

HFC Access Plant and DOCSIS Cable Modem

This lecture describes a modern Hybrid Fibre-Co-ax (HFC) cable TV access network and the most commonly used aspects of the DOCSIS version 1.0 cable modem specification.

After this lecture you should be able to:

- (1) distinguish between trunk, distribution and drop cables, explain the location and role of optical nodes and distribution amplifiers, convert signal levels between mV, dBm and dBmV, compute the output noise thermal power for an RF amplifier, identify the source of CTB, CSO, gain flatness and group delay distortions and their units,
- (2) distinguish between upstream and downstream links based on frequency, transmitting device (CM or CMTS), and direction of information flow (US or DS); list four factors that result in lower data rates and more complex processing on the upstream; compute the maximum correctable error burst length for a given value of t and block interleaver size; compute the raw US and DS data rates for a particular combination of modulation and symbol rate;
- (3) list the main purpose for each of the (5) protocol layers used by a DOCSIS CM; explain two reasons for the differences between upstream and downstream MAC protocols; decode the information contained in the MAP frame time allocation information element, and list the PHY parameters adjusted during ranging.

HFC Cable System Architecture

Figure 1 below¹ shows a block diagram of a typical network used by “cable companies” to provide broadcast TV, Internet access and telephone service using VoIP protocols. This HFC (Hybrid Fibre-Co-ax) access plant uses both fiber-optic and co-axial cables.

The headend contains links to other service providers: satellite receivers for broadcast services and fiber optic (FO) links to Internet and PSTN gateways. The signal is distributed over fiber optic cables first, then trunk cables, then distribution cables and finally over one drop cable for each subscriber.

The initial distribution is over one fiber in each direction to “optical nodes”. FO cables have much lower loss than co-ax cables — about 0.3 dB/km compared to about 30 dB/km for 1” co-ax — and are smaller, lighter and less expensive. These fiber optic cables extend the range of the system by distributing the signal with minimum degradation in quality.


Trunk cables further extend the range of the system. Since each amplifier degrades the signal by introducing noise and distortion, trunk cables are designed to minimize loss. Trunk cables are usually 1” or 3/4” diameter low-loss co-ax. Such a co-ax trunk cable requires an amplifier approximately every 1 km to make up for the attenuation of the cable.

Distribution cables are co-ax cables that run down each street and past each house. The distribution ca-

ble is typically 1/2” co-ax. Each distribution cable will have many taps (splitters) along with distribution amplifiers to make up for the loss of signal power due to the splitters.

Each subscriber connects to a tap in the distribution cable with a “drop” cable. This cable is typically flexible co-ax less than 50 m long.

Exercise 1: In the example in the figure below, a city block is 100m long and there are 10 houses on the block. The distribution cable runs down a side street and the along the street being serviced. What total length of distribution cable is required? If each drop is 20m long, how much drop cable is required? what is the ratio of distribution cable to drop cable?



To transmit the signal over fiber, the frequency-multiplexed TV and cable modem signal is used to amplitude-modulate a laser-generated optical carrier at 1550 or 1310 nm. At the receiver a photodiode recovers the original electrical signal. Note that even though the signal itself contains many modulated carriers at frequencies up to about 1 GHz, to the optical modulator/demodulator the FDM DS signal is a single low-pass baseband signal.

¹From the Wikipedia [HFC](#) article.

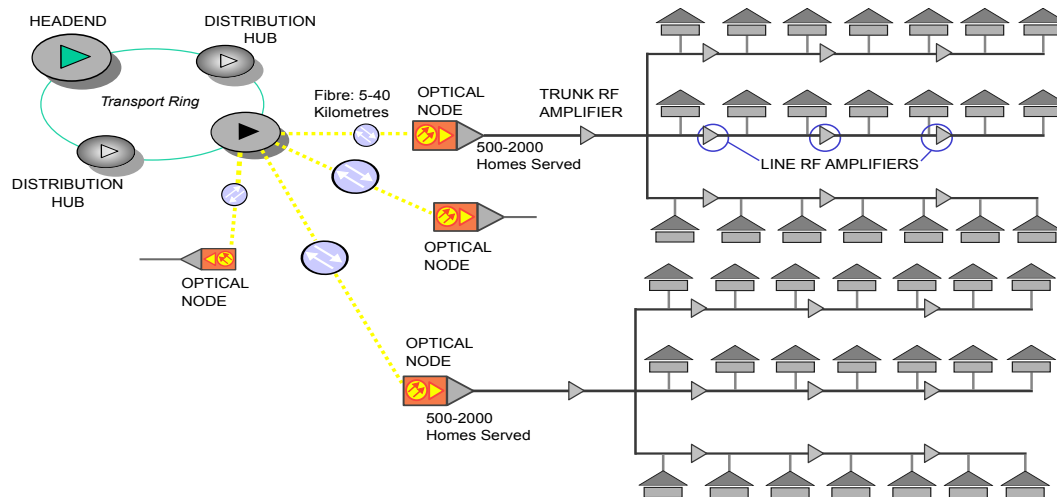


Figure 1: Typical HFC Network.

The hardware component that converts between the RF and optical signals (in both directions) is called an optical node. An optical node may serve several hundred subscribers and costs about \$2000.

Components

Not surprisingly, the largest capital expense in setting up a CATV system is the co-ax cable itself. CATV systems always use co-ax with a 75 ohm characteristic impedance.


The loss of co-ax cables, measured in dB, is proportional to distance and depends on the dimensions and on the dielectric and conductor materials used (e.g. copper vs aluminum conductors, foamed vs solid polyethylene dielectric and braid, foil or solid shield). Manufacturer's data sheets need to be consulted to obtain the loss per unit distance (x dB/100m) at various frequencies. The loss also varies with temperature.

Exercise 2: A manufacturer specifies that one of their co-ax cables has a loss of 2.5 dB/100m at 100 MHz. What would be the attenuation, at 100 MHz, of 20 m of this cable?

Trunk and distribution cables are typically copper "semi-rigid" cable with a shield of thin-walled copper tubing while drop cables are flexible aluminum cable with braided aluminum wire and foil shield.

Another expensive component is the distribution amplifiers and optical nodes. Distribution amplifiers and optical nodes are mounted in sealed enclo-

tures with external heat sinks that look like this when mounted on cables between poles:

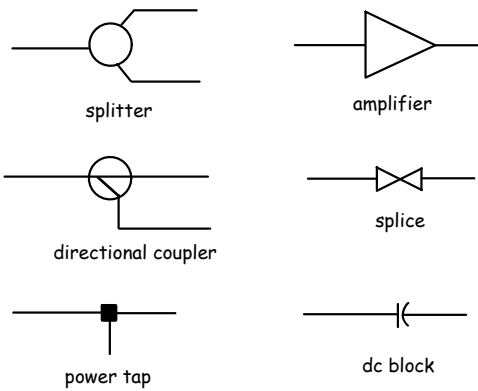


Inside the case are the amplifier electronics, test points, configuration jumpers and adjustments for gain.

Power supplies, often with battery backup, are used to inject power, typically 90VAC, into different places in the cable system to supply the amplifiers and optical nodes. Power for the amplifiers and optical nodes is usually distributed over a co-ax cable center conductor. Power distribution, including battery backup systems, is an important part of the HFC plant.

Optical nodes and amplifiers are usually designed to pass significant current (up to 15A) to supply downstream optical nodes and distribution amplifiers.

Here are some symbols used in diagrams of CATV systems:



20 dB amplifier. What is the C/N at the output of the second amplifier?

Signal Levels

Signal levels cable TV systems are usually measured in dBmV. This is the signal level relative to 1 millivolt (mV) (not dBm which is relative to 1 mW). This is a convenient unit because the input to a TV receiver is typically 0 dBmV. When computing power levels an impedance of 75Ω is used.

Exercise 3: What is voltage in mV of a 0 dBmV signal? What is the power in mW if measured across an impedance of 75 ohms? What is this power in dBm?

Noise and Distortion

Since the signal may have to pass through many amplifiers before reaching the subscriber, a CATV system is particularly sensitive to noise and distortion caused by distribution amplifiers.

The thermal noise level at the output of an amplifier is given by $N = G \cdot kTBF$ where G is the amplifier gain, k is Boltzman's constant, T is the temperature (typically assumed to be 290K), B is the bandwidth in Hz and F is the amplifier noise figure. When this value is computed in dB, kT is -174 dBm-Hz.

Exercise 4: A 0 dBmV signal is applied to the input of a distribution amplifier with a gain of 20 dB and a 2 dB noise figure. What are the output signal and noise powers? What is the output C/N (carrier to noise power ratio) measured over a 6 MHz bandwidth?

The noise figure of amplifiers closest to the head end are most critical since the noise output by an amplifier is amplified by every downstream amplifier.

Exercise 5: The signal and noise then travel down a length of co-ax that introduces 20 dB loss and then through another

Distortion

Non-linearities in amplifiers causes inter-modulation distortion (IMD) that results in additional frequency components. These distortion products may cause interference to other signals.

The most significant distortion products are those due to second- and third-order products. Operators specify second- and third-order distortion using measures called "composite second order" (CSO) and "composite tripe beat" (CTB). These are defined as the difference (in dB) between the carrier level and the levels of these distortion products. We will not go into more detail on how these are measured.

Linear distortion must also be controlled. Since the loss of co-ax cable increases with frequency, the gain of distribution amplifiers must increase with frequency to compensate. The gain and the slope of the gain versus frequency curve can be adjusted at each distribution amplifier to compensate for different types and lengths of cable. This is called "equalization".

A CATV system will have requirements for gain flatness (maximum loss difference in dB) which is specified over the upstream or downstream frequency range and maximum group delay (ns) which is specified over the 6 MHz bandwidth of a TV channel.

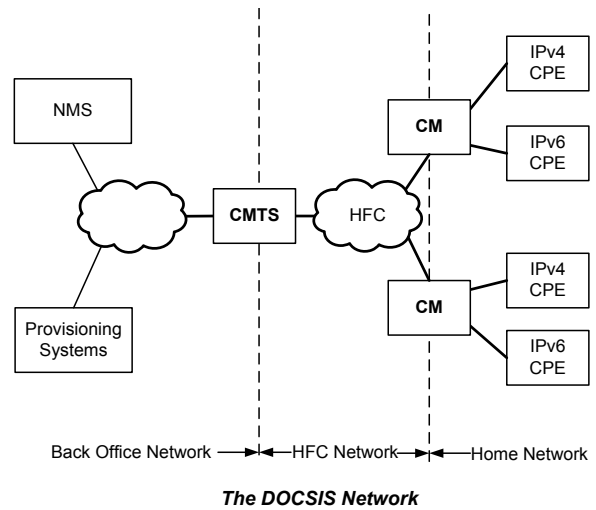

Frequencies

All modern cable systems are bidirectional. This allows a signal to be returned from subscribers to the head-end to support services such as internet access and telephony. Signals above a certain frequency (either 54 or 108 MHz) propagate from the head end to the subscriber. This is called the forward or downstream (D, DS) direction. Frequencies from 5 to 30 MHz (or 42 or 85 MHz) propagate from the subscriber toward the head end. This is called the reverse, return, or upstream (U, US) direction.

The actual frequency ranges in each direction are chosen by the cable company and depend on cost

considerations. Systems with larger bandwidths require more and/or better amplifiers and are thus more costly. The maximum downstream frequency may be between 300 and 1002 MHz.

In a bidirectional system all of the distribution amplifiers need to be bi-directional. A pair of analog filters called a diplexer (diplex filter) is used on each end of the amplifier to separate the two directions based on frequency. A simplified diagram is shown below:



EuroDOCSIS is a version of DOCSIS for the 8 MHz channel spacing used in Europe as opposed to the 6 MHz spacing used in North America.

A CMTS is typically a rack-mounted router with plug-in cards supporting multiple upstream and downstream RF channels such as the following Cisco uBR10012 Universal Broadband Router (CMTS):



A CM is typically a small desktop device that connects to the cable TV co-ax cable and an Ethernet port such as the following Motorola SB5100 Cable Modem (CM):




DOCSIS

DOCSIS (Data Over Cable Service Interface Specification) is a standard defined by CableLabs, an R&D and standards organization funded by the cable industry.

DOCSIS has progressed through various versions (1, 1.1, 2 and 3), each one providing improved data rates (primarily upstream), enhanced QoS (Quality of Service, delivery of time-sensitive data), and security.

DOCSIS defines the interface between cable modems (CMs) at the subscriber's premises and a cable modem termination system (CMTS) at the head end as shown in these diagrams from the DOCSIS 3.0 MAC and OSSI (Operations Support System Interface) specs²:



²CPE is customer premises equipment, NMS is network management system

Channel Specifications

DOCSIS systems make use of the bidirectional HFC access plant with upstream transmission at frequencies from 5 to 42 MHz (sometimes higher) and up-

stream transmission from 52 to 860 MHz (often lower).

The following table shows the assumptions about the downstream channel as specified in DOCSIS version 1.0:

Table 2-1. Assumed Downstream RF Channel Transmission Characteristics

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 msec (typically much less)
Carrier-to-noise ratio in a 6-MHz band (analog video level)	Not less than 35 dB (Note 4)
Carrier-to-interference ratio for total power (discrete and broadband ingress signals)	Not less than 35 dB within the design bandwidth
Composite triple beat distortion for analog modulated carriers	Not greater than -50 dBc within the design bandwidth
Composite second order distortion for analog modulated carriers	Not greater than -50 dBc within the design bandwidth
Cross-modulation level	Not greater than -40 dBc within the design bandwidth
Amplitude ripple	0.5 dB within the design bandwidth
Group delay ripple in the spectrum occupied by the CMTS	75 ns within the design bandwidth
Micro-reflections bound for dominant echo	-10 dBc @ ≤ 0.5 μ sec, -15 dBc @ ≤ 1.0 μ sec -20 dBc @ ≤ 1.5 μ sec, -30 dBc @ > 1.5 μ sec
Carrier hum modulation	Not greater than -26 dBc (5%)
Burst noise	Not longer than 25 μ sec at a 10 Hz average rate
Seasonal and diurnal signal level variation	8 dB
Signal level slope, 50-750 MHz	16 dB
Maximum analog video carrier level at the CM input, inclusive of above signal level variation	17 dBmV
Lowest analog video carrier level at the CM input, inclusive of above signal level variation	-5 dBmV

Notes to Table 2-1

1. Transmission is from the headend combiner to the CM input at the customer location.
2. For measurements above the normal downstream operating frequency band (except hum), impairments are referenced to the highest-frequency NTSC carrier level.
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 NTSC carrier.
4. This presumes that the digital carrier is operated at analog peak carrier level. When the digital carrier is operated below the analog peak carrier level, this C/N may be less.
5. Measurement methods defined in [NCTA] or [CableLabs2].

Exercise 6: Assuming the carrier power is the signal power, what is the Shannon capacity of this channel? What is the maximum symbol rate if the symbol period is to be more than 10 times the group delay ripple? What is the maximum area (in km²) that can be serviced by one DOCSIS CMTS assuming a velocity factor of 0.66?

The following table shows the corresponding assumptions about the upstream channel:

Table 2-2. Assumed Upstream RF Channel Transmission Characteristics

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 msec (typically much less)
Carrier-to-noise ratio	Not less than 25 dB
Carrier-to-ingress power (the sum of discrete and broadband ingress signals) ratio	Not less than 25 dB (Note 2)
Carrier-to-interference (the sum of noise, distortion, common-path distortion and cross-modulation) ratio	Not less than 25 dB
Carrier hum modulation	Not greater than -23 dBc (7.0%)
Burst noise	Not longer than 10 μ sec at a 1 kHz average rate for most cases (Notes 3, 4 and 5)
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 μ sec -20 dBc @ ≤ 1.0 μ sec -30 dBc @ > 1.0 μ sec
Seasonal and diurnal signal level variation	Not greater than 8 dB min to max


Notes to Table 2-2

1. Transmission is from the CM output at the customer location to the headend.
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 [CableLabs1].
3. Amplitude and frequency characteristics sufficiently strong to partially or wholly mask the data carrier.
4. CableLabs report containing distribution of return-path burst noise measurements and measurement method is forthcoming.
5. Impulse noise levels more prevalent at lower frequencies (< 15 MHz).

Exercise 7: Again compute the Shannon capacity based on the CNR and the maximum symbol rate based on a period equal to 10 times the group delay ripple.

Upstream data transmission in the HFC plant is significantly more difficult than downstream because:

- much less bandwidth is available (e.g. 37 MHz vs 810 MHz)
- CMs cannot hear each other's transmissions so the CMTS must coordinate upstream transmissions to avoid collisions
- high-power HF transmitters operating at the upstream frequencies and can leak into ("ingress") the cable system
- noise and interference sources appear combined at the head end, forming a "funnel" effect as shown below:




Revisions to the DOCSIS spec since the first version in 1997 have therefore concentrated on improving upstream performance while remaining compatible with the large number of CMs already installed.

Other changes have included the ability to “bond” adjacent channels to increase the channel bandwidths and thus the maximum data rates available to users.

The following sections briefly describe the physical (PHY) and MAC layers of the DOCSIS protocol.


Upstream PHY

The following block diagram shows the sequence of operations performed on the US data before it is transmitted by a DOCSIS 3.0 CM:



A Reed-Solomon FEC block code is used. This code uses elements from GF(256) rather than GF(2). This means the code operates on bytes rather than bits. The RS code will correct any number of errors in each byte. The number of parity bytes (not bits) ($n - k$) can be configured to be between 0 to 32 bytes per block (correctable errors, $t = 0$ to 16) and the number of data bytes per block (k) can be configured to be between 16 and 253.

The purpose of the byte interleaver is to spread out noise bursts so that if the burst is short enough relative to the interleaver depth the number of errors per codeword will be sufficiently small that the FEC code can correct them. It operates as a block interleaver where the bytes are written row-wise and read out column-wise. The interleaver width is set to the RS codeword length ($n = k + 2t$) and the depth is configurable with a maximum of 2048 bytes total memory.



Input sequence: $C_1(1) \dots C_1(N), C_2(1) \dots C_2(N), C_3(1) \dots C_3(N), \dots, C_n(1) \dots C_n(N)$
 Output sequence: $C_1(1), C_2(1) \dots C_n(1), C_1(2), C_2(2) \dots C_n(2), C_1(3) \dots C_n(3), \dots, C_1(N), C_2(N) \dots C_n(N)$

Figure 6-5 - Byte Interleaver Operation

Exercise 8: Assuming 256 bytes per codeword, what is the maximum number of interleaver rows? If an error burst starts with the first byte, what is the longest error burst that will result in 2 or fewer errors per codeword? t or fewer errors per codeword? How many errors will appear at the output of the decoder for error burst of this length or shorter?

A 15-bit LFSR-based additive scrambler is used.

The modulation symbol rates are 160 to 5120 kHz with QPSK and 8- through 256-QAM modulation.

Exercise 9: What are the minimum and maximum bit rates for these symbol rates and modulation formats, not including FEC or other overhead?

The synchronous CDMA (S-CDMA) mode uses orthogonal spreading sequences to allow multiple users to transmit simultaneously on the upstream channel. We will not cover the S-CDMA mode. DOCSIS version 3.1 also has an upstream OFDM mode which we will also not cover.

Downstream PHY

The downstream signal format was designed for transmission of digital video and predates the design of the upstream. The downstream is based on ITU standard ITU-T J.83. The downstream parameters are as follows (from ITU-T J.83):

Table B.3 – Cable transmission format

Parameter	64-QAM format	256-QAM format
Modulation	64-QAM, rotationally invariant coding	256-QAM, rotationally invariant coding
Symbol size	3 bits for "I" and 3 bits for "Q" dimensions	4 bits for "I" and 4 bits for "Q" dimensions
Transmission band	54 to 860 MHz (Note)	54 to 860 MHz (Note)
Channel spacing	6 MHz (Note)	6 MHz (Note)
Symbol rate	5.056941 Msps \pm 5 ppm (Note)	5.360537 Msps \pm 5 ppm (Note)
Information bit rate	26.97035 Mbps \pm 5 ppm (Note)	38.81070 Mbps \pm 5 ppm (Note)
Frequency response	Square root raised cosine filter (Roll-off \approx 0.18)	Square root raised cosine filter (Roll-off \approx 0.12)
FEC framing	42-bit sync trailer following 60 RS blocks (see B.5.3)	40-bit sync trailer following 88 RS blocks (see B.5.3)
QAM constellation mapping	6 bits per symbol (see B.5.5)	8 bits per symbol (see B.5.5)

NOTE: These values are specific to 6 MHz channel spacing. Additional sets of values for differing channel spacing are under study.

The following diagram, also taken from the standard, shows the signal processing involved:




Figure B.7 – Layers of processing in the FEC

It supports 64- and 256-QAM modulation at symbol rates of approximately 5.05 and 5.36 MHz (for 256-QAM). A Reed-Solomon (128,122) code over GF(128) (7-bit symbols) is used that can correct $t = 3$ errors per block.

A convolutional interleaver is used rather than a block interleaver:

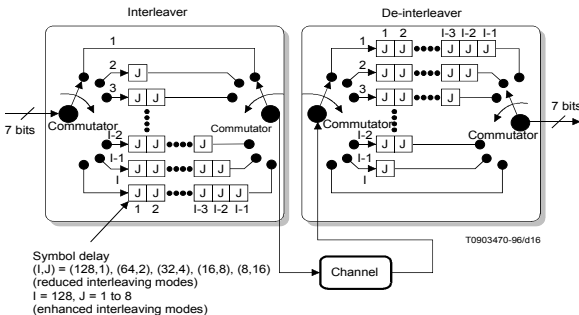


Figure B.8 – Interleaving functional block diagram

as well as a scrambler (called a “randomizer” in the spec) based on a LFSR operating on symbols from GF(128).

The modulator uses a technique called Trellis Coded Modulation (TCM) which uses a convolutional encoder to protect the bits in the constellation that are most likely to have errors. The use of two FEC codes (“inner” and “outer”) is called concatenated coding.

MAC Overview

The operation of the network is controlled by the CMTS which transmits continuously on one or more downstream (DS) channels. CMs receive from the CMTS and respond to CMTS commands by transmitting on one or more upstream (US) channels. The DOCSIS Medium Access Control (MAC) protocol controls transmissions by CMs.

The CMTS periodically transmits a MAP (“Upstream Bandwidth Allocation Map”) message allocating upstream “mini-slots” to the different CMs. The remaining un-allocated time slots can be used by CMs on a contention basis.

The “mini-slot” durations are a power of 2 (1, 2, 4, ... 128) multiples of 6.25 μ s, which is between 6.25 μ s (1 \times) and 800 μ s (128 \times).

In addition, the CMTS provides time and frequency synchronization and ranging (measurement of the distance to the CMTS) for CMs.

Exercise 10: Ignoring guard times, FEC and other overhead, how many bits could be transmitted at the lowest PHY rate during the shortest mini-slot? At the highest PHY rate during the longest min-slot?

DOCSIS MAC Layers

The DOCSIS Layer 2 (Data Link Layer) and Layer 3 (Network) protocol layers sit above the Physical layer (Layer 1) and are responsible for converting IP frames to/from frames that can be transmitted/received as a DOCSIS physical-layer signal.

DOCSIS uses the IEEE 802 Layer 2 and IETF IP Layer 3 protocols. This means that CMs have 48-bit MAC addresses and 32-bit IPv4 addresses and make use of the same bridging and routing algorithms as Ethernet LANs and TCP/IP networks.

The CM acts as a L2 bridge between the cable and Ethernet interfaces. This bridge can be configured with filtering rules by the service provider. For examples these rules would typically prevent bridging Ethernet traffic that was local to the customer’s network.

The following diagram, from the DOCSIS 1.0 specification, shows the various protocol layers involved and that will be discussed in more detail below:

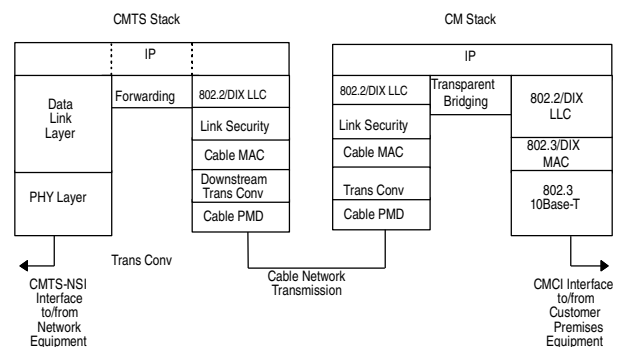


Figure 3-2. Data Forwarding Through the CM and CMTS

- the IP layer is used to interface to customer equipment (CPE) and for CM configuration and management.

- the 802.2/DIX frame format, used internally only, defines a L2 framing format similar to Ethernet (802.3).
- the “Cable MAC” is a DOCSIS-specific protocol layer described below.
- the “Trans Conv” layer represents the MPEG transport stream protocol in the downstream direction.
- the PMD (Physical Medium Dependent) layer is the “PHY” which was described in the previous lecture.

Upstream and downstream are on different frequencies so, unlike Ethernet, CMs do not receive each others’ frames. All broadcast communication including coordination to avoid collisions, has to be done by the CMTS.

The description below is a high-level overview of the DOCSIS 1.0 MAC layer functions. Later versions added backwards-compatible features to improve performance, particularly for the upstream and for constant-bit-rate data sources such as cable telephony.

Provisioning and Management

In addition to using IP to provide service to the customer premises equipment (CPE) such as computers and VoIP phones the CM also implements an internal L3 IP protocol stack which is used for configuration and management of the CM.

The CM first configures its IP stack using DHCP (dynamic host configuration protocol). It then uses a file transfer protocol such as TFTP (trivial file transfer protocol) to download a configuration file. This configuration file can specify service IDs, US and DS frequencies, security settings, bandwidth restrictions, packet filtering, etc.

Once the modem is in service the network management system can monitor CMs’ status (e.g. signal levels, bit rates, error rates, ...) using SNMP (Simple Network Management Protocol).

DOCSIS MAC Framing

In addition to the IP and 802 frames, DOCSIS defines an additional layer of framing. These DOCSIS MAC frames may themselves contain data frames (802 frames, for example) or they may contain DOCSIS-specific frames carrying timing information or management frames. These DOCSIS MAC frames include frames that describe the upstream channel and allow for request/response exchanges for registration, ranging (distance/delay measurement), encryption key distribution, etc. as well as the upstream channel allocation (MAP) frame described below.

Downlink Framing

The downstream has a single transmitter for each channel so there is no contention for media access.

The downstream PHY was originally designed to carry video in an MPEG transport stream consisting of a sequence of 188-byte frames containing a 4-byte header and 184-byte payload. DOCSIS uses frames whose MPEG header has a PID (packet ID) value of 0x1FFE to distinguish DOCSIS frames from video frames that might be transmitted on the same data stream.

DOCSIS MAC frames are split up to fit the fixed-length MPEG frames and may span multiple MPEG frames.

Exercise 11: List the protocol layers (headers) between the PHY and the payload of a downstream SNMP frame. *Hint: Assume SNMP uses UDP.*

Uplink Framing

The upstream has one receiver and multiple transmitters so the upstream channel time is divided into time slots that are allocated to different CMs’ uplink transmissions. The CMTS provides timing reference for upstream transmissions and allocates time slots to the CMs.

Some slots are left unallocated for contention-based transmission of requests for upstream channel time. The contention protocol is slotted Aloha – transmissions are aligned to start at mini-slot boundaries. An exponentially-increasing “contention win-

“window” size is used – the backoff is chosen randomly within a contention window that doubles for each unacknowledged transmission.

Various techniques (e.g. “piggybacking” and “concatenation”) are used to improve the efficiency of the upstream MAC for various types of traffic. We will not cover these.

Time slots are allocated to 14-bit “Service IDs” (SIDs). A CM may be allocated more than one SID. For example, one SID may be assigned for constant-bit-rate traffic for cable telephony and another for bursty data.

A DOCSIS MAC “MAP” frame describes the allocation of subsequent uplink channel time. The CMTS may allocate the uplink time for ranging and contention access as well as contention-free time allocations for each SID. The allocation for each SID is encoded as a 32-bit descriptor with a 14-bit SID, a 4-bit “interval usage code” and a 14-bit offset in units of mini-slots.

The following diagram, also from the DOCSIS 1.0 spec, shows the format of the MAP information:

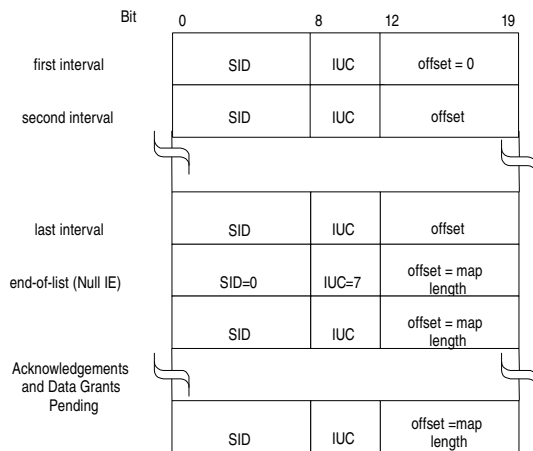


Figure 6-19. MAP Information Element Structure

Exercise 12: Assuming the maximum number of CMs per uplink is determined by available SIDs, each CM is allocated two SIDs and only SIDs from 0x0001 to 0x1FFF are available, how many CMs could be supported per uplink channel? What is the longest time offset that can be specified in a MAP frame assuming the length of a mini-slot is 25µs?

Uplink Ranging

When a CM first starts up it searches for a downstream channel. Once it has found a suitable channel

it must measure the delay between it and the CMTS so that upstream transmissions are properly timed and do not interfere with other upstream transmissions.

To do this the CM transmits a Ranging Request frame during a portion of the uplink channel time allocated specifically for these requests. The CMTS measures the arrival time, frequency offset and power level of the ranging request frame. This information is transmitted on the downlink in a Ranging Response frame which also allocates a temporary SID. The CM adjusts its transmit parameters and transmits further Ranging Request frames and uses the received Ranging Response frames to fine-tune the PHY parameters.