



The Wideband Protocol for a DOCSIS® Network

By:

John T. Chapman
Distinguished Engineer
Cisco Systems

jchapman@cisco.com

March 7, 2005

Originally published on Jan 14 for Emerging Technologies 2005
Reprinted and revised for VON / Fast Net FUTURES Spring 2005



Table of Contents

Introduction	5
Overview	5
Data Capacity Potential of Today's HFC Plant	6
Price and Density Trend Lines	9
The Wideband Protocol	10
The Concept	11
<i>DOCSIS 1.1 and 2.0 Systems with a Wideband Overlay</i>	11
<i>DOCSIS 3.0 with a Modular CMTS System</i>	12
Wideband in the Downstream	13
<i>The Downstream Wideband Channel</i>	13
<i>QAM Adjacency</i>	15
<i>Wideband Signaling Channel</i>	16
<i>Timing</i>	16
<i>Skew</i>	17
Wideband in the Upstream	18
<i>The Packet Streaming Protocol</i>	18
<i>The Packet Streaming Queue</i>	19
<i>Comparison with the Traditional DOCSIS Upstream</i>	20
<i>Configuration</i>	21
<i>Extensions to the Basic Idea</i>	22
Wideband Signaling	22
<i>Wideband Channel Descriptor</i>	22
<i>Registration</i>	23
<i>Addressing</i>	23
<i>Security</i>	23
<i>MIBs</i>	24
<i>Wideband Signaling Summary</i>	24
Wideband Deployment	25
WCMTS System Configuration	25
WCMTS	25
<i>Edge QAM Device</i>	26
<i>High Availability</i>	26
WCM Design	27
The Big Pipe Strategy	28
<i>What About DOCSIS Load Balancing?</i>	29
Wideband Migration Strategies	30
Wideband in the Marketplace	32
The DSL and FTTx Threat	32

Closing Remarks	35
5 Year Goal	35
<i>Big Pipes</i>	35
<i>What about the Asymmetry?</i>	36
Impact to the DOCSIS Specification.....	37
<i>PHY</i>	37
<i>MAC Management Messages</i>	37
<i>Registration</i>	37
<i>MIBs</i>	37
Impact to Implementation.....	37
<i>WCMTS</i>	37
<i>WCM</i>	37
Summary.....	38
About the Author	39

List of Figures

Figure 1: CMTS Port Density and Cost Trends.....	9
Figure 2: Relationship of Wideband and Traditional Downstream Channels.....	11
Figure 3: Relationship of Wideband and Traditional DOCSIS for M-CMTS.....	12
Figure 4: Downstream Wideband MPEG-TS Bonding	13
Figure 5: Wideband Sequence for Bonding	14
Figure 6: Multiplexing of Wideband DOCSIS with Traditional DOCSIS	14
Figure 7: Multiplexing of Multiple Wideband Channels.....	15
Figure 8: Association of Wideband and Traditional DOCSIS Downstreams	16
Figure 9: WCM Upstream with two PSQs.....	19
Figure 10: PSQ State Machine	19
Figure 11: WCM Upstream with two PSQs.	20
Figure 12: WCMTS System Configuration	25
Figure 13: WCM System Configuration	27
Figure 14: Bonded ADSL2+ Loop Lengths	33

List of Tables

Table 1: Data Capacity of Today's HFC Plant	6
Table 2: Data Capacity in use on a 100K HHP HFC Plant Today	8
Table 3: Over-Subscription Example	29
Table 4: Achieving 100x Increase in Downstream Capacity	35

Source Filename: Wideband-DOCSIS-VON-Spring-2005-WP.doc

Document Control: EDCS-410406

Introduction

Overview

Service providers and cable operators around the world are gearing up for the battle of the century. Triple-play offerings, commercial services, and peer-to-peer (P2P) applications are putting demand on networks, bandwidth and data throughput capacities, and resources like never before. In this charged environment, the number of announcements about increased data rates by cable operators is only eclipsed by LEC (Local Exchange Carrier) announcements of service price reductions. However, in spite of this frenzy for dominance, stockholders, analysts, and boards are closely watching the bottom line for signs of weakness or strategic miscalculations. This jockeying for position and dominance has the potential of forcing operators to take the proverbial long drive off a short “broadband” pier—that is, unless a new approach is adopted.

For cable operators, while additional RF bandwidth will be available from digitization of the entire spectrum and judicious management of the spectrum to remove interference, it is anticipated that these methods together are not enough to address the long-term bandwidth consumption growth, projected to be doubling every 12-18 months. It turns out that the only thing constant about bandwidth is the constant demand for more of it and at a lower cost of ownership.

When DOCSIS®-based cable modems (CMs) arrived in the marketplace, they offered 10x to 100x the throughput of dial-up modems with a price that was originally less than twice the price of the dial-up offering. Today, DOCSIS cable modem service has exceeded everyone’s expectations for efficiency, cost, and profitability. However, as with analog modems, the channel capacity of the DOCSIS modem has an upper limit, set by the original specification, which focused primarily on addressing upstream traffic. The current generation of Cable Modem Termination Systems (CMTSs), using the same approach, also has limits to downstream performance because of the tight coupling of upstreams to downstreams. For future growth and network flexibility, this is a significant roadblock to sustaining cable’s bandwidth advantages and meeting the demands of anticipated services and profitability requirements. To date, the cable industry has only scratched the surface of its bandwidth capacity potential.

A new technology promises to deliver the bandwidth and data throughput that will sustain the cable industry’s advantage for the foreseeable future, as well as allow operators to fully leverage their existing DOCSIS and video transport infrastructures, increase network flexibility, and drive down the per-stream costs to new levels. The Wideband protocol for a DOCSIS Network will solve the bandwidth challenge by adding one or more additional downstreams to the standard broadband DOCSIS system. This new set of downstream channels is grouped into one larger channel,

The Wideband Protocol for a DOCSIS® Network

and is known as a Wideband channel. The Wideband protocol can also be delivered on the upstream for return data traffic and signaling.

Wideband DOCSIS enables cable operators to bundle multiple downstreams and upstreams to build cost-effective high-bandwidth pipes that scale to Gigabit Ethernet and above. The de facto standard for video transport is rapidly shifting away from DVB-ASI (Digital Video Broadcast-Asynchronous Interface) to Gigabit Ethernet. In contrast to traditional DVB-ASI infrastructures, Gigabit Ethernet will deliver more data throughput per interface and significantly lower costs because of reduced interface costs and fewer required wavelengths and fibers.

Wideband DOCSIS blows away today's speeds and capacity limitations, offering higher throughput pipes at significantly lower cost points per downstream. It does this by allowing downstreams to be added independent of upstreams, leveraging deployed DOCSIS CMTS platforms, and by taking advantage of the decline in pricing of external Edge QAM devices. The goal of the Wideband protocol for a DOCSIS Network is to provide 10x the throughput at 1/10th the cost.

This paper outlines the market needs driving Wideband protocol adoption, Wideband protocol highlights, and strategy for adoption and implementation. The paper contrasts the Wideband DOCSIS approach and proposed architecture with non-DOCSIS and hybrid alternatives being discussed today. This paper emphasizes the advantages of adopting Wideband DOCSIS and how it fits into CableLabs® efforts and enables cable operators to address new opportunities and compete with xDSL and ETTx/FTTx advancements.

Data Capacity Potential of Today's HFC Plant

HFC Plant Scenario	Data Capacity
HFC Downstream is ~131 channels. At 256-QAM: 131 channels * 38 Mbps	~ 5 Gbps per FN
At one optical transmitter per 500 HHP FN: 5 Gbps / 500 HHP	~ 10 Mbps per HHP
# of HDTV stream per HHP with H.264 AVC 10 Mbps / 7 Mbps	~ 1.4 HD Streams per HHP
Data Capacity of a 100K HHP Plant: 5 Gbps per FN * (100K HHP / 500 HHP per FN) = 5 Gbps per FN * 200 FN =	~ 1 Tbps per small plant

Table 1: Data Capacity of Today's HFC Plant

There is an incredible amount of bandwidth and data capacity in today's HFC plant, most of which are vastly under utilized. The challenge is to mine this bandwidth by

The Wideband Protocol for a DOCSIS® Network

re-organizing the bits on the HFC plant. Table 1 shows an example of the available capacity.

The RF spectrum on a North American 860 MHz plant carries approximately 131 channels at 6 MHz each. If each channel were modulated with a 256-QAM carrier, which had a data capacity of approximately 38 Mbps, the net downstream data capacity of a fiber node is 5 Gbps. With 1024-QAM, this number increases to 6.25 Gbps. However, a newer modulation technique isn't really needed until the first 5 Gbps is used efficiently:

- 131 channels * 38 Mbps per channel for 256-QAM = 5 Gbps
- 131 channels * 48 Mbps per channel for 1024-QAM = 6.25 Gbps

If that 5 Gbps were used to unicast data to every single home on that fiber node of 500 households passed (HHP), there would be an average of 10 Mbps per HHP:

- 5 Gbps per FN / 500 HHP per FN = 10 Mbps per HHP

At 7 Mbps for a H.264 part 10 AVC (Advanced Video Coding) encoded high definition stream, each HHP could receive 1.4 HDTV streams on average:

- 10 Mbps per HHP / 7 Mbps per HD stream = 1.4 HDTV streams per HHP

The throughput per household passed could be improved by a factor of four by assuming a 50% subscription rate with only 50% of the subscribers watching at any one time. They could be improved by another factor of four by subdividing the fiber node to 125 HHP.

Now let's take a small city of 100,000 HHP. At 500 HHP per fiber node, this would be 200 fiber nodes. At 5 Gbps per fiber node, that small city has 1 Terabit per second (Tbps) of downstream throughput!

- 100,000 HHP per city / 500 HHP per FN = 200 FN per city
- 200 FN per city * 5 Gbps per FN = 1 Tbps per city

The irony is that that data capacity of that size of city today is probably serviced by something between several Fast Ethernets or a Gigabit Ethernet – something that is 1000 times less than the actual capability of the HFC plant.

The bandwidth and data capacity of the HFC plant largely go unrealized. Seventy-nine of the 131 channels are still analog which is 1/10 as efficient and has less quality than an MPEG-2 digital video signal encoded at 3.75 Mbps. In addition, all the downstreams are combined and placed in parallel as part of a broadcast architecture (except for the DOCSIS and VOD feeds). A capacity analysis of a 100K HHP HFC Plant is shown in Table 2.

Service	# Channels		Digital BW	# SG	Capacity	
Analog	79	60%	3.75 Mbps	1 SG	0.3 Gbps	3%
Digital	43	33%	38 Mbps	1 SG	1.7 Gbps	15%
VOD	8	6%	38 Mbps	50 SG	15 Gbps	72%
DOCSIS	1	1%	38 Mbps	50 SG	2 Gbps	10%
BW Today	131				19 Gbps	
Capacity	131		38 Mbps	200 SG	1000 Gbps	
Efficiency					1.9%	

Table 2: Data Capacity in use on a 100K HHP HFC Plant Today

If there were one unicast DOCSIS feed and eight VOD unicast feeds per four fiber nodes, then the equivalent capacity that today's plant is being used at is 19 Gbps:

- $(79 \text{ ch} * 3.75 \text{ Mbps}) + (43 \text{ ch} * 38 \text{ Mbps}) + (9 \text{ ch} * 200/4 \text{ fiber nodes} * 38 \text{ Mbps})$
 $= 300 \text{ Mbps} + 1.6 \text{ Gbps} + 17 \text{ Gbps}$
 $= 19 \text{ Gbps}$

There are some interesting conclusions to draw from this analysis. First, today's HFC plant, in this example of 100K HHP, is only being used at 2% or less of its capacity:

- $19 \text{ Gbps per } 200 \text{ FN} / 1 \text{ Tbps per } 200 \text{ FN} = 1.9\%$

This number drops well below 1% when it is realized that the downstream combining is usually common for several hundred thousand to several million homes passed.

Second, 60% of today's HFC plant is used to create 1.6% of its data capacity while 7% of today's HFC plant is used to create 90% of its data capacity:

- $79 \text{ ch} / 131 \text{ ch} = 60\%$
 $43 \text{ ch} / 131 \text{ ch} = 33\%$
 $9 \text{ ch} / 131 \text{ ch} = 7\%$
- $300 \text{ Mbps} / 19 \text{ Gbps} = 1.6 \%$
 $1.6 \text{ Gbps} / 19 \text{ Gbps} = 8.4\%$
 $17 \text{ Gbps} / 19 \text{ Gbps} = 90\%$

The goal is clear. The need is not to increase the bandwidth and data capacity of the plant, either through extending the RF spectrum, increasing modulation rates, or driving fiber deeper. The goal is to take the available bandwidth and data capacity and use them more efficiently by changing the connectivity of the backbone to the plant. In doing so, we will need a new generation of CMTSs that can provide 10

times, 100 times, or even 1000 times the data capacity at a price point that is significantly lower than today's per port CMTS costs.

Price and Density Trend Lines

CMTSs have received a bad rap as being too expensive. The truth is two fold. First, the revenue for data services from a CMTS is usually enough to pay for a CMTS in 1-2 weeks. However, a CMTS cannot compete today with Edge QAM devices on cost per downstream. Edge QAM devices in the market today can offer downstream connectivity at 5% to 10% the cost of an equivalent downstream port on a CMTS.

There is a good reason for this. Edge QAM devices only contain downstream ports and were optimized for downstream connectivity. CMTSs contain downstream and upstream ports and were optimized for upstream connectivity. For every downstream port bought on a CMTS, you get 4-8 upstream ports. If a port-to-port comparison were done (upstream and downstream ports), the difference would be closer to 25% to 50%.

The answer is obvious. The CMTS has to be designed so that downstreams may be added, and thus purchased, independent of upstreams. Figure 1 shows data for cost and density trends. This data is normalized to the first DOCSIS systems shipped in 1997.




Figure 1: CMTS Port Density and Cost Trends

The reality is that costs have come down already through integration and increased density. The next generations of CMTSs that employ Wideband and eventually a Modular CMTS™ architecture with separate MAC and PHY will continue the curve downwards.

The Wideband Protocol

The Wideband protocol for a DOCSIS network, originated by the author in 2001 and in development by Cisco Systems, solves the bandwidth and data capacity challenges by adding one or more additional channels to the standard broadband DOCSIS configuration. This new set of channels grouped into one larger “Wideband” channel may be applied in either the downstream direction or the downstream and upstream direction, although the details of the implementation are different.

Wideband downstream channels can utilize external QAM modulators and up-converters via Gigabit Ethernet spigots or internal DOCSIS QAM modulators on existing RF line cards. The cost challenge is addressed by allowing downstreams to be added independent of upstreams, and by taking advantage of the decline in pricing of external Edge QAM devices. The goal of the Wideband protocol for a DOCSIS Network is to provide 10x the throughput at 1/10th the cost of today’s CMTSs.

A CMTS that implements the Wideband protocol is referred to in this paper as a WCMTS, and a CM which implements the Wideband protocol is referred to as a WCM. The term “a traditional DOCSIS channel” will be used to refer to the single QAM channel currently defined in DOCSIS 1.1 or 2.0.

The Wideband protocol is being targeted for the next version of DOCSIS – DOCSIS 3.0 – and is backwards compatible with all earlier versions of DOCSIS.

The Concept

DOCSIS 1.1 and 2.0 Systems with a Wideband Overlay

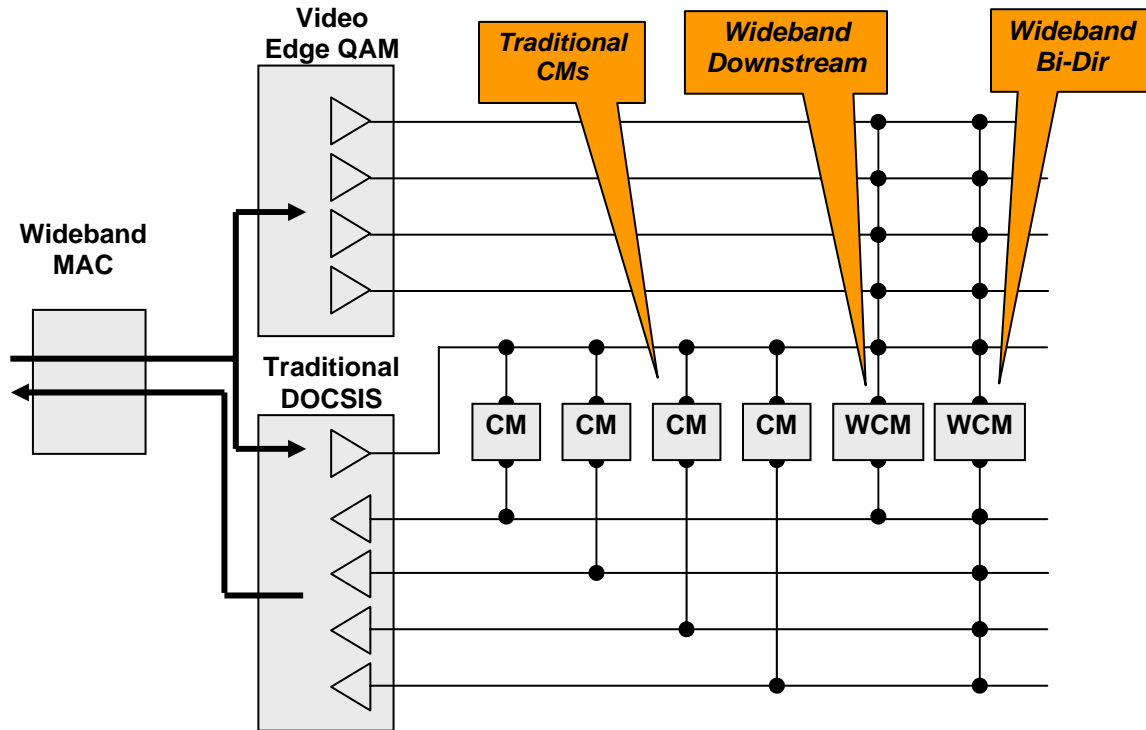


Figure 2: Relationship of Wideband and Traditional Downstream Channels

The relationship of Wideband and traditional QAM modulators is shown in Figure 2. Traditional DOCSIS CMs share a common downstream and are spread across several upstreams. The WCM uses the same traditional infrastructure and adds more downstream or upstream capacity.

The multiple QAM carriers within a Wideband channel are coupled together to form what looks like a multi-carrier PHY. The coupling does not actually occur in the PHY; it occurs within the transmission convergence layer between the MAC and PHY. This is very significant because it implies that large pipes – upwards of a Gigabit – can be constructed out of today’s inexpensive QAM technology. The Wideband protocol will support any number of QAM carriers, although the implementations of the WCMTS and WCM will set operational limits.

This scenario is targeted at the maximum re-use of existing equipment, including current CMTSs and current Edge QAMs.

DOCSIS 3.0 with a Modular CMTS System

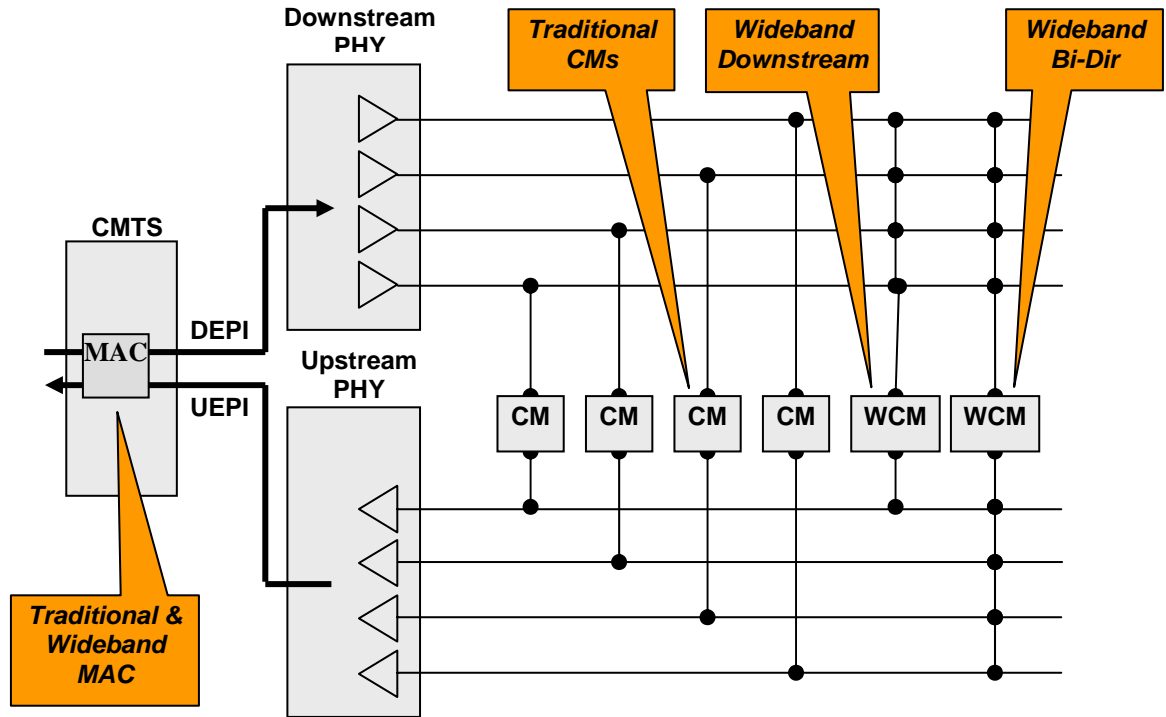


Figure 3: Relationship of Wideband and Traditional DOCSIS for M-CMTS

The Modular CMTS (M-CMTS™) architecture work is underway at CableLabs as part of the DOCSIS 3.0 effort. It is proposed in this paper that for the M-CMTS architecture, the traditional DOCSIS MAC and the Wideband MAC be placed in the CMTS and then interface to the Edge QAM with the Downstream External PHY Interface (DEPI) protocol. This Edge QAM would be synchronous and would support both traditional DOCSIS CMs as well as Wideband CMs. This concept is shown in Figure 3.

With this separate MAC-PHY approach, the MAC can be replaced and scaled to larger and larger channel capacities without disturbing the Edge QAM. Further, the Wideband channel may be bonded across Edge QAMs in separate chassis so that if an Edge QAM fails, the Wideband channel can reconfigure around the remaining Edge QAMs.

Wideband in the Downstream

The Wideband protocol has a separate bearer channel and signaling channel. The bearer channel is the collection of QAM carriers, which are bonded together as a Wideband channel. The signaling channel is a traditional, non-bonded DOCSIS channel. The bearer channels and signaling channels may actually share the same QAM carriers, but are kept logically separate through the use of different MPEG-TS (Moving Pictures Experts Group-Transport Stream) PIDs (Program Identifiers).

The Downstream Wideband Channel




Figure 4: Downstream Wideband MPEG-TS Bonding

The downstream Wideband channel is created by taking DOCSIS frames, putting them into MPEG-TS packets, and placing those MPEG-TS packets onto QAM carriers. However, instead of placing those MPEG-TS packets “horizontally” in time along a single QAM carrier as is done in traditional DOCSIS, the Wideband protocol places those MPEG-TS packets “vertically” across the QAM carriers assigned to a Wideband channel. This is conceptually shown in Figure 4. A DOCSIS frame is literally tipped on its side and striped out across a group of QAM channels.

The Wideband Protocol for a DOCSIS® Network




Figure 5: Wideband Sequence for Bonding

Figure 5 shows the bonding of QAM carriers containing MPEG-TS streams in more detail. The technique that was chosen embeds a sequence number in each MPEG-TS packet. This allows the transmit Wideband framer the maximum flexibility of laying down MPEG-TS packets as it needs to. It also keeps the reconstruction information for the receive Wideband framer in band for robustness.

The segmentation of the DOCSIS frame into smaller units could have been done at the bit, byte, MPEG-TS packet, or even at the packet level. The MPEG-TS level is a convenient choice because the bonding of the channels occurs at the transmission convergence layer, which is above the PHY layer and below the MAC layer. That allows the PHY and MAC (in particular the framing portion of the MAC) to remain unchanged. The Wideband protocol is transparent to the traditional DOCSIS protocol. This is very powerful as it has the potential to maximize the re-use of the existing DOCSIS environment.

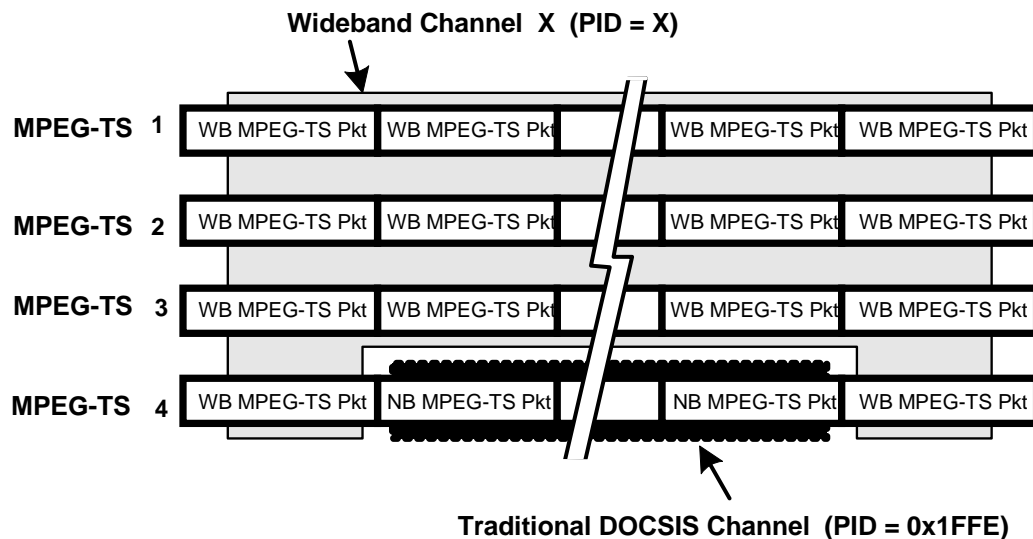


Figure 6: Multiplexing of Wideband DOCSIS with Traditional DOCSIS

The MPEG-TS packets used for the Wideband channel are assigned a unique PID. This PID is different from the DOCSIS reserved PID of 0x1FFE. This allows the

Wideband channel to be multiplexed with traditional DOCSIS channels on any or all QAM carriers within the Wideband channel. This is shown in Figure 6. Since the Wideband PID is operator assigned, the Wideband channel may also be multiplexed with other Wideband channels. This is shown in Figure 7. Since the Wideband channel is PID based, it may also be multiplexed with other MPEG-TS services such as video.




Figure 7: Multiplexing of Multiple Wideband Channels

QAM Adjacency

The Wideband protocol does not require the QAM carriers to be adjacent. This allows QAM carriers from different parts of the spectrum to be bonded together without having to re-juggle the channel line-up. There are some operational considerations.

The first constraint is the Wideband tuner in the WCM, which is based upon a block downconverter and followed by an A/D converter. The bandwidth of today's A/D converters allows about 100 MHz of spectrum to be captured. That capture window of 100 MHz is the equivalent of sixteen 6 MHz channels, or twelve 8 MHz channels, and can be located anywhere in the DOCSIS downstream spectrum of 88 to 860 MHz. To keep the WCM low cost, one such block downconverter and A/D assembly should be sufficient.

The second constraint is introduced by frequency stacking and direct RF synthesis at the CMTS and Edge QAM device. In frequency stacking, multiple QAM carriers are created in the digital domain. These are then sent through shared D/A converters, shared block upconverters, and through RF power amplifiers. The result is QAM carriers that are adjacent. However, if some of those QAM carriers are used for non-DOCSIS applications such as VOD, the remaining QAM carriers used for DOCSIS do not need to be adjacent.

Wideband Signaling Channel

The CMTS already sends signaling messages such as MAPs on the traditional DOCSIS downstream. Rather than duplicate these messages on the Wideband downstream, the WCM uses the traditional DOCSIS downstream as its signaling channel. This permits the legacy CMs and the new WCMs to exist on the same control plane. Thus, a WCM must connect to both the traditional DOCSIS downstream and upstream channel as well as the Wideband DOCSIS downstream and upstream channel.

There does not have to be a one-to-one correspondence between the Wideband channels and the traditional DOCSIS channels. The Wideband channel may be shared with video QAM carriers that have a different connectivity to the HFC plant than the DOCSIS QAM carriers. Also, because of the higher throughput and initially lower subscription rate for Wideband, it may make sense to combine more RF downstreams together into one Wideband channel. The resulting combinations are shown in Figure 8.

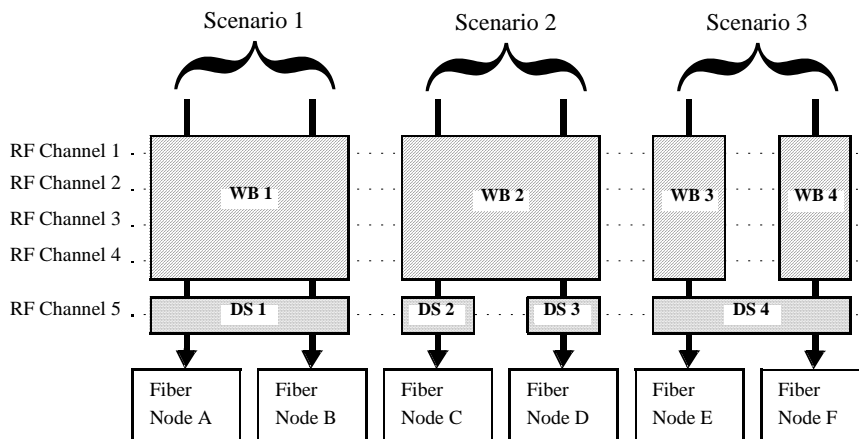


Figure 8: Association of Wideband and Traditional DOCSIS Downstreams

This fact, combined with the scalability of the size of the Wideband channel, provides considerable flexibility and scalability of the protocol.

Timing

The Wideband channel is an asynchronous channel in that it does not contain any timing information for the WCM. The traditional single DOCSIS downstream channel, by contrast, can be considered a synchronous channel in that it does carry timing information for the CM in the form of DOCSIS SYNC messages and, for DOCSIS 2.0, the frequency of the underlying QAM symbol rate.

This means that any additional delay over the Wideband channel incurred from the disassembly and reassembly of the DOCSIS frames will not impact the signaling messages, such as MAPs, since they don't traverse the Wideband channel.

Skew

As discussed later in this document, the Wideband protocol has built-in high availability features that allow it to repair itself. This is made possible by locating the Wideband MAC in a separate location from the Edge QAM device. By having the Wideband MAC separate, the Wideband channel can be striped across multiple Edge QAM chassis. If one Edge QAM chassis fails, then the Wideband protocol can shrink the size of the Wideband channel to path around the failed Edge QAM device.

This means that the skew between QAM carriers either within an ASIC (Application Specific Integrated Circuit), within a line card, between line cards within a chassis, and between line cards in different chassis has to be accommodated. There is also network or backplane skew and jitter that might exist between the Wideband MAC and the Edge QAM devices. Another source of skew exists if the downstream QAM carriers are run at different modulation rates such as some channels at 64-QAM and some channels at 256-QAM.

The final amount of skew has to be quantified, minimized, and accommodated with the depth of the reconstruction buffer in the WCM. Buffer depths that have a nominal reconstruction delay of 0.5 to 2 ms with peak delays of up to 5 to 10 ms might not be unreasonable.

Wideband in the Upstream

The upstream Wideband protocol shares similarities with the downstream Wideband protocol, but differs in its details. Both directions obtain higher throughput by bonding together multiple QAM channels. Both directions have a default signaling channel that is a traditional single DOCSIS QAM channel and that is independent of the Wideband channel. In addition, both directions overlay the Wideband channel on top of traditional DOCSIS channels.

What differs is that the upstream does not have an MPEG-TS framing structure. Also, there has to be a way to request upstream capacity, and do so across multiple upstream carriers.

The diagrams for the upstream follow the convention of having the subscriber premises on the right hand side and the cable headend on the left hand side. This causes upstream data flow to go from right to left.

The Packet Streaming Protocol

IP packets are first received by the WCM on its CMCI (CM to CPE Interface) port. These IP packets are placed into a Packet Streaming Queue (PSQ). The Packet Streaming Queue is part of the Packet Streaming Protocol (PSP) and is a new concept being introduced in this paper. The Packet Streaming Protocol allows packets to be appended together in a queue on an ongoing basis. The queue may then be broken up into arbitrary chunks of data and transmitted.

This is different from the concatenation and fragmentation protocols that are found in DOCSIS. The primary differences are:

- Concatenation is done for a fixed number of DOCSIS frames at a time. The Packet Streaming Protocol is done on a continuous number of IP packets. Hence the use of the name “streaming”.
- The protocol headers are completely different.
- Fragmentation and concatenation are implemented at Layer 2. The Packet Streaming Protocol is implemented at Layer 3.
- Fragmentation and concatenation are per Service Flow and within an upstream QAM carrier. The Packet Streaming Protocol may be per Service Flow, across multiple Service Flows, and/or across multiple QAM carriers.
- DOCSIS Requests and Grants (MAPs) are managed independently of the Packet Streaming Queues.

These differences will become more obvious in the next section as the operation of the upstream Wideband channel is described.

The Packet Streaming Queue




Figure 9: WCM Upstream with two PSQs.

Figure 9 shows an example of the upstream portion of a WCM. Incoming packets are sorted into input queues called Packet Streaming Queues (PSQ). There would be one PSQ per Quality of Service (QoS) queue. This allows the WCMTS to manage QoS at the WCM (as is done in a traditional DOCSIS system with multiple SIDs) and

prevents head-of-line blocking, a scenario where a higher priority packet may get stuck behind a lower priority packet. The output of the PSQ is then sent to one Service Flow (SF) chosen from a group of Service Flows, which in turn are located on one QAM carrier chosen from within a group of upstream QAM carriers. There could be any number of PSQs, each with any number of QAM/SF combinations.

As the PSQs begins to collect bytes, the PSQ State Machine shown in Figure 10 will issue PSQ-REQs to the DOCSIS REQ-GNT State Machines, which in turn launch requests (REQs) to the CMTS for a given number of bytes. One outstanding REQ is permitted per Service Flow per upstream QAM carrier. The use of multiple upstream QAM/SF combinations allows multiple PSQ-REQ to be outstanding at any one time.

The PSQ Traffic Manager will choose which upstream QAM channel and on which DOCSIS Service Flow to launch the REQ. This may be based upon the perceived least busy upstream channel, upon some weighting criteria supplied by the WCMTS, or even a simple round-robin approach.

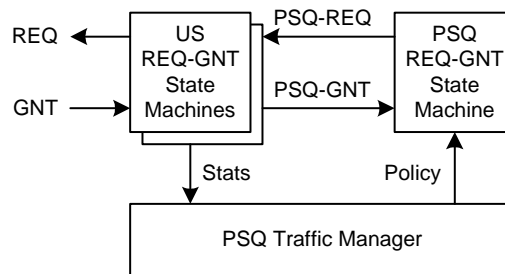


Figure 10: PSQ State Machine

The size of PSQ-REQ is a configuration parameter from the WCMTS to the WCM. The PSQ-REQ size does not

need to line up with a packet boundary. It may be a fraction of a packet length, multiple packet lengths, or both. It may be a pre-determined value, a multiple of a predetermined value, or an arbitrary value. One choice is to have a value of PSQ-REQ chosen so that the Wideband payload would fill the upstream FEC Blocks efficiently.

Each PSQ, as it receives more packets, continues to generate PSQ-REQs without waiting for PSQ-GNTs, up to the maximum number of outstanding PSQ-REQ allowed. Eventually the WCMTS sends back GNTs in the form of MAPs. (The GNT is really a transmit opportunity in a MAP message.) Now here is the interesting part.




Figure 11: WCM Upstream with two PSQs.

Notice that the GNTs are out of order. This does not matter because the PSQ State Machine dynamically creates a PSQ-GNT to match the first GNT received and sends that PSQ-GNT to the PSQ. The PSQ then de-queues the number of bytes indicated by the PSQ-GNT – not the number of bytes indicated by PSQ-REQ – and delivers it to the output queue corresponding to the QAM/SF pair prior to the time that the data is needed on the wire. This provides the effect of having multiple outstanding REQs, each of which may be a different size.

The DOCSIS GNTs may be part of a DOCSIS layer 2 fragmentation or concatenation. In either event, the PSQ State Machine translates the payload of the GNT to an appropriate PSQ-GNT.

The net effect is that a stream of IP packets is received by the WCM from the home network. The WCM launches a series of REQs and gets a series of GNTs. The GNTs are used as received to service the output queuing structure based upon QoS. In the end, the number of bytes requested equals the number of bytes granted, the law of conservation of bytes is upheld, and all packets make it through the WCM with their QoS intact.

Comparison with the Traditional DOCSIS Upstream

Both the REQs and GNTs are identical to the REQ and GNTs in the DOCSIS 2.0 protocol. REQs may be contention based or piggybacked on a DOCSIS frame. GNTs may be solicited or unsolicited grants. The Wideband protocol can contain policy information that steers certain grants towards certain PSQs. For example, there could be a PSQ for VoIP that receives all the unsolicited grants, while the solicited grants are distributed to the other PSQs. Alternatively, VoIP GNTs could be sent to

The Wideband Protocol for a DOCSIS® Network

dedicated Service Flows on traditional DOCSIS upstreams, and data GNTs could be sent to the Wideband channel.

Since there is only one outstanding REQ per Service Flow per QAM carrier, all the current CM and CMTS upstream scheduling mechanisms and local REQ-GNT state machines can be preserved. This extremely important feature allows interoperability of WCM and CM within not only the same upstream, but overlaid on top of the existing CMTS and CM designs.

The Packet Steaming Protocol used by Wideband will work in conjunction with the DOCSIS layer 2 services of fragmentation and concatenation. These two services are algorithms that translate DOCSIS REQs into DOCSIS GNTs. PSP will always deliver the amount of bytes asked for by the DOCSIS GNT.

Since BPI+ keys are per WCM, it would make sense to use the same BPI+ key on all upstreams within a Wideband flow from a particular WCM. PHS could be applied to all upstream Wideband channels by suppressing the Ethernet/IP/UDP header since all PSP packets from a WCM will be sent to the same source.

The Packet Streaming Protocol is not Wideband specific. It could be applied to a single DOCSIS upstream channel by utilizing multiple Service Flows. With only two to four Service Flows and DOCSIS concatenation, this would allow a CMTS to provide to a CM up to 100% of the throughput of an upstream. Although this might not be practical for an extended period of time on a shared network, it may be useful for providing lower latencies to bursts of upstream traffic.

Configuration

The upstream QAM/SF combinations would each be assigned a DOCSIS Service Flow. This would be the child Service Flow. The aggregate upstream Wideband channel would be assigned another Service Flow. This would be the parent Service Flow. The parent Service Flow would have knowledge of the child Service Flows. This would result in a two-tier hierarchical Service Flow. An example of the usage would be to assign the parent flow 40 Mbps and each of four child Service Flows would be assigned 20 Mbps. This would permit the Wideband channel to be dynamically allocated by the WCM to different upstream QAM/SFs as resources are available.

In a traditional DOCSIS upstream, Service Flows and SIDs are used to manage different QoS scheduling algorithms for different sessions. Since each PBQ uses multiple Service Flows, PBQs will tend to be assigned to an entire class of traffic. For example, all data may go on one PBQ that is given best effort scheduling by the CMTS, while all VoIP may go on another PBQ that is given unsolicited grant service (UGS) scheduling.

For the case of UGS, the WCM would not send a PSQ-REQ or a REQ. Instead, the WCMTS would distribute UGS GNTs across the upstream QAM/SF pairs as it felt appropriate.

Extensions to the Basic Idea

DOCSIS upstreams are considered smart pipes where the CMTS manages the QoS of the upstream flow by prioritizing bandwidth GNTs to different WCM Service Flows/SIDs – as opposed to a dumb pipe model where the CMTS would provide a homogenous QoS pipe, and the CM would schedule traffic locally to that pipe based upon QoS. So on a smart pipe, the CMTS would manage data and voice separately, whereas on a dumb pipe, the CM would manage data and voice separately.

As an extension to the basic upstream wideband idea and to facilitate a smart pipe model, the received GNTs do not even have to be applied to the PSQ that originated the REQ. The received GNTs may be applied to the PSQs depending upon QoS policy, GNT size, or other criteria, and not upon which PSQ generated the REQ. The PSQ State Machine could even split up the GNT between two PSQ-GNTs if necessary. This would allow a common pool of SIDs to be shared among multiple PSQs. However, if there is sufficient SID space, this extra layer of complexity may not be necessary.

If the WCM could be configured with queuing policies to manage upstream QoS, then the WCM could have multiple QoS queues that would then feed into one PBQ, which then drove multiple upstream QAM/SF combinations.

The concept of multiple outstanding REQs could be taken one step further and applied to the actual REQ-GNT state machines which exist per SID. In this scenario, multiple REQs would be issued per SID. If a REQ were lost due to contention, the use of the Packet Stream Protocol would allow the current REQ to be used and the dropped REQ could be resent.

Wideband Signaling

Wideband Channel Descriptor

The Wideband Channel Descriptor (WCD) is a message sent on the traditional DOCSIS downstream that contains information about the associated Wideband channel. There may be more than one WCD as there may be more than one Wideband channel associated with a particular traditional DOCSIS channel.

Each WCD contains a list of frequencies of the QAM carriers for its Wideband channel. It also contains the PID value that the WCM should use for receiving the

Wideband channel. The WCD also contains operational status information for each QAM carrier.

Registration

The WCM identifies itself as being Wideband capable during the configuration process. This allows the DOCSIS TFTP provisioning server to enable or disable Wideband mode, and to choose the appropriate configuration parameters.

Two design options exist for managing registration: single or double registration. In single registration, the same TFTP configuration file contains all the information for Wideband and traditional DOCSIS channels. This may be a convenient way of keeping all configuration information in one place. In double registration, the traditional CM receives one TFTP configuration file, while the WCM entity receives a second TFTP configuration file. This technique treats the Wideband network as an overlay network over the traditional DOCSIS network. It also minimizes the changes to a standard TFTP configuration file.

Addressing

All addressing for packets is done with Ethernet and IP addresses. PIDs are used within the MPEG-TS structure to denote a particular Wideband channel, but are not used to address individual WCMs.

The WCM uses DHCP to obtain an IP address. Two design options exist for obtaining IP addresses: single or dual IP addresses. For single IP address, that infers that the Wideband DOCSIS MAC and the traditional DOCSIS MAC are on the same subnet and within the same CMTS, perhaps even the same line card. For dual IP address, it presumes that the Wideband DOCSIS MAC may have a different connectivity to the HFC plant than the traditional DOCSIS MAC (as shown in “Figure 6: Multiplexing of Wideband DOCSIS with Traditional DOCSIS” on page 14), and that the traditional DOCSIS MACs may be assigned to different IP subnets.

Security

The Wideband channel uses the same DOCSIS frame format and the same Baseline Privacy (BPI+) for its link level encryption as the traditional DOCSIS channel. It just has more capacity. The Wideband channel runs BPI+ for its link level encryption. Two design options exist: same or different keys.

With the first option, all SFs on the Wideband DOCSIS channel and the traditional DOCSIS channel would use the same BPI+ keys. This makes sense from a WCM perspective since BPI+ keys are intended to be per CM and the traditional DOCSIS channel and the Wideband DOCSIS channel are terminating on the same WCM.

The Wideband Protocol for a DOCSIS® Network

With the second option, the Service Flows on the Wideband DOCSIS channel would use a separate set of keys from the traditional DOCSIS channel. This may make sense from the implementation viewpoint of the WCMTS if the Wideband MAC is located separately from the traditional DOCSIS MAC.

MIBs

The Wideband MIB for the WCM would be an extension to the existing CM MIB or the same MIB instantiated twice: once for the CM entity and once for the WCM entity.

Wideband Signaling Summary

These options will get debated as part of the DOCSIS 3.0 working group. The solution in each case could be the first option, second option, or both options.

Wideband Deployment

WCMTS System Configuration

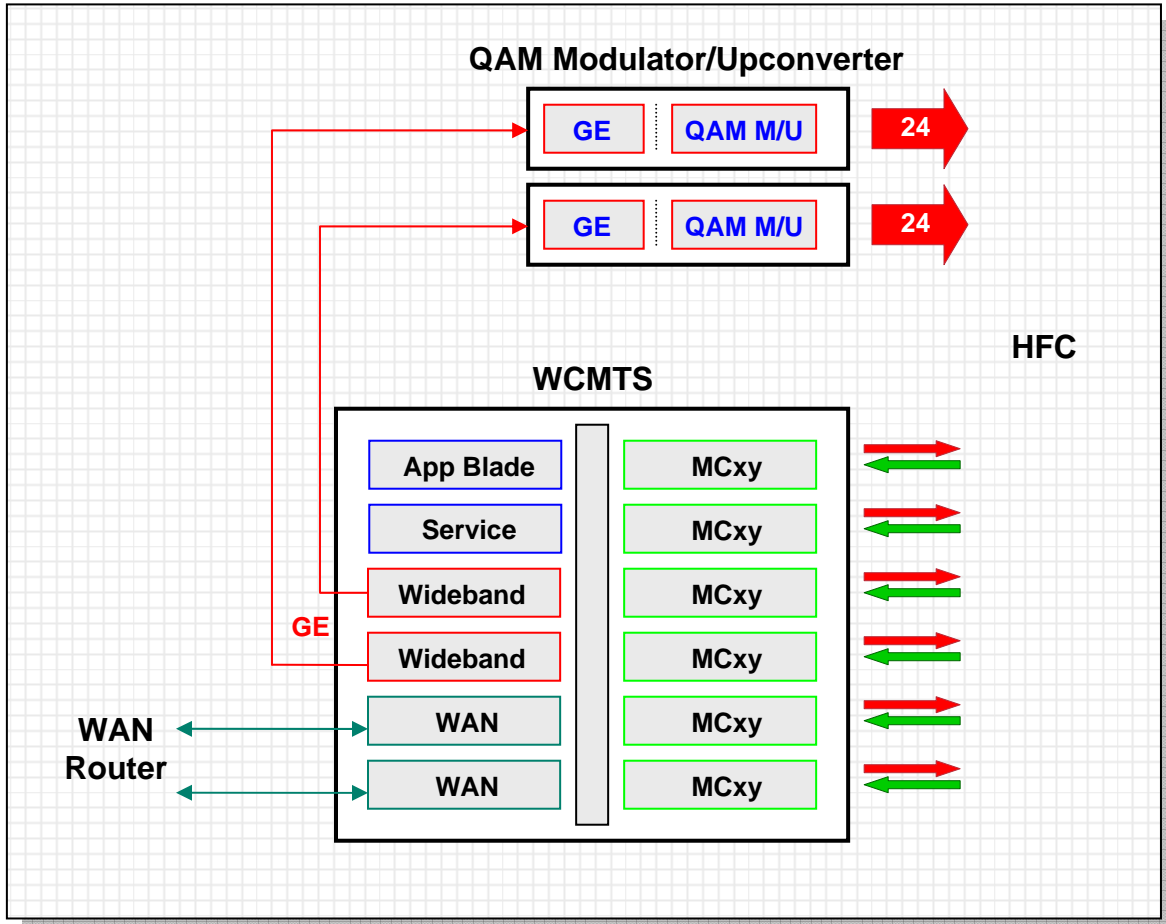


Figure 12: WCMTS System Configuration

The packaging of the WCMTS has many options. One proposed plan is shown in Figure 12.

WCMTS

The WCMTS would completely manage and own the Wideband protocol. The MAC and framer portion of the Wideband protocol would be implemented in a Wideband line card in a conventional CMTS. Each line card would initially manage some number of downstream and upstream QAM carriers, say, 24 or 48 carriers. For the downstream, the Wideband line card would strip downstream packets across local

MPEG-TS queues. The output of these MPEG-TS queues would be encapsulated in a UDP datagram and sent to the Edge QAM device. For the upstream, the Wideband line card would Receive UDP datagrams from either internal or external burst demodulators, remove the UDP encapsulation, reassemble the payload, and send the resulting packets into the backbone network.

Edge QAM Device

An external Edge QAM device would be used to couple the Wideband MPEG-TS packet onto the HFC plant. However, the Edge-QAM device would not be involved in the active management of data capacity of the Wideband channels. The Edge QAM device would not even be aware of the IP addressing used by the WCMs. It would simply terminate the UDP Tunnel from the WCMTS, take the MPEG-TS payload, and put it onto the HFC plant.

Because the Wideband protocol is asynchronous, and the Wideband framer is added to an existing CMTS to make it a WCMTS, standard Edge QAM devices that exist today could be used in conjunction with the WCMTS to generate a Wideband channel. This is a huge advantage for early deployment.

There is no technical reason that the QAM modulators have to be external to the WCMTS. They could be internal, external, or both. External Edge QAM devices are used because they have a declining price point, and because they provide a system configuration for the cable operator, which would allow one bank of QAM modulators to be used by two services – native MPEG-TS video and DOCSIS-based IP services.

High Availability

By placing the Wideband framer in the WCMTS, downstream Wideband channels may be striped across multiple Edge QAM devices. If an Edge QAM device fails, the WCMTS can patch around the failed device by dynamically removing those QAM carriers from the Wideband channel line-up. When the Edge QAM device comes back on line, the WCMTS can dynamically add those QAM carriers back into the Wideband channel line-up.

This is a significant feature of the Wideband protocol. It allows the Wideband protocol to survive the failure of an Edge QAM device without redundancy and without an RF Switch. The trade-off is less throughput during the failure condition. This is usually an acceptable trade-off. If not, the Edge QAM devices can be made more reliable with an RF switch and redundancy, and the Wideband channel will maintain the same throughput during a failure. Note that if the Wideband protocol were implemented in the Edge QAM device and the Edge QAM device failed, the entire Wideband channel would be lost.

The same is true with the upstream Wideband protocol. If the upstream Wideband channel is striped across multiple DOCSIS line cards and one line card fails, the Wideband channel can be dynamically reconfigured around the failed line card. This is only possible if the Wideband framer is separate from the DOCSIS receivers.

WCM Design

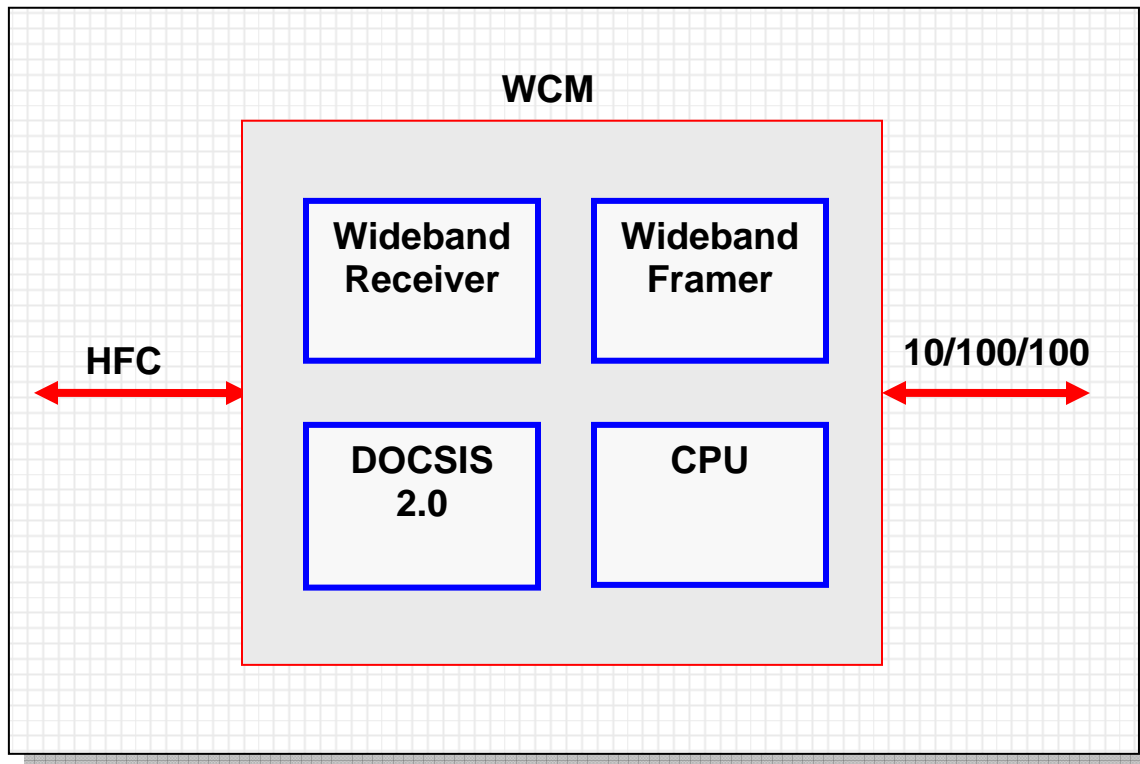


Figure 13: WCM System Configuration

The WCM design is backward compatible with DOCSIS 1.1/2.0 and contains a complete DOCSIS 2.0 core. If the Wideband channel is not available, the WCM can operate strictly as a CM.

The WCM adds a Wideband receiver that receives the multiple QAM carriers. It also adds a Wideband framer that decodes the signal from the Wideband receiver and extracts the packets for the home network. This is illustrated Figure 13.

As time progresses, the Wideband framer may integrate into either the DOCSIS chip or the Wideband receiver. This would bring the cost of the WCM in close alignment with current CM pricing.

The Big Pipe Strategy

The larger the data channel and the smaller the data stream per subscriber, the more subscribers a data channel will support. There are two reasons for this. First, the simple division of the two numbers is higher. Second, and more importantly, is that the natural statistical multiplexing of the subscriber's traffic is more efficient because the more users that exist on a downstream, the more random and distributed the overall traffic pattern appears.

For example, typical CMTSs today might use 38 Mbps on a downstream channel (256-QAM, 6 MHz channel bandwidth) and support 1,000 subscribers, each, which may have an offered data throughput of 2 Mbps. This would suggest the following metrics:

- Offered_Throughput_to_Channel_Throughput_Ratio
= 2 Mbps / 38 Mbps \approx 1/20
= 5%
- Over_Subscription_Ratio
= Total_Offered_Throughput / Channel_Throughput
= 1000 subscribers * 2 Mbps per sub / 38 Mbps per channel
 \approx 50

So today's deployments with today's data profiles (bursty, latency-tolerant data), with an offered data throughput which is 5% of the channel's data throughput, will support an over-subscription ratio of 50. If the channel throughput remains the same but the offered data throughput per subscriber is increased, then the natural statistical multiplexing becomes less efficient and the over-subscription ratio would have to be decreased.

As an extreme example, if the offered throughput were the same as the channel throughput of 38 Mbps, then the offered throughput could not be delivered if two subscribers transferred data at the same time. The probability that two subscribers will simultaneously transmit is intuitively much higher than 20 subscribers (from the previous example) simultaneously transmitting. Also intuitively, this probability dramatically increases further as burst sizes increase as a result of the content moving from bursty data to more constant multi-media traffic.

Thus, it becomes extremely inefficient to use a single channel for large throughput streams. This problem must be managed by reducing the over-subscription ratio, and thus the number of subscribers for a given chunk of throughput, and/or increasing the channel throughput.

Table 3: Over-Subscription Example" shows an intuitive approach to keeping the statistical multiplexing constant by holding the ratio of the subscriber offered

throughput to the channel throughput constant, and calculating the required width of the Wideband channel in terms of QAM carriers.

Parameter	Today	0-2 Years	3-5 Years
<i>Offered Throughput</i>	2 Mbps	10 Mbps	30 Mbps
<i>Channel Throughput</i>	38 Mbps	200 Mbps	600 Mbps
<i>Ratio</i>	1/20	1/20	1/20
<i>QAM Carrier(s)</i>	1	5	16

Table 3: Over-Subscription Example

To make the scenario even worse, as the content on the network moves from a bursty data profile to more of a variable bit rate video profile, there will again be less efficiency in the averaging of traffic. This will further decrease the over-subscription ratio, making it even more important to keep the ratio of channel throughput to peak subscriber throughput high.

What About DOCSIS Load Balancing?

DOCSIS contains a comprehensive suite of load balancing features through the DCC command (which was originally invented by the author¹). DCC was intended to allow multiple traditional DOCSIS channels, either downstream or upstream, to appear to look like one large pipe by dynamically moving traffic between QAM carriers.

To load balance well, load balancing software needs to be able to predict what is going to happen before it happens. For VoIP, which is what DCC was initially designed to manage, all calls are booked so the load balancing software has complete visibility into what is about to happen and can make intelligent decisions about what is going to happen next. There are still issues with moving new calls while other calls are in progress, but at least for VoIP, that is manageable.

It is a different story for data traffic. Data traffic is more random. By the very nature of randomness, it is not easily predicted. For a small number of channels with a high number of modems with low traffic loads that are data only, general observations may be made regarding bulk traffic patterns that allow some reasonable load balancing. However, for larger channel groups, for non-data traffic, or for large streams that are greater than 10% of the channel capacity, the imperfections with load balancing and the resulting inefficiencies are noticeable.

In addition, load balancing presumes that IP streams may be interrupted. VoIP is actually the most tolerant of this because VoIP sessions run concealment, which

¹DOCSIS ECO RFI-O-00021

covers up dropped voice packets. Data sessions recover, but do so with end-to-end retransmits. It has always worked to drop a packet or two, but this is being tolerated less and less. Finally, there are video flows. When a DCC is done that move a video flow, there will be packet loss. That packet loss will usually show up as noticeable artifacts on a video screen. The amount of load balancing needed usually exceeds the allowable number of screen artifacts allowed over a period of time.

In summary, DOCSIS load balancing is a good solution for two to four channels as a practical maximum, and for traffic that is VoIP or non-critical data. Large numbers of QAM carriers and video traffic would be better served by a single large pipe that did not interrupt exiting IP flows when new IP flows are added.

Wideband Migration Strategies

With the introduction of DOCSIS 2.0, the upstream data throughput was increased from 10 Mbps to 30 Mbps, a 3x increase. Meanwhile, the downstream maximum raw data rate has stayed at 42 Mbps (per 6 MHz channel). In fact, with today's typical 1x4 DOCSIS configurations of downstreams to upstreams, there is now more upstream throughput than downstream throughput! A 1x4 DOCSIS 2.0 domain has 40 Mbps down x 120 Mbps up. Thus, with today's CMTS configurations, it is not really possible to use the data throughput provided by DOCSIS 2.0 (ATDMA or SCDMA).

The irony of this is that the next wave of IP applications, such as video over IP, are primarily downstream based. Thus, the first market need for the Wideband protocol will be to address the upcoming shortage of downstream data capacity.

Today's ratio of offered downstream throughput to offered upstream throughput is typically:

- Offered_Downstream_to_Upstream_Throughput = 2 Mbps / 384 kbps \approx 5x

If this ratio were kept at 5x, then the number of downstreams that those same four DOCSIS 2.0 upstreams could support would be:

- $(120 \text{ Mbps} * 5) / 38 \text{ Mbps} \approx 15$ QAM carriers.

This is incredible! Rounded up for implementation, this example matches 16 downstream QAM carriers with 4 upstream QAM carriers. Instead of a ratio of 1x4 that exists with today's CMTS equipment, DOCSIS 2.0 really requires a ratio of 4x1. That is a 16x increase in downstream throughput needed to match the advancements that have occurred in the upstream over the last few years.

As the market begins to accept and use Wideband in the downstream and the number of QAM carriers in the Wideband channel begins to approach eight to 16 channels, fiber nodes have been split to the limit of available fiber, a time will come

The Wideband Protocol for a DOCSIS® Network

when DOCSIS 2.0 will not provide enough throughput. At this time, the Wideband protocol may be applied to the upstream. The deployment of upstream Wideband would be further accelerated both by the movement to symmetrical offered throughput (commercial services) and by when the upstream spectrum is expanded.

Wideband in the Marketplace

The cable technology pot is simmering with alternatives to increase data capacity for next-generation services—both DOCSIS and non-DOCSIS based – to increase the flexibility of the HFC plant. Despite operators being in the early stages of deploying DOCSIS 2.0-based equipment, CableLabs is already working on DOCSIS 2.0 extensions and DOCSIS 3.0 to address such features as MAC layer improvements, roaming, committed data rates, and better commercial service capabilities. Microsoft is initiating a major IPTV push without any DOCSIS component, based on the fact they believe DOCSIS downstreams are too expensive. Numerous cable operators have expressed an interest in an economical and orderly transition from traditional DOCSIS to Wideband DOCSIS services based on the following considerations:

- CapEx minimization—cable operators would like to reuse their traditional DOCSIS infrastructure for wideband services, leveraging existing CMTS chassis and RF line cards when possible
- Spectrum scarcity—cable operators would like to mix wideband and traditional DOCSIS services on common downstreams as they transition customers to wideband services.
- OpEx minimization—cable operators would like to avoid rewiring existing facilities and moving existing customers off current downstreams.

The DSL and FTTx Threat

While the development of advanced DOCSIS technology has been going on, the competition has not been sitting idly by. The telcos, PTTs, and CLECs in particular have continued to evolve and develop their own brand of data throughput nirvana for the last/first mile. However, the growth of broadband services and the promise of the triple play of data, voice and video are driving telcos and PTTs around the world to push their DSL network capacities and/or dramatically drop prices to stay competitive.

However, while the twisted pair DSL solutions have achieved success in basic broadband access and VoIP markets, it is increasingly being viewed as a partial solution that must be combined with a deep fiber strategy if the long-term triple play service is to be realized. Even with an array of alphabet soup solutions including ADSL, VDSL, RADSL, HDSL, G.Lite, ADSL2, ADSL2+, etc., the delivery of true triple play network using twisted-pair-based approach has been elusive, except in high density, short loop length environments.

This is now changing for DSL. Prices for DSL equipment have come down rapidly so that DSL now shares the same economics as DOCSIS. Outside of North America, DSL is now winning over cable for broadband connections. In Korea, 13 Mbps VDSL is routinely deployed against DOCSIS. Thanks to Moore's Law that silicon chips double in density every 18 months, there are newer and more complex DSL standards.




Figure 14: Bonded ADSL2+ Loop Lengths

One of the recent popular developments is called “Bonded ADSL2+”. Two ADSL2+ lines can be bonded to provide 44 Mbps on lines shorter than 5,000 feet. On longer phone lines, two ADSL2+ lines can be bonded to provide nearly 8 Mbps out to 12,000 feet². This is shown in Figure 14. Next up is UDSL (Uni-DSL) which can support 200 Mbps aggregate on short loop lengths (100 Mbps symmetrical), 180 Mbps aggregate at 1,000 feet, and 50 Mbps aggregate at 4,000 feet.³

Even in markets such as Japan, FTTP is an official government program where a number of operators have announced 1 Gbps residential services threatening to destabilize the cable offering. In Korea, the story is the same. Korean Telecom has announced plans to launch an aggressive FTTP program, committing to deliver between 50 Mbps and 100 Mbps to an estimated 73% of Korean households by 2010. In India, Videsh Sanchar Nigam Limited (VSNL), India's leading telecommunications and Internet service provider has announced they will build an ETTN solution to deliver 10/100 Mbps service in the first phase of deployment to over 1 million subscribers.⁴

In the U.S., telcos have turned to partnerships with direct broadcast satellite TV providers to fill the triple play gap, but it is admittedly a temporary solution. The long-term goal for service providers around the world continues to be deep fiber

² <http://www.aware.com/products/dsl/bonded.htm>

³ <http://www.ti.com/udsl>

⁴ http://biz.yahoo.com/bw/041103/35804_1.html

The Wideband Protocol for a DOCSIS® Network

penetration in the form of FTTN, FTTC, or FTTH. SBC's announcement⁵ of plans to deliver FTTH with IP-based data, video, and voice to over 18 million households by 2007 with a service delivery objective of 15 to 20 Mbps,⁶ underscores this objective. This \$6+ billion dollar initiative is surely to be followed by similar announcements and more aggressive expansion programs by the rest of the former RBOCs.

In Europe, service providers do not display the same timidity as the U.S. telcos regarding video or deployment of triple play services via DSL technology. For instance, France Telecom has announced 18 Mbps service and a number of other PTTs with short, average loop length plants and is aggressively deploying triple play services over ADSL2+ or via deep fiber networks. FastWeb⁷, a subsidiary of eBiscom, of Milan, Italy, now offers an integrated package of IP-based broadcast video, on-demand video and network personal video recording (PVR) services, plus telephony and high-speed Internet access, over both fiber and ADSL connections.

So while the U.S. cable industry has spent almost \$80 billion dollars over the last 10 years upgrading and expanding networks, it looks like the race is just about to begin. Whether it is a telco or cable operator, the pipeline is only part of the solution. The driver for all of this is optimizing the average revenue per user (ARPU). Having a highly adaptive and cost effective architecture that adjusts to changing traffic patterns, services, and customer needs is critical. On this front, the cable industry has an obvious advantage with the evolving DOCSIS specification, a large pipeline already in place, and its historic focus on service bundling and compelling content. Nevertheless, this leadership can be squandered if operators and their vendors take their eye off the ball and fail to quickly evolve existing HFC networks to new levels of efficiency and flexibility.

Competition is likely from other sources as well. Wireless and power companies are evaluating their opportunities and strategies in how to get their fair share of broadband customers. The power companies in particular got a recent boost when the FCC approved the delivery of broadband services over the power grid. Broadband over powerline (BPL), while still in its development stages, has had some spotty success in trials, but has potential as a long-term threat if the power industry decides to move forward.

⁵ <http://www.sbc.com/gen/press-room?pid=4800&cdvn=news&newsarticleid=21427>

⁶ http://www.sbc.com/Common/files/pdf/EEW_Super_Comm.pdf

⁷ <http://www.fastweb.it>

Closing Remarks

5 Year Goal

Big Pipes

I would like to propose a 5-year goal for the cable industry to have at least 1 Gbps of downstream data capacity and 100 Mbps of upstream data capacity for IP services.

A gigabit of downstream data capacity with 256-QAM would be approximately 24 QAM carriers. Although this sounds like a lot of QAM carriers, it is, in fact, not an aggressive goal. That is only 24 out of 131 channels, which is only 20% of the downstream channel capacity. Five years from now, 20% of the downstream channel capacity for IP services is not unreasonable.

One gigabit per second shared across 500 HHP with 250 subscribers would be a dedicated or minimum of 4 Mbps per subscriber. While this type of throughput could be easily scaled to a 50 or 100 Mbps service for data, more throughput will eventually be needed when the content is more steady state in nature, such as video streaming.

There are a few interesting aspects to the choice of a GE on the downstream.

- First, if the Wideband channel was on the order of 1 Gbps, then the entire 131 channel downstream could be reduced to about 5 GE pipes.
- Second, the GE on the HFC Plant would map well to the GE backhauls, which would eliminate media rate changes at the CMTS.
- Third, as shown in the first column of Table 4, this path provides a four times, 16 times and finally a 100 times boost in bandwidth when compared to a typical DOCSIS deployment today (1x4 card with 1 FN per upstream) with only 20% of the RF Bandwidth used.

CMTS to HFC Plant Connectivity	Relative Capacity	
1 DOCSIS DS across 24 FN	--	1x
1 DOCSIS DS across 4 FN	1 x	6x
4 DOCSIS DS across 4 FN	4 x	24x
4 DOCSIS DS across 1 FN	16 x	100x
24 DOCSIS DS across 1 FN	100 x	600x

Table 4: Achieving 100x Increase in Downstream Capacity

The Wideband Protocol for a DOCSIS® Network

The second column in Table 4 shows relative capacity from the first DOCSIS deployments seven years ago (1998, 1x6 card with 4 FN per upstream). That means that if the cable industry can deploy four downstreams of DOCSIS per fiber node within 3 years (2008), then the industry would have increased its deployed data capacity by a factor of 100x over a period of 10 years!

On the upstream, 100 Mbps of data capacity can be achieved with four DOCSIS 2.0 upstreams running at 6.4 MHz, 32-QAM, 25 Mbps with a total spectral usage of 25.6 MHz. With an increase of the upstream spectrum from 42 MHz (USA) to 65 MHz or even 85 MHz, this number could come close to doubling.

What about the Asymmetry?

This is a 10:1 ratio of downstream data capacity to upstream data capacity. Some might argue a 2:1 ratio would be better. One of the issues with that is that the capacity of the downstream channel will become artificially constrained by the capacity of the upstream channel. With 2:1, if the upstream maxes out at 100 Mbps, then the downstream will be artificially limited to 200 Mbps.

The HFC plant today is asymmetrical. It has a downstream to upstream bandwidth ratio of 22:1.

$$(862 \text{ MHz} - 54 \text{ MHz}) / (42 \text{ MHz} - 5 \text{ MHz}) = 808 / 37 = 22:1$$

This is because the downstream service model for cable requires this kind of bandwidth. If services such as video move to IP, a high asymmetrical ratio will be required as well. Current statistics show that the upstream and downstream data throughput is almost symmetrical. This is not really true and is an artifact of peer-to-peer networking. Fifteen percent of the upstream subscribers use about 85% of the data capacity. (It is not the subscribers, of course, it is the servers in their homes.) The other 85% of the subscribers generally have asymmetrical throughput with more downstream usage than upstream usage because they are using services from the network instead of delivering services to the network.

In summary, the 5-year goal is to achieve at least 1 Gbps in the downstream and 100 Mbps in the upstream as a shared bandwidth medium. The HFC Plant is asymmetrical, and this will be the service model for some time to come.

Impact to the DOCSIS Specification

PHY

The PHY specs are identical to earlier versions of DOCSIS. There are no RF changes. The protocol just uses more than one PHY.

MAC Management Messages

A new downstream message called a Wideband Channel Descriptor (WCD) is being proposed. The WCD describes the downstream frequencies used for the Wideband channel along with other operating attributes.

Registration

The current proposal uses the existing traditional registration followed by a wideband registration. The Wideband Cable Modem (WCM) identifies its capability, which allows the provisioning server to enable or disable those capabilities.

MIBs

Since Wideband uses Service Flows and the rest of the DOCSIS protocol, the MIB structures for the WCMTS and WCM should be similar to the CMTS and CM.

Impact to Implementation

The Wideband protocol was designed to have minimum impact on existing implementations. As such, the hardware changes supplement, rather than replace, existing hardware.

WCMTS

A new hardware blade within a standard CMTS would be required to support the downstream Wideband framer and the additional QAM carriers. The existing traditional DOCSIS line cards would be re-used.

WCM

A WCM could use the same DOCSIS silicon as a standard CM. It would be supplemented with a multi-QAM carrier tuner and a downstream Wideband framer,

which could be FPGA based. With time, integration can occur to drive the price down.

Summary

The only thing constant about data capacity and throughput is the constant demand for more of it and at a lower cost of ownership. The Wideband protocol for a DOCSIS network re-uses existing DOCSIS PHY and supplements the DOCSIS MAC to build a low cost, scalable, gigabit class DOCSIS downstream. It allows DOCSIS to increase in bandwidth by a factor of 10 times, while allowing the cost of downstream bandwidth to drop inversely by the same factor of 10 times.

There is an incredible amount of bandwidth on today's HFC plant. The challenge is to mine this bandwidth by re-organizing the bits on the HFC plant. The Wideband protocol for a DOCSIS network achieves this goal and sets the stage for a new level of performance for IP connectivity.

During the last eight years, the Cable industry outpaced its competition based upon good economics and solid, standards-based technology. Round one of the fight was dial-up, which was won by the telcos who owned the transport network for the ISPs. Round two of the fight was Broadband and was won initially by the cable industry. Now, it is time for Round three, the Wideband wars with multi-megabit and gigabit speed connections. Who is going to win this time?

About the Author



John T. Chapman is currently a Distinguished Engineer and the Chief Architect for the Cable Business Unit at Cisco Systems in San Jose, California. As a founding member of the Cisco Cable BU, John has made significant contributions to Cisco and the cable industry through his pioneering work in DOCSIS and development of key technologies and concepts critical to the deployment of IP services over HFC plants. Included in these achievements are being the primary author of significant portions of the DOCSIS and PacketCable specifications as well as the originator of DOCSIS Set-top Gateway (DSG) and evolving specifications for DOCSIS Wideband and Modular CMTS architectures for the industry's Next Generation Network Architecture (NGNA) initiative. John has also published a number of ground breaking whitepapers on Multimedia Traffic Engineering (MMTE), DSG, QoS, and high availability and is a respected and frequently requested speaker at industry events.

John graduated from the University of Alberta in Canada with a Bachelor in Science of Electrical Engineering in 1984. John began his formal career designing voice line, trunk, and ISDN cards for ROLM/IBM/Siemens and joined Cisco Systems in 1989. At Cisco, before joining the Cable BU in 1996, he was responsible for the development of HSSI (High-Speed Serial Interface) which is now an ANSI and ITU-T standard (V.12). Later projects included the industry's first ISDN BRI router, the design of the multi-protocol serial interfaces and HDLC ASICs used across the Cisco router product line, and co-architecting the VIP card in the Cisco 7500 product line.

John has 18 patents issued and 27 patents pending in a variety of technologies including telephony, VoIP, wide area networking, and broadband access for HFC cable networks. In his spare time, John enjoys spending time with his wife and two daughters. John is a 6th Degree Black Belt Master in Tae Kwon Do and enjoys white water canoeing and skiing.

Previous papers by John may be found at <http://www.johntchapman.com>.