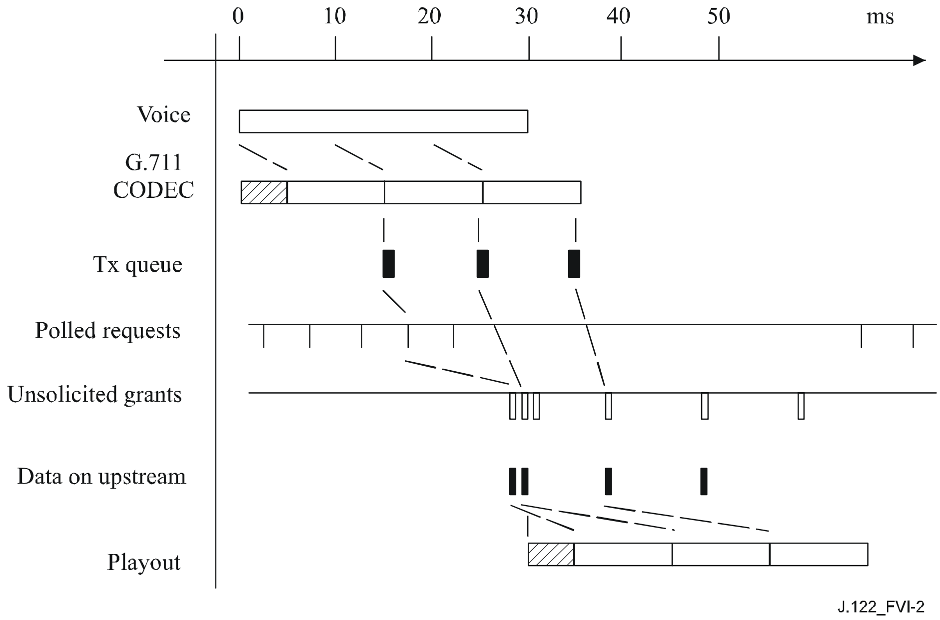

Figure VI.2 – VAD start-up and stop

Figure VI.2 shows an example of a single G.711 (64 kbit/s) voice call with a packet size of 10 ms, and a receive jitter buffer that requires a minimum of 20 ms of voice (thus, 2 packets) before it will begin playout.

Assume voice begins at time zero. After a nominal processing delay and a 10-ms packetization delay, the DSP CODEC generates voice packets that are then transferred to the upstream transmit queue. The next polled request is used, which results in the start of the unsolicited grants some time later. Additional unsolicited grants are immediately issued to clear out the upstream queue.

These packets traverse the network and arrive at the receive jitter buffer. The 20-ms minimum jitter buffer is met when the second packet arrives. Because the packets arrive close together, only an additional few milliseconds of latency is been added. After a nominal processing delay, playout begins.

When the voice spurt ends, the CM sends one remaining packet with no payload and with the active grants field of the UGSH set to zero grants. Some time later, UGS stops, and real-time polling begins.

VI.2.5 Talk spurt grant burst

The extra burst of unsolicited grants when a flow becomes active is necessary because the jitter buffer at the receiving CODEC typically waits to have a minimum amount of voice samples before beginning the playout. Any delay between the arrival of these initial packets will add to the final latency of the phone call. Thus, the sooner the CMTS recognizes that the CM has packets to send and can empty the CM's buffer, the sooner those packets will reach the receiver, and the lower the latency that will be incurred in the phone call.

It is an indeterminate problem as to how many grants must be burst. When the CM makes its request for an additional grant, one voice packet has already accumulated. The CM has no idea how many extra grants to request as it has no idea of the round-trip response time it will receive from the CMTS, and thus how many packets may accumulate. The CMTS has a better idea, although it does not know the far-end jitter buffer requirements.

The solution is for the CMTS to choose the burst size, and burst these grants close together at the beginning of the talk spurt. This occurs when moving from real-time polling to UGS, and when increasing the number of UGS grants per interval.

A typical start-up latency that will be introduced by the request to grant response time is shown in Table VI.1.

Table VI.1 – Example request to grant response time

Variable		Example value
1	The time taken from the voice packet being created to the time that voice packet arrives in the CM upstream queue.	0-1 ms
2	The time until a polled request is received. The worst case time is the polled request interval.	0-5 ms
3	The request-grant response time of the CMTS. This value is affected by MAP length and the number of outstanding MAPS.	5-15 ms
4	The round-trip delay of the HFC plant, including the downstream interleaving delay.	1-5 ms
	Total	6-26 ms

This number will vary between CMTS implementations, but a reasonable number of extra grants to expect from the example above would be:

Table VI.2 – Example extra grants for new talk spurts

UGS interval	Extra grants for new talk spurts
10 ms	2
20 ms	1
30 ms	0

Once again, it is worth noting that the CMTS and CM cannot and do not associate individual subflows with individual grants. That means that when current subflows are active and a new subflow becomes active, the new subflow will immediately begin to use the existing pool of grants. This potentially reduces the start-up latency of new talk spurts, but increases the latency of the other subflows. When the burst of grants arrives, it is shared with all the subflows, and restores or even reduces the original latency. This is a jitter component. The more subflows that are active, the less impact that adding a new subflow has.

VI.2.6 Admission considerations

Note that when configuring the CMTS admission control, the following factors must be taken into account.

VAD allows the upstream to be over-provisioned. For example, an upstream that might normally handle 24 VoIP sessions might be over-provisioned by as much as 36 (50%) or even 48 (100%). Whenever there is over-provisioning, there exists the statistical possibility that all upstream VoIP sessions may become active. At that time, the CMTS may be unable to schedule all the VoIP traffic. Additionally, the talk spurt grant bursts would be stretched out. CM implementations of VAD should recognize this possibility, and set a limit as to how many packets they will allow to accumulate on its queue.

Occasional saturation of the upstream during VAD can be eliminated by provisioning the maximum number of permitted VoIP sessions to be less than the maximum capacity of the upstream with all voice traffic (24 in the previous example). VAD would cause the channel usage to drop from 100% to around 40% for voice, allowing the remaining 60% to be used for data and maintenance traffic.

Appendix VII

S-CDMA framing

(This appendix does not form an integral part of this Recommendation)

This appendix is informative. In case of conflict between this appendix and any normative part of this Recommendation, the normative part takes precedence.

Please note that the pseudo-code below is specific to the case of a single burst using all spreading codes.

VII.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 that 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 initialized 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 );
     /* 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 );
          /* Store the coded portion of the symbol (or preamble) to
(row,coded_col) */
     }
}
```

VII.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).

VII.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.

```
for( col=0; col < framelen; col++ )
     for( row=0; row < ACTIVE_CODES; row++ )
          outsym = get_data( row, col );
```

VII.4 Comments

In all of the samples, the number of iterations for the loop is not always correct since an allocation can be less than the number of codes. In clause VII.2, the listing supports the case of a shortened subframe.

Appendix VIII

Ambient temperature and wind loading effects

(This appendix does not form an integral part of this Recommendation)

This appendix discusses possible ambient temperature change and dynamic wind loading effects relevant to operating a system with DOCS 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 are 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.

VIII.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 seconds apart. Table VIII.1 lists plant delay variations that can be accommodated in S-CDMA and TDMA modes for a 1-dB degradation under example conditions.

Table VIII.1 – Allowable plant timing drift

Constellation	E_s/N_o for 1e-8 BER (dB)	Allowable peak-to-peak plant delay variation (ns) S-CDMA mode	Allowable peak-to-peak plant delay variation (ns)TDMA mode
Fully-coded QPSK	5	90	800
TCM QPSK	9	79	N/A
TCM 8QAM	12	57	N/A
Uncoded QPSK	15	38	800
Fully-coded 64QAM	17.7	24	800
TCM 32QAM	19	18	N/A
Uncoded 16QAM	22	9	800
Uncoded 32QAM	25	6	800
TCM 128QAM	25	6	N/A
Uncoded 64QAM	28	2	800

Defined conditions:

- 1 dB degradation at 1e-8 BER;
- uniform ranging offset over ±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-second 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 4-2.

DOCS requires station maintenance at least every 30 seconds (T4 time-out has a maximum value of 35 seconds). 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 advanced TDMA mode.

The following clauses discuss the relationship of temperature changes and wind loading on the propagation delay in coaxial and HFC cable plants.

VIII.2 Change in propagation delay due to temperature changes

VIII.2.1 Fibre 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 fibre. A typical value for propagation delay variation due to temperature change of the fibre is 44 picoseconds per km per degree C [b-TRISCHITTA]. The delay variation comes mostly from the change in refractive index of the glass with temperature, not the change in fibre 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 fibre from stresses in the cable itself. The fibre usually sits loosely in a tube inside the cable and considerable relative movement is possible. This construction allows normal cable handling and aerial deployment without resulting in high stress on the optical fibre.

Assuming 44 picoseconds per km per degree C, any product of optical cable length and temperature change that equals 50 results in approximately a 2-nanosecond change in the fibre propagation delay. For example, 25-km fibre 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 DOCS of 100 miles, or roughly 160 km, the temperature change needed for a 2-ns second change in one-way propagation delay is 0.3 degree C.

Obviously, the issue is how fast the cable core (where the fibre 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 considerably higher than ambient temperatures 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.

VIII.2.2 Coaxial cable delay changes due to temperature

The coaxial cable has a blown foam between the centre conductor and the solid shield, and nominal propagation velocity is about 87% of free space velocity according to one manufacturer. 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.

VIII.2.3 Delay change due to wind

Aerial cable stretches with wind loading, 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 typical cable was investigated with a programme from a coaxial cable manufacturer. These calculations showed that length changes in the range 0.01% to 0.05% are possible for various amounts of wind loading. This converts to significant propagation delay variation. As an example, with 5 miles (8 km) and 0.02% length variation, the change in propagation delay is:

$$\frac{8\ \text{km}}{3\times10^{5}\ \text{km/s}}\times\frac{1}{0.87}\times2\times10^{-4}=6\ \text{ns}$$

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 writing, wind gusts of up to 40 mph are reported, while average wind speed is about 10 mph. Hence, over a period of 1 to 4 seconds (the typical wind gust measurement interval), the wind speed changes by 30 mph. Much stronger wind gusts are frequently measured in locations prone to windy conditions.

Bibliography

[b-Cab 960702]	CableLabs SP-CMTS-NSII01-960702 (1996), *Data-Over-Cable Service Interface Specifications: Cable Modem Termination System – Network Side Interface Specification*.
[b-Cab 010829]	CableLabs SP-CMCI-I06-010829 (2001), *Data-Over-Cable Service Interface Specifications: Cable Modem to Customer Premise Equipment Interface Specification*.
[b-ETSI EG 201 212]	ETSI EG 201 212 (1998), *Electrical safety; Classification of interfaces for equipment to be connected to telecommunication networks*. (Also available from CENELEC as ROBT-002.)
[b-EN 50081-1]	CENELEC EN 50081-1 (1992), *Electromagnetic compatibility – Generic emission standard – Part 1: Residential, commercial, and light industry*.
[b-EN 50082-1]	CENELEC EN 50082-1 (1997), *Electromagnetic compatibility – Generic immunity standard – Part 1: Residential, commercial, and light industry*.
[b-EN 50083-1]	CENELEC EN 50083-1 (1993), *Cabled networks for television signals, sound signals, and interactive services – Part 1: Safety requirements*.
[b-EN 50083-2]	CENELEC EN 50083-2 (1995), *Cabled networks for television signals, sound signals, and interactive services – Part 2: Electromagnetic compatibility for equipment*.
[b-EN 50083-7]	CENELEC EN 50083-7 (1996), *Cabled networks for television signals, sound signals, and interactive services – Part 7: System performance*.
[b-EN 50083-10]	CENELEC EN 50083-10 (1999), *Cabled networks for television signals, sound signals, and interactive services – Part 10: System performance for return paths*.
[b-EN 60950]	CENELEC EN 60950 (1992), *Specification for safety of information technology equipment*.
[b-US 47CFR 18.305]	*Code of Federal Regulations*, Title 47, Parts 15 and 76, October 2000.
[b-CommScope]	CommScope catalog for Parameter III cable (QR cable is listed as 88%).
[b-ISO/IEC 10038]	ISO/IEC 10038:1993, *Information Technology – Telecommunications and information exchange between systems – Local area networks – Media access control (MAC) bridges*.
[b-RFC 1700]	IETF RFC 1700 (1994), *Assigned Numbers*.
[b-RFC 2210]	IETF RFC 2210 (1997), *The Use of RSVP with the IETF Integrated Services*.
[b-RFC 2211]	IETF RFC 2211 (1997), *Specification of the Controlled-Load Network Element Service*.
[b-TRISCHITTA]	TRISCHITTA (P.R.), VARMA (E.L.): Jitter in Digital Transmission Systems, *Artech House*, Norwood, MA, 1989.

[b-SpanMaster] SpanMaster: available at www.commscope.com/html/community_access_cable.shtml.

[b-VCCI-1] Voluntary Control Council for Interference by Information Technology Equipment (VCCI-1): http://www.vcci.or.jp/vvci_e/index.html.

[b-Ethernet] Ethernet Protocol Version 2.0, Digital, Intel, Xerox, 1982.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	**Cable networks and transmission of television, sound programme and other multimedia signals**
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Terminals and subjective and objective assessment methods
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects and next-generation networks
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2009