DOCSIS™ Cable Modem Technology

David Fellows, Pilot House Ventures
Doug Jones,YAS Corporation

ABSTRACT

An overview of the cable modem interface standard is presented. This standard is the foundation for high-speed access to content on the Internet being offered by North American cable operators. It also facilitates retail sales of modems and integration into computers, is being adopted by international operators, and provides a path for convergence of video, data, and voice services on a broadband cable system. The details of the protocol stack and evolution of the standard are discussed.

INTRODUCTION

At the end of 1995 a number of cable operators agreed to create, under the auspices of Cable Television Laboratories (CableLabs®), a set of Data Over Cable Service Interface Specifications (DOCSIS™)[1]. Two proprietary systems were in widespread use at the time, but neither was interoperable with the other. This was already creating inventory and service provision issues, and never would have allowed a cable operator's dream of having cable modems built into new computers and available at retail. The group set a goal of having such a set of specifications available in six months. They failed miserably at that, but succeeded eventually, as over 100 modem models have been certified as interoperable by CableLabs at this writing, from thirty-six different companies. Over two million modems have been shipped to customers, and another 10 million are expected in 2001.

The premise behind DOCSIS is having interoperable cable modems built specifically for cable operator business plans, which at the time was best-effort connections of Ethernet-connected computers to the World Wide Web. The main specification work for DOCSIS 1.0 was completed in March 1997. The first prototypes came into CableLabs for testing that summer.

A cable system consists of a headend, where video signals are collected and modulated in a frequency stack known as the channel lineup. This signal is sent to every home in the franchise area over coaxial cable. In the last 10 years, an architecture known as hybrid fiber coax is typically used, where fiber cables, in a star arrangement, carry the cable signal to pockets of homes (between 100 and 2000), and coax carries the signal the rest of the way to the home. A cable data system consists of multiple cable modems (CMs), in subscriber locations, and a cable modem termination system (CMTS), all connected by a CATV plant. The CMTS can reside in the headend, distribution hub, or some other location depending on the plant layout. In a two-way system, lower frequencies carry signals from the home to the headend (the upstream direction), and higher frequencies carry signals from the headend downstream to the home. Note that the final coaxial link to the home is a bus architecture — the signals, which go into one home, also go into the neighboring homes.

DOCSIS PROTOCOL STACK

The protocols chosen for DOCSIS, from the modulation schemes up to the networking protocols, all have specific functions. The relationships between these functions are shown in a protocol stack.

Figure 1 shows such a stack for a DOCSIS cable modem. The lower four layers are specific to multiple system operator (MSO) cable data networks in that they are only present between the CM and CMTS over the cable network. The higher-layer protocols, IP and above, are carried by the DOCSIS layers across the cable network and are used for communication with the Internet.

The CM performs the lower four layers shown in Fig. 1. The CM will receive an IP packet from host customer premises equipment (CPE), typically over Ethernet, and will add link encryption, mediate access to the return path, and finally modulate the data onto the cable network. On the forward path, the CMTS also adds an MPEG-2 framing layer.

Above the DOCSIS protocol layers is the Internet Protocol (IP) layer. This protocol is the "glue" that binds together the current Internet as we know it. The layers shown at the very top of Fig. 1 are standard Internet service protocols (mail, news, Web, etc.). These protocols are out of the scope of this discussion; however, DOCSIS was specifically designed for use with these

0163-6804/01/$10.00 © 2001 IEEE

Internet protocols, and hence is expected to grow as the Internet grows.

The lower four layers of DOCSIS are:

- The physical layer includes the modulation schemes used on the coax network. For DOCSIS, these are 64-quadrature amplitude modulation (QAM) and 256-QAM on the forward, quadrature phase shift keying (QPSK) and 16QAM on the reverse.
- The MPEG-2 Transmission Convergence layer (present in the downstream only) means DOCSIS CM data is encapsulated in 188-byte MPEG-2 frames. This allows the data to be multiplexed with other MPEG streams on the same carrier on the forward path. An example would include sending MPEG-2 video and audio on the same carrier as MPEG-2 encapsulated DOCSIS data.
- The media access control (MAC) layer controls CM access to the return path. Since there can be many CMs trying to transmit simultaneously on the return path, the MAC protocol provides an orderly method for the CMTS to tell a CM when it gets to transmit and for how long.
- The data link encryption layer protects user data. Since the cable network is shared, there has to be a method to protect user data from malicious users. DOCSIS has defined the Baseline Privacy Interface (BPI) for this function.

PHYSICAL LAYER PROTOCOL

The physical layer protocol describes the modulation formats used on the cable network, on both the forward and reverse paths. The starting point for the design was to describe both the forward and return path channel models that are essentially the worst case RF channels for a DOCSIS cable data system.

In the forward direction, the cable system is assumed to have a passband with a lower edge at 50 MHz and an upper edge, which is implementation-dependent but typically in the range of 300–860 MHz. Within that passband, NTSC analog television signals in 6-MHz channels are assumed to be present on either of the standard, HRC, or IRC frequency plans (EIA Interim Standard IS-6) as well as other narrowband and wideband digital signals. In the upstream direction, the cable system may have a subsplit (5–30 MHz) or extended subsplit (5–42 MHz) passband. NTSC analog television signals in 6-MHz channels, as well as other signals, may be present on the reverse path.

A CM, configured with at least one set of defined physical layer parameters — modulation, forward error correction (FEC), symbol rate, and so on — from the range of configuration settings described in DOCSIS, must be capable of operating with a packet loss rate of less than 1 percent (assuming 1500-byte packets) while forwarding at least 100 packets/s on cable networks having characteristics as defined in the specification.

Forward Path Modulation Format — The downstream physical layer is based on the International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) Recommendation J.83B, "Digital Transmission of Television Signals, Annex B." This revision of ITU-T J.83B includes not only the original 64-QAM modulation and a fixed-depth interleaver used to deliver digital video, but also includes 256-QAM for higher downstream channel data rates as well as a variable-depth interleaver. DOCSIS-compliant downstream channels may occupy any 6 MHz band between 88 MHz and 860 MHz.

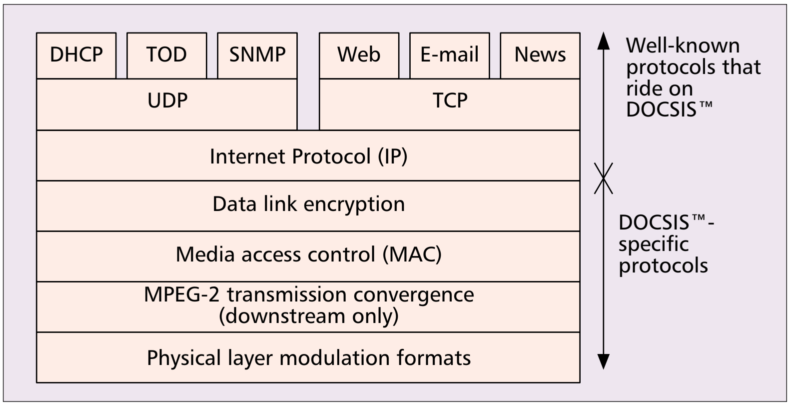

Figure 1. *The DOCSIS™ cable modem protocol stack.*

Parameter	Value
Center frequency	91–857 MHz ±30 kHz
Level range (one channel)	–15 dBmV to +15 dBmV
Modulation type	64-QAM and 256-QAM
Symbol rate (nominal)	5.056941 Msym/s (64-QAM) and 5.360537 Msym/s (256-QAM)
Bandwidth	6 MHz
Total input power (40–900 MHz)	<30 dBmV
Input (load) impedance	75 ohms
Input Return Loss	> 6 dB (88–860 MHz)

Table 1. *Electrical input to the CM.*

64-QAM versions of these downstream channels have been widely deployed with digital set-top boxes, and are successfully delivering digital video services to cable subscribers. 256-QAM versions of these downstream channels have been proven in extensive rigorous tests; this technology is ready for use. The reliability of QAM-modulated downstream channels is helped by the powerful concatenated FEC provided by ITU-T J.83B. Multiple layers of error detection and error correction, coupled with variable-depth interleaving to provide variable-length burst error resilience, deliver error rates ensuring customer satisfaction. The high data rates together with the low error rates provide a bandwidth-efficient delivery mechanism for digital data delivery.

Forward Channels — The CM must accept an RF modulated signal with characteristics as described in Table 1. The output signal level of

I (# of Taps)	J (Increment)	Burst protection 64-QAM/256-QAM	Latency 64-QAM/256-QAM
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/16 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

Table 2. *Interleaver characteristics.*

the CMTS is variable over the range of 50–61 dBmV.

Variable-Depth Interleave — The forward path modulation formats support a variable-depth interleaver with the characteristics defined in Table 2.

Variable-depth interleaving is a form of error protection. Based on the level of the interleaving, the data on the forward path is protected from noise burst lengths as mentioned in the table. Even with reduced noise margins, the downstream channels are designed to deliver 64-QAM signals with a bit error rate (BER) of less than 10^{-8} at a carrier to noise (C/N) ratio of 23.5 dB. 256-QAM channels can be expected to deliver a similar BER at 30 dB C/N. This translates to 1 bit error/3–5 s. Downstream channels are shared between subscribers; therefore, errors will be distributed randomly to different subscribers. An additional benefit from the strength of this FEC is that it permits operation of the downstream digital data channels 10 dB lower than the nominal level of video carriers on the system. This helps minimize system loading while still delivering robust digital data services.

One of the side effects of the interleaver is added latency in the downstream channels. The process of interleaving the outgoing symbols, shuffling the position of the symbols so that normally adjacent related symbols are now separated by unrelated symbols which would otherwise be transmitted later, delays the delivery time of related symbols. The benefit is that a burst of noise damages only unrelated symbols. The FEC can then correct the damaged symbols once they are reshuffled back into their normal order as long as the burst damage did not span too many related symbols. There is an intrinsic relationship between the depth of the interleaving and the latency incurred by the interleaving. The deepest interleaving depth available in the DOCSIS RF specification provides 95 µs burst protection at the cost of 4 ms of latency. Four milliseconds of latency is insignificant when watching digital video or using standard Internet data services such as Web browsing, e-mail, and file transfer. However, near-real-time constant bit rate services (e.g., IP telephony) that have tight end-to-end latency requirements may be impacted.

The variable-depth interleaver enables the system engineer to trade between how much burst error protection is required and how much latency can be tolerated by the services being delivered. The depth of the interleave can also be dynamically controlled by the CMTS based on RF channel conditions.

Return Path Modulation Format — DOCSIS-compliant upstream channels provide for both frequency-division multiple access (FDMA) and time-division multiple access (TDMA). FDMA is provided by the ability to have multiple upstream channels simultaneously supporting multiple modems. Return path channels reside in the spectrum between 5 MHz and 42 MHz.

TDMA is provided by the use of "slotting" on the upstream channels. Each upstream channel is divided into equal-time segments called mini-slots. A mini-slot is nominally 16 bytes long, but can be up to 128 bytes. The use of each mini-slot is controlled by the CMTS. The CMTS assigns contiguous intervals of mini-slots to individual CMs, or makes them available for contention by groups of CMs, or opens them up for contention by all modems. Additionally, the type of communication within the assigned interval is dictated by the CMTS. All DOCSIS-compliant CMs will time-coordinate all their upstream transmissions so that they only transmit within the appropriately allocated mini-slots. This provides the mechanism for multiple access in the time domain.

Return Channels — The CM modulator must provide both QPSK and 16-QAM, at symbol rates and channel widths described in Table 3.

The CM return path modulator must operate over the following power ranges:

- QPSK: +8 to +58 dBmV
- 16-QAM: +5 to +55 dBmV

As can be seen in Table 3, each upstream channel has an assigned bandwidth associated with it, and the occupied bandwidth is directly related to the data rate of the channel. DOCSIS-compliant upstream channels occupy bandwidths of 200, 400, 800, 1600, or 3200 kHz. This corresponds to channel data rates of 160, 320, 640, 1280, or 2560 ksym/s. The specification provides for both QPSK and 16-QAM transmissions upstream, which allow either 2 b/sym or 4 b/sym, respectively. Therefore, channel data rates are available between 320 kb/s (QPSK at 160 ksym/s) and 10.24 Mb/s (16QAM at 2560 ksym/s). While the bandwidth is fixed for any upstream channel because the symbol rate is defined for that channel, the bit rate on the channel is variable because independent transmissions on that channel can be either QPSK modulated or 16-QAM modulated.

Forward Error Correction — The DOCSIS flexible upstream FEC coding enables the system operator to set the size of the error protected data blocks and to set the number of correctable errors within each block. In previous proprietary CM systems, when impairments in a data channel caused too many errors, the only solution was to abandon that frequency and hop the channel to a (hopefully) cleaner portion of the spectrum. While DOCSIS-compliant systems can operate this way, the flexible FEC coding option enables the system operator to choose to stay on the same frequency by dynamically changing the error protection on that channel. Even though the additional few bytes of error protection reduces the channel information rate a small amount, it allows the upstream spectrum to be managed more efficiently.

TRANSMISSION CONVERGENCE PROTOCOL (MPEG-2)

A transmission convergence (TC) protocol is one that allows multiple types of services to share the same RF carrier. For DOCSIS, the TC layer is MPEG-2, a protocol seeing more and more use in CATV networks. Using MPEG-2 means that other types of information also encapsulated with MPEG-2 (e.g., voice and video) can all be sent on the same RF carrier as DOCSIS CM data. That is, voice, video, and data can all be sent in the same forward path channel as long as they are encapsulated with MPEG-2.

MPEG-2 provides a mechanism to identify individual packets within an MPEG-2 stream such that a CM or set-top box can identify which packets it should decode. This mechanism, called a program identifier (PID), is present in all MPEG-2 frames. DOCSIS has declared the value 0x1FFE to be the well-known PID for all CM traffic on that channel. DOCSIS CMs will only operate on MPEG packets with this PID.

In addition, MPEG-2 provides a frame structure that facilitates channel lock. MPEG-2 frames start every 188 bytes with a synchronization byte. Searching for this MPEG sync byte, which repeats at a regular and fairly fast interval, makes it easier to synchronize with the overall bit stream on the channel. Since the CM may have to search several forward path channels for DOCSIS data, any means that helps find the proper channel sooner is beneficial.

MAC LAYER PROTOCOL

The MAC layer protocol controls access to the return path. Since several modems may have data to transmit at any given time, the MAC protocol provides a means for the CMTS to indicate which modem can transmit when and for how long. If CMs transmit individually, data "collisions" won't occur on the return path and the system will operate efficiently. Collisions cause CMs to retransmit data, which is inefficient.

The DOCSIS MAC protocol uses a request/grant mechanism. Using this type of protocol, a CM requests from the CMTS an opportunity to transmit a certain amount of data. As the CMTS receives requests from all the modems with data to transmit, it reserves mini-slots (essentially transmit opportunities) on the return path accordingly. Periodically (every few milliseconds), the CMTS sends a message to the CMs (over the forward path) indicating the specific mini-slots granted to each modem. As a result of reserving bandwidth, modems are guaranteed a collision-free interval in which to transmit. The CMTS is responsible for allocating bandwidth to modems based on the types of service provided to the user.

In order for a modem to request bandwidth without first having been allocated a time to send a message, the CMTS periodically allocates a portion of the return channel for any modem to send a request. These requests may collide with a request from another modem; however, the CMTS can vary both the length and frequency of the contention request intervals to keep the number of collisions to a minimum. If a collision is detected, the CMs that sent the messages will back off random intervals and then retransmit their requests.

As an example of the intricacies of working on a cable plant, the following sections describe what the MAC layer of the modem goes through after being connected to the network, before lighting the green LED.

Symbol rate (ksym/s)	Channel width (MHz)[1]	QPSK rate (Mb/s)	16-QAM rate (Mb/s)
160	0.20	0.32	0.64
320	0.40	0.64	1.28
640	0.80	1.28	2.56
1280	1.60	2.56	5.12
2,560	3.20	5.12	10.24

■ **Table 3.** *Maximum channel width.*

Channel Acquisition — The modem first scans for a downstream channel, obtains QAM lock, and finds MPEG packets with the DOCSIS well-known PID. A CM may have to search several QAM channels before finally finding one with DOCSIS data.

Channel acquisition is successful when the CM is decoding MPEG-2 frames that bear the DOCSIS well-known PID. The MPEG-2 framing will be stripped away and the resultant MAC frames passed to the MAC layer for processing.

Obtain Upstream Parameters — During this phase, the CM is looking for three MAC messages that are repeatedly sent by the CMTS on all DOCSIS downstream channels.

The first such message is the time synchronization message (SYNC). The SYNC message is sent by the CMTS to provide a common time reference to all modems.

The next such message is the upstream channel descriptor (UCD). The CM must find a UCD that describes an upstream channel that matches its own capabilities. For instance, if the CM can only transmit at certain symbol rates or is limited in its frequency range and modulation types, the CM may not be able to transmit on a particular upstream channel.

The final such message to find is an upstream bandwidth allocation map (MAP) that matches the upstream channel described in the UCD. The MAP message contains the mini-slot information that indicates when a modem can transmit and for how long. (The SYNC provides the time reference for these transmissions.) For its initial transmission, the CM will look in the MAP for a transmit opportunity reserved for modems just connecting to the network, called the *initial maintenance* transmit opportunity.

The ranging process is first performed during an initial maintenance transmit opportunity when a CM is first connected to the network. Once booted, the ranging process is repeated for each CM at regular intervals during periodic maintenance opportunities that the CMTS schedules.

At the conclusion of obtaining upstream parameters, the modem has:

- A rough time reference
- Knowledge of the details of an upstream transmit channel
- Knowledge of initial maintenance transmit opportunities for newly connected cable modems on that upstream channel

These three pieces of information are all used for the next process.

Ranging — Ranging describes the next three processes the CM must accomplish. These include:

- Fine adjust on the time reference
- Fine tuning the transmit frequency
- Fine adjust on the transmit power

Since each CM is a unique distance from the CMTS, each will have unique settings for these parameters.

To begin the ranging process, the CM transmits a ranging request message to the CMTS during an initial maintenance opportunity (as described by MAP messages) on an upstream channel. The CM begins transmitting this message at the starting boundary of the initial maintenance opportunity, based on its rough time SYNC and its interpretation of the MAP. Upon receipt of this message, the CMTS sends a ranging response message addressed to that CM.

If the CM does not receive a ranging response message from the CMTS within a timeout period, one of two things may have happened. First, since the initial maintenance opportunity is available to any CM just attaching to the network, ranging request messages may collide. Second, the power level at which the modem is transmitting may be too low to be detected at the CMTS. Therefore, if the CM does not receive a ranging response message, it will both increase transmit power and wait a random number of initial maintenance opportunities before sending another ranging request.

In preparing the ranging response message, the CMTS notes:

- The time offset of when it received the ranging request from the actual beginning of the initial maintenance transmit opportunity
- The exact frequency of the transmission
- The incoming power level of the transmission

Based on these data, the CMTS determines corrections and sends these back to the CM in the ranging response message. The CM adjusts its parameters based on these corrections and transmits a second ranging request to the CMTS. The CMTS again returns a ranging response with time, frequency, and power corrections, as needed. This process continues iteratively until the CMTS is satisfied with the timing, frequency, and power settings being used by the CM. When finished, the timing is synchronized to within less than 1 μs, transmit frequency to within 10 Hz, and transmit power to within 1/4 dB.

The ranging process is first performed during an initial maintenance transmit opportunity when a CM is first connected to the network. Once booted, the ranging process is repeated for each CM at regular intervals during periodic maintenance opportunities the CMTS schedules. Periodic tuning of timing, transmit frequency, and transmit power ensures continued reliable communications between the CMs and the CMTS.

Time synchronization is critical to the function of the MAC protocol. Time synchronization requires that the two-way round-trip transmission delay be negated because all upstream transmissions must align with the mini-slot timing as viewed at the CMTS. Communication is delayed in the downstream direction by latency, downstream propagation delay, and processing overhead. When the CM sends a transmission upstream, it is delayed by upstream propagation and processing overhead. The ranging process effectively offsets the sum of these delays. The timing offset actually causes the CM to transmit earlier than the assigned mini-slot time (as measured at the CMTS) in order to negate delays caused by interleaving latency in the downstream, propagation in the system, and fixed processing overhead in both the CMTS and CM.

All further transmissions on the return path take place based on requests from the CM to the CMTS, and grants from the CMTS to the CM. Whenever a modem has data to send, it will request bandwidth on the return path during a request region (as defined on a MAP). The CMTS will grant this request and schedule that modem a unique opportunity to transmit the data. The MAP request slot is unique to a modem, but if a modem has not been assigned a slot, it will use a slot in a contention region.

Initial ranging can be considered the DOCSIS "analog" to AutoDiscovery. While this term is not used by DOCSIS, AutoDiscovery is essentially an opportunity given the modem to transmit without first being given a grant to transmit. The DOCSIS MAC allows this by offering periodic initial maintenance transmit opportunities on the return channel that can be used by any modem just connecting to the network.

IP Layer Establishment — Once timing, frequency, and power are set, the modem now must establish the IP protocol. It does so by invoking the Dynamic Host Configuration Protocol (DHCP) which causes the CM to be assigned an IP address. DHCP runs between the CM and a DHCP server administered by either the MSO or some other service provider. As long as that CM is active, it will be "leased" an IP address to use. After a period of inactivity, that IP address can be reclaimed and given out to another active modem, thereby conserving IP address space.

Registration — Registration begins with the CM downloading a configuration file. The IP address of a configuration file server and the name of the configuration file the modem is to download are both included in the DHCP response to the modem. The CM uses the Trivial File Transfer Protocol (TFTP) to download the configuration file from the server. The configuration file contains information the CM uses to operate, such as how much bandwidth it is allowed to use and which services it is allowed to provide. These service-provisioning items are taken when the subscriber first calls in to request

CM service. This information is fed into a back-end system to create the configuration file.

During the final phase of registration, the modem sends a registration message to the CMTS confirming the configuration file it received. The CMTS also retrieves a copy of the configuration file from the configuration file server. The CMTS compares the file from the server with the data from the CM to ensure the modem will only be using services for which it was authorized. Only after the configuration file data is cross-checked is the CM finally allowed to transmit real user data onto the network.

It is at this time that the green LED comes on!

DATA LINK ENCRYPTION

Once the modem is registered and sending customer data, DOCSIS provides for data link encryption (encryption between the CM and CMTS only). Known as Baseline Privacy Interface (BPI), the specification describes a simple data privacy function. BPI is intended to provide the minimum level of encryption support needed to protect user data from theft. The goal of BPI is to provide CM users with data privacy across the RF network by encrypting traffic flows between a CM and the CMTS.

Since data privacy is the principal service goal of BPI, and given neither CM nor CMTS authentication are prerequisite for providing user data privacy, the BPI encryption key distribution protocol does not authenticate either the CM or CMTS (i.e., it does not employ authentication mechanisms such as passwords or digital signatures). In the absence of authentication, BPI provides basic protection of service by ensuring that a modem, uniquely identified by its 48-bit IEEE MAC address, can only obtain keying material for services it is authorized to access. Since BPI does not authenticate CMs, it cannot protect against an attacker employing a cloned modem (i.e., one masquerading as an authorized modem). Other methods are available to detect such clones.

BPI uses the Cipher Block Chaining (CBC) mode of the Data Encryption Standard (DES) algorithm [2] to encrypt data in both upstream and downstream frames. CMs use the protocol to obtain authorization and encryption keys from the CMTS, and to support periodic reauthorization and changing the keys, which are single DES (1-DES). The key management protocol uses RSA, a public-key encryption algorithm [3], and the Electronic Codebook (ECB) mode of DES [2] to secure key exchanges between the CM and CMTS. CMs must have factory-installed RSA private/public key pairs, or provide an internal algorithm to generate such key pairs dynamically. The frequency of change of the shared secret keys is an operator-determined parameter. These keys can be established for very long periods of time (e.g., many months to weeks) or changed as often as every 30 s.

DOCSIS 1.1

In the spring of 1998, a program was initiated to add functionality for new systems. This effort has been named DOCSIS 1.1. The business plans of the cable operators had moved beyond best-effort Web cruising to include services such as voice over IP. The main DOCSIS 1.1 specifications were published in March 1999.

In the spring of 1998, a program was initiated to add functionality for new systems. This effort has been named DOCSIS 1.1. The business plans of the cable operators had moved beyond best-effort Web cruising to include services such as voice over IP.

DOCSIS 1.1 enhances the existing DOCSIS 1.0 services. Several areas were addressed, including:

- Quality of service (QoS)
- IP multicast
- Security
- Operational support

With the specification effort for DOCSIS 1.1 substantially complete, CableLabs is turning its attention to the DOCSIS 1.1 test effort. Certification Wave (CW) 16, held in the fourth quarter of 2000, was the first time a DOCSIS 1.1 product was tested for compliance (this testing is continuing into CW 17, in the first quarter of 2001).

QUALITY OF SERVICE

DOCSIS 1.0 QoS is limited to providing a tiered best effort service. That is, all data packets from a modem are given the same priority and are throttled to the same throughput limits. A CM can be assigned maximum throughputs on both the upstream and downstream. On a lightly loaded network, data through the CM is throttled to the maximum. On a heavily loaded network there is no guarantee the maximum will be met.

Consider the following DOCSIS 1.0 example. Two people are working through a single CM; one is Web surfing, the other doing file transfer. The CM has a maximum downstream bandwidth of 1 Mb/s. The Web surfing and file transfer combine against the 1 Mb/s downstream limit; therefore, downloading Web material takes bandwidth away from file transfer, and vice versa. DOCSIS 1.0 does not define how to differentiate between Web surfing and file transfer to give them separate bandwidth limits.

DOCSIS 1.1, on the other hand, includes methods to guarantee QoS to different services going through the same CM or on the same RF channel. These new methods include:

- Packet classification — Packets are classified based on where the traffic originated, where its going, what kind of traffic it is, and so on. Once classified, packets are then assigned to service flows.
- Service flows — Flows define the specific guarantees and limits of types of trafficct. Flows define maximum and minimum throughputs, latency, jitter, tolerances, and so on.
- Dynamic service assignment — This allows the system to dynamically create and delete service flows, although static service flows can also be created at the time of initial provisioning.

A CM can be assigned multiple service flows, and in this way can give differentiated treatment to different types of Internet traffic. That is, certain traffic can be given differentiated priorities based on the type, origin, destination, and so forth of that traffic.

Extending the DOCSIS 1.0 example to DOCSIS 1.1, Web surfing would count against one traffic limit, and the file transfer traffic could count against another traffic limit, even though

The protocol that controls access to IP multicast services is IGMP. DOCSIS 1.1 defines specific rules for how users can access multicast services based on IGMP, and how the CM should treat these requests.

both are using the same CM. These traffic limits could be guaranteed, regardless of how busy the network is. Based on these simple tools, DOCSIS 1.1 QoS tools provide many methods to offer new services to users.

MAC LAYER FRAGMENTATION

Quality of service goes a long way in allowing for differentiated service. But in a heavily loaded cable data network, QoS alone may not be enough to ensure that service guarantees are met. Cable networks are special due to the potential of the return path being shared by many users. When return path traffic is heavy, it makes sense to limit the size of data packets to ensure the CMTS can schedule transmit opportunities for certain CMs when those opportunities are needed in order to maintain service guarantees. In order to limit the size of upstream data packets, MAC-layer fragmentation gives the CMTS the option to instruct a CM to fragment a large upstream packet into multiple smaller packets. Each of the smaller packets is then scheduled its own individual transmit time. By breaking up large data packets, the CMTS has more flexibility in scheduling return path transmissions for other modems.

ENHANCED IP MULTICAST

DOCSIS 1.1 has been enhanced to efficiently handle the Internet Group Membership Protocol (IGMP) and related IP multicast services. Multicast packets are sent only once, but can be received by multiple users. Since they are only sent once, they consume a lot less bandwidth than if they were sent to each user individually. Examples of services that can be delivered with IP multicast are audio, video, stock tickers, news headlines, and weather. Multicast provides a bandwidth-efficient method to offer these services to subscribers.

The protocol that controls access to IP multicast services is IGMP. DOCSIS 1.1 defines specific rules for how users can access multicast services based on IGMP, and how the CM should treat these requests.

BASELINE PRIVACY PLUS

Baseline Privacy Interface Plus (BPI+) enhances the existing BPI service defined for DOCSIS 1.0. BPI+ was designed specifically to handle the new services defined in DOCSIS 1.1. Specifically, BPI+ describes how data encryption works when CMs are able to do different services. This includes upstream packets that may have been fragmented. Also, BPI+ describes how to perform secure IP multicast services. Finally, BPI+ adds certificate-based authentication to DOCSIS cable data systems. Using certificate-based authentication, BPI+ provides strong security from modem clones and theft of service.

EURODOCSIS

The DOCSIS specification has been adopted by many cable operators in Europe, and has been adapted for use on their cable systems to form a specification known as EuroDOCSIS. The main differences are in the modulation format, where ITU-T J83 Annex A is used, a different frequency split is used to separate upstream and downstream directions, and slightly different management information bases (MIBs) are used to describe these different parameters. An entity known as tComLabs has been formed to test for EuroDOCSIS interoperability and specification compliance, and the first modems and manufacturers have passed certification.

FUTURE DIRECTIONS

The DOCSIS specification will continue to evolve as technology and the market requires. A great deal of work has been done on CPE computer-controlled modems (CCCMs). This specification would allow the CPU of a personal computer to perform some of the functions of the DOCSIS CM, permitting lower-cost implementation. Special specifications were required so that the resulting device could not be used as a pirate box in a cable system.

Two proposals have also been made to DOCSIS for advanced physical layers, both outgrowths of work done in the standards body IEEE 802.14. One, known as frequency-agile TDMA, is a relatively straightforward extension to current DOCSIS. The other, known as synchronous code-division multiple access, is hoped to provide a more powerful toolset for ensuring reliable transmission at higher bit rates in a cable environment. Testing is underway on these techniques to determine if they will be included in future DOCSIS releases. In addition, a specification for one-way cable plants, which uses a telephone channel as the return path, has been written, and such modems have been deployed.

Evolution of the interface between the CM and the CPE is also allowed. The initial specification was for Ethernet, but USB and direct PCI plug-in cards have been specified and certified. Modems with built-in firewalls and other functionality, and with home networking interfaces (HPNA, HomeRF, 802.11, etc.) have also been tested.

A significant technical challenge facing cable operators is the ability to allow a customer to choose from different Internet service providers (ISPs) — because either the government requires it, a cable company has agreed to it, or customers want it. Currently it is thought that DOCSIS, because it is an interface specification, is not directly affected by this movement, but the DOCSIS effort is ready to become involved in modifying the interfaces if that turns out to facilitate the solution.

Finally, it is worth mentioning work collectively known as PacketCable™ (or IP Cablecomm in the standards arena). This work includes the specifications needed for operation of streaming services (especially voice over IP).

CONCLUSION

The widespread cooperation of vendors, cable operators, and technical experts has made the DOCSIS set of specifications a success. DOCSIS modems are being deployed by cable operators at a rate of approximately 1000 per week. DOCSIS 1.1 will allow voice and other streaming services to be deployed over cable systems, and the technology is being deployed in Europe as well as other parts of the world. In addition to digital video, the high-speed access to the Internet enabled by DOCSIS is becoming synonymous with broadband.

DOCSIS 1.1 will allow voice and other streaming services to be deployed over cable systems, and the technology is being deployed in Europe as well as other parts of the world.

REFERENCES

[1] The DOCSIS references are available at www.cablemodem.com, and include:
- Data-Over-Cable Service Interface Specifications, Radio Frequency Interface Specification, SP-RFI-I05-991105.
- Data-Over-Cable Service Interface Specifications, Cable Modem to Customer Premise Equipment Interface Specification, SP-CMCI-I02-980317.
- Data-Over-Cable Service Interface Specifications, Baseline Privacy Interface Specification, SP-BPI-I02-990319
- Data-Over-Cable Service Interface Specifications, Radio Frequency Interface Specification, SP-RFIv1.1-I05-000714
- Data-Over-Cable Service Interface Specifications, Operations Support System Interface Specification, SP-OSSIv1.1-I02-000714
- Data-Over-Cable Service Interface Specifications, Baseline Privacy Plus Interface Specification, SP-BPI+-I05-000714

[2] The Data Encryption Standard is described in the following documents:
- U.S. National Bureau of Standards, "Data Encryption Standard," Federal Information Processing Standard Publication 46, Jan. 1977.
- U.S. National Bureau of Standards, "Data Encryption Standard," Federal Processing Standard Publication 46-1, Jan. 1988.
- U.S. National Bureau of Standards, "Guidelines for Implementing and Using the Data Encryption Standard," Federal Information Processing Standard Publication 74, Apr. 1981.
- U.S. National Bureau of Standards, "DES Modes of Operation" Federal Information Processing Standard Publication 81, Dec. 1980.

[3] The RSA approach and algorithms are described in these references:
- R. Rivest, A. Shamir, and L. M. Adleman, "Cryptographic Communications System and Method," U.S. Patent 4,405,829, 20 Sept. 1983.
- R. Rivest, A. Shamir, and L. M. Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems," *Communi. ACM*, vol. 21, no. 2, Feb. 1978, pp. 120–26.
- R. Rivest, A. Shamir, and L. M. Adleman, "On Digital Signatures and Public-Key Cryptosystems," MIT Lab. for Comp. Sci., tech. rep. MIT/LCS/TR-212, Jan. 1979.
- RSA Laboratories, "The Public-Key Cryptography Standards," RSA Data Security, Inc., Redwood City, CA.
- RSA Laboratories, "PKCS #1: RSA Encryption Standard. Version 1.5," Nov. 1993.
- B. Schneier, *Applied Cryptography*, 2nd ed., New York: Wiley, , 1996.
- NIST, FIPS PUB 180-1: Secure Hash Algorithm Standard, Apr. 1995.

BIOGRAPHIES

DAVID FELLOWS (DFellows@PilotHouse.com) and two partners manage a venture capital fund, and he consults and is a board member of several companies. Prior to this, he was the interim CTO of RoadRunner, the high-speed data joint venture between Time Warner and MediaOne. He has also served as chief technology officer of US West Media Group, and senior vice president of engineering and technology for Media One, and prior to that for Continental Cablevision. Prior to joining Continental, he worked for Scientific-Atlanta as president of the Transmission Systems Business Division. He was vice president of technology for Siemens Transmission Systems. He is chair of the SCTE Digital Standards Subcommittee and head of CableLabs' DOCSIS certification board. He has been named *Telephony* magazine's Fiber Optic Visionary and *CED* magazine's Man of the Year for 1996. He also serves on the Board of Directors of Anadigics, BigBand Networks, Inc., and Celox Networks and the Board of Advisors for Broadcom Corporation, Mirror Worlds, Emperative, and Compaq. He holds two patents and is a trustee for the North Shore Music Theatre, and is vice-chair of the Friends of Harvard Rowing. He graduated from Harvard College with an S.B. degree in engineering and applied physics, holds a Master's degree from Northeastern University, and completed the PMD program at Harvard Business School. He was a member of seven U.S. rowing teams, captain of the undefeated intercollegiate championship Harvard rowing team, captain of the 1976 Olympic rowing team, and won silver and bronze medals at the Pan-American games in 1979 and 1975, respectively.

DOUG JONES (doug@yas.com) is a chief architect with YAS Corporation, Andover, Massachusetts. YAS Corporation provides venture capital financing and strategic and technical consulting services, all focused on broadband technologies and services. He has been part of the DOCSIS cable modem effort since the beginning in 1996, and currently acts as a consultant to CableLabs supporting DOCSIS and other projects there. Prior to YAS, he worked with MediaOne on DOCSIS cable modem deployments and with U S WEST on their video on demand trial in Omaha, Nebraska. He began his career with Bellcore on 5ESS switch systems engineering. In all, he has over 14 years in the communications industry. He holds M.S. and B.S. degrees in electrical engineering from the University of Texas in Austin (1986), and a B.S. degree in chemistry from the State University of New York at Albany (1982).