UNIFIED DATA AND VIDEO CMTS: ONE SYSTEM FOR ALL SERVICES

Scott Cummings and Victor T. Hou
Broadcom Corporation

Abstract

This paper opens with a quick history of the traditional Head End (HE) and current CMTS platforms in the HE. The "loosely coupled, tightly bolted" architecture of the traditional HE is exposed. The issues in adding new services and technologies will become apparent.

Then an overview of several current and new technologies entering the next generations of HEs is listed. Each technology is described. Then the Modular CMTS (M-CMTS) architecture is applied to the Next Generation Network Architecture (NGNA) for the HEs. The traditional CMTS is subdivided into three modules: CMTS core, Edge QAM, and Upstream Receiver. A new component to the HE is defined: a System Resource Manager (SRM).

The incoming technologies are applied to the SRM and M-CMTS. As these technologies are incrementally added to the SRM and M-CMTS, the changes to support them are simple. Adding servers, SRM interfaces, and possibly CMTS cores and the system is ready for the new technology or service. The combination of the SRM and M-CMTS convert this HE architecture into the unified CMTS, a system for all services.

Traditional Head End (HE)

The traditional HE is composed of many devices that offer several services. Video, Data, and Voice services are all supported by the HE. Each of these services and the equipment used to deliver these services will be quickly examined. The term HE is used a bit loosely here. There are distribution hubs and other types of nodal locations that may provide the functions described in this paper as a HE.

Cable Modem Termination System (CMTS)

The current DOCSIS 1.x, and 2.0 CMTS platforms predominantly deliver data service for cable modems. These systems operate with a downstream modulation type of either 64-QAM or 256-QAM in the downstream. The downstream channels support either 6 MHz (North American plants) or 8 MHz (European plants). The CMTS also provides some networking services for the cable modems (DHCP, ToD, Gateway). Some CMTSs incorporate these servers, others do not. DOCSIS 1.1 introduced Quality of Service features to the CMTS such as rate shaping and per user rate limiting.
The CMTS also provides upstream data services. DOCSIS 1.1 is limited to a maximum bandwidth for a single channel of 10 Mbps. To provide symmetric bandwidth with the downstream, the DOCSIS 1.1 CMTS platform are designed for a 1:N downstream to upstream ratio, where N would be 4, 6, or 8. DOCSIS 2.0 introduces many upstream modulation types which now allows the CMTS to offer QPSK, 8-QAM, 16-QAM, 32-QAM, 64-QAM, or 128-QAM (SCDMA only). However, the downstream to upstream ratio has not changed.

The current CMTS platforms are composed of a rack of blades. Typically the CMTS contains a network blade, a processor blade and a physical layer blade. Additional blades can be added to the rack to meet the needs of the node or head end supported. The physical layer blade maintains the 1:N ratio. These are standalone systems that merely share the Hybrid Fiber Coaxial (HFC) cable plant.

The current CMTS platforms also delivers Voice over Internet Protocol (VoIP) service. The VoIP services are bidirectional services that have restrict latency and jitter requirements relative to the general data services. Due to these requirements, the VoIP services are often segregated from the data services. This segregation may be as simple as dedicated channels for VoIP, or as rigorous as a dedicated CMTS for the VoIP services.

The CMTS systems make up only a small part of the Traditional HE.

<u>Video Distribution</u>

The majority of equipment in a traditional HE is for video distribution. There are two distinct types of video distribution in a HE.

The first is broadcast video. The broadcast video may be delivered with traditional NTSC analog signals, or with digital video MPEG transmissions. In most HE today, a combination of both analog and digital signals are transmitted on the same HFC plant.

The second type of video is narrowcast video. The narrowcast video is composed of Video on Demand (VoD) and other selectable video services. These services are not streaming video all the time. A given video stream may be streamed on multiple RF channels. Billing systems for the individual purchase of video streams are also required for these services. The narrowcast video services often also require a return path for ordering and billing confirmation. Narrowcast video may also be used to send programs only when they are viewed. Rather than broadcast all the channels available to the subscriber, a future model for narrowcast video would send only the channel requested in either multicast or unicast streams.

Ignoring the equipment just to receive video (multiple satellite dishes, local antennas, high speed fiber connections, downconverters), the HE is composed of multiple QAMs, upconverters and MPEG multiplexing stations. The MPEG multiplexing stations may be simple multiplexing units or very sophisticated and expensive units that are capable of transcoding and transrating MPEG streams in real time.

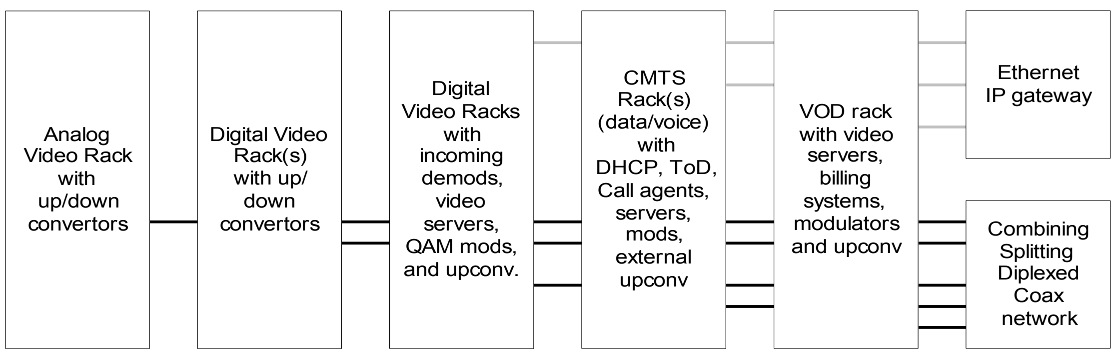

Traditional HE system level diagram

The distribution of broadcast video and narrowcast video both use the same type of equipment. The billing, switching, and return path requirements of the narrowcast video generally create implementations for these systems that are stand alone systems relative to the broadcast video systems.

The CMTS systems providing data and VoIP services also tend to be independent systems from the video distribution systems. This creates a collage of equipment in the HE. This also creates operational issues for maintaining a HE. Each system requires individual expertise to be maintained and operated. As new technologies and/or new solutions for existing services become available, these systems tend to contain requirements that lead to stand-alone implementations as well.

New Technologies

There are several new technologies that are currently available or are soon to be available to the next generation of HEs. These technologies include DOCSIS Set-top Gateway (DSG), Internet Protocol (IP) video, Channel Bonding, the Personal Video Recorder (PVR), MPEG-4, part 10 (AVC) digital video compression and 1024-QAM downstream modulation. Each of these innovations has targeted applications and benefits.

DSG

The DSG technology provides a DOCSIS tunnel to any device connected to the HFC plant. This tunnel is a 1 way (downstream only). However if the CPE device contains some form of an embedded CM (eCM), a 2 way communication path can be implemented. (both downstream and upstream). This opens a development path for the HE to control devices other than Cable Modems (CMs) with DOCSIS. DOCSIS has provided the MSO with the certification process to guarantee compatibility across multiple vendor devices. The certification and qualification process also delivers a very robust and reliable system. The Consumer Premise Equipment (CPE) connected to the HFC plant can now all be controlled using this DSG tunnel.

The DSG technology is available today. The DSG effort has been active at CableLabs and the specifications are complete.

IP video

The concept of delivering video over IP on an HFC plant has been around for quite some time. IP video dominates video delivery to Personal Computers (PCs). QoS, buffering issues, and overhead issues have all hampered the acceptance of IP video. IP version 6 (IPv6) has addressed some of the QoS issues. Memory and disk space technology continue to follow Moore's law. The cost for a bit of RAM or disk space continues to drop. This has greatly reduced the buffering issues with IP video, and will continually reduce buffering issues well into the future. The additional overhead for the IP headers can be addressed with Payload Header Suppression (PHS) in DOCSIS CMTS platforms.

The power of IP video comes from a couple of sources: economies of scale, and the robust feature set of the IP protocol. The IP technologies are ubiquitous. These solutions are extremely cost effective from an equipment standpoint. IP video also automatically provides several avenues for advanced video services. These services include stopping, pausing, forwarding, rewinding, book marking, and other video control features. The IP video can be served to a wide range of devices. PCs, STBs, and even Hand Held Devices (HHDs) are all capable of processing IP video.

PVR

The PVR has been available to the consumer for over a year now. The original PVRs provided storage to hold recorded video programming and a simple software Graphical User Interface (GUI) to help program and manage stored recordings. The early PVRs had limited storage and therefore did not offer many features beyond a mechanism to record and store video.

This is similar to the MP3 player. The MP3 player started as a simple device to store and record audio. Over time, the MP3 player has expanded to be a general purpose storage device that has the capability of playing and recording audio.

The PVR has the opportunity to take a similar migration from a dedicated video device to at general purpose storage devices that can play and record video, audio, still photos, and even display web pages. The PVR actually could be the link between the STB and the PC. The STB has had a processor, memory, an Operating System (OS), and a GUI for several years. Many of the later models of STB have Ethernet and/or USB ports on them. Keyboard style remote control units exist for both STB and high end TVs. The lone missing piece of a PC in an STB is the disk drive. The convergence of the Web and Video, the STB and the PC, could all be in the PVR.

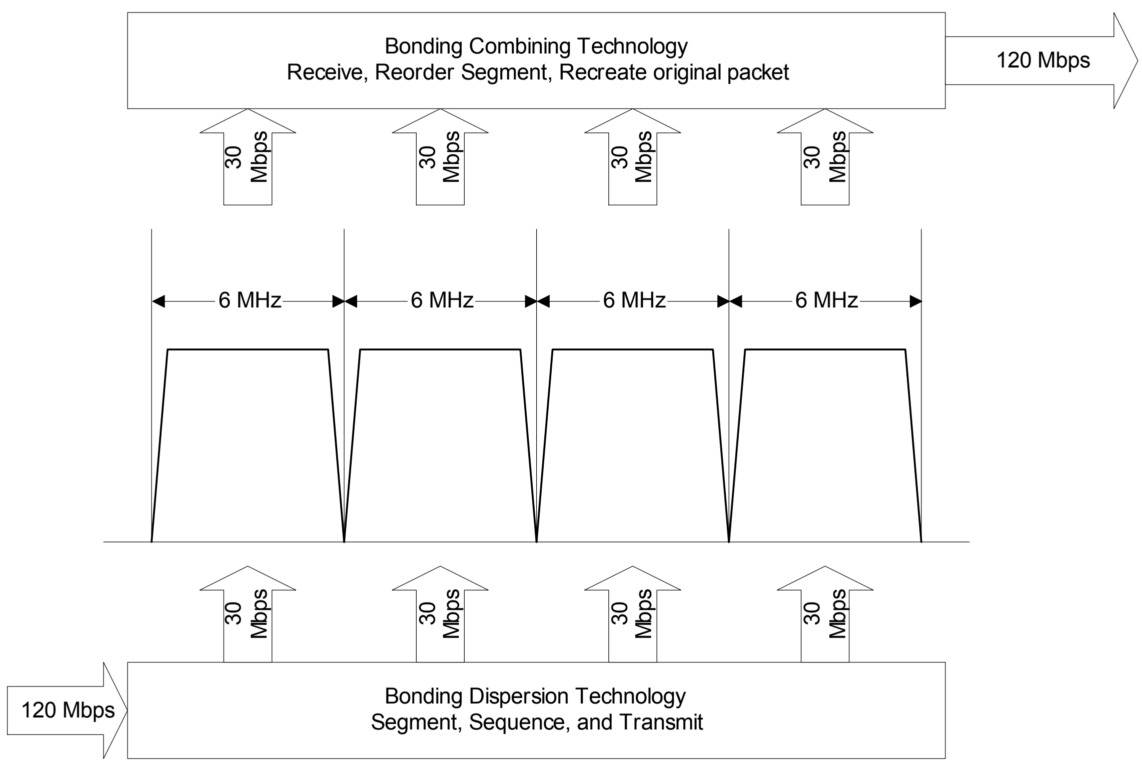

Diagram of channel bonding a 120 Mbps load across 4 downstream channels

Channel Bonding

Channel Bonding is a new technology that is part of the discussions for DOCSIS 3.0. The concept of channel bonding is fairly simple. Channel Bonding bonds N number of existing channels into 1 larger virtual channel. This provides several features. First, the maximum throughput to a single CPE will be increased by a factor on N. An MSO could offer a high speed service of well over the 38.8 Mbps that a single 256-QAM downstream channel can provide today.

Another feature of Channel Bonding is the statistical multiplexing gain of having more users share a wider channel. Statistical multiplexing benefits occur when a channel gets oversubscribed. Service on that channel has to be reduced or stopped for some users. If another channel was undersubscribed, the excess load could be taken by the second channel. Today, instantaneous load balancing is not available in CTMS platforms. If both channels could be virtually combined, the new bonded channel would not be oversubscribed and no reduction or loss of service would be observed by the users.

The granularity of the Channel Bonding solution may also provide other benefits. For Channel Bonding to work, the CPE must have the ability to simultaneously receive multiple channels. With more channels to select where and when to place downstream packets, latency and jitter could be better controlled. This is similar to the statistical multiplexing gain with the expectation that the system has not been loaded to the point of over subscription. As the system is approaching the maximum subscription rate, this level of the loading or congestion begins to create an increase in packets that request the transmission at the same transmission time. Clearly in a single channel, two packets cannot be transmitted at the same time. One packet will be transmitted, and the other packet will either be latent or get advanced. In this case, jitter has been introduced. As the system approaches the maximum subscription rate, jitter and latency begin to increase.

A Channel Bonding solution that has MPEG layer granularity and the freedom to select the channel at transmission time can send multiple packets at the same time. If a channel is at the maximum subscription rate, the channel bonding system can use the other channels to help get critical packets downstream with minimal jitter and latency effects.

<u>1024-QAM modulation</u>

Many new technologies have been applied to the HFC plant. No advancement in DOCSIS from its inception has provided improved bandwidth in the downstream. PHS can be applied in the downstream to save some bandwidth. The only improvements in bandwidth in the downstream over the past 5+ years has been the expansion of the HFC plants frequency spectrum. He original HFC plants started with an upper frequency limit of 550 MHz. This has expanded to 650, 750, 850, and now there are discussions of opening up to more than 1 GHz. The frequency expansion has hardly been a new technology. Replacing cable, connectors and amplifiers with cleaner, equipment specified up to the higher frequencies is not an architectural technology improvement.

A more efficient modulation scheme can be introduced to the HFC plant. Currently, the highest modulation scheme available to an HFC plant is 256-QAM. This provides 8 bits per symbol. 1024-QAM provides 10 bits per symbol. Upgrading the modulators and demodulators to 1024-QAM would improve the raw bandwidth by 25%.

Beyond the straight bandwidth gain of the modulation order, a statistical multiplexing gain will also be accrued. To get the 1024-QAM modulation to operate, an HFC plant with no margin will have to supply 6 dB more Signal to Noise Ratio (SNR). This begs the question: Is it easier to extend the frequency in the HFC plant by 25% or cleanup up the HFC plant for 1024-QAM? Actually, it would be synergistic to do both and gain (1.25 * 1.25 = 1.5625) 56.25 % more bandwidth.

MPEG 4 Part 10 (AVC) video compression

The MPEG-4 Part 10, advanced video compression (AVC) standard provides superior compression versus MPEG-2 with the same resulting video quality. The improved suppression will allow more AVC streams per RF channel versus MPEG-2. AVC, also known as MPEG-4 Part 10, MPEG-4 AVC or ITU-T H.264, is an extension to MPEG-4. AVC offers many new compression techniques to provide equivalent quality to MPEG-2 at half or less than half the bandwidth. The most effective tool for better compression is the new mechanisms provided for advanced prediction techniques. Clearly, for digital video, both MPEG Transport Streams (TS) and IP video, the additional compression saves bandwidth.

AVC offers another bandwidth saving technique. The compression of video can be coded with a Multiple Description coding technique. Rather than code the video into a single stream, the video can be coded into multiple streams. Each stream is then sent to the decompression device. The decompression device can then use all the threads to create the resulting video. However, the decompression device may not use all of the streams to decompress the video. The resulting video will have lesser quality relative to the video using all the streams. The Multiple Description technique could be used to compress the video at different pixel resolutions. The first stream would produce video for a device with 180 lines of video (PDA, HHD, …). The second stream could then be decompressed to generate video with 360 lines of video. The next stream could then generate a resulting stream of 720 lines of video.

Each stream then is sent downstream. A device with limited screen resolution would decode the only the first stream. A device with moderate screen resolution would use the first two streams, and a device with HDTV screen resolution would use all of the streams. The sum of the three streams would likely be larger than the size of a single stream targeted for an HDTV application. However, the sum of the three streams is likely to be smaller than the sum of the three streams required to support video across the range of three devices in this example.

The Multiple Description feature of AVC is still under academic exploration. The technology of this technique is not available today.

Digital TV and HDTV

The large screen digital TV phenomenon is here. High Definition Television (HDTV) is the consumer buzzword attached to these monster TVs. HDTV signal transportation will load the existing HFC plants. Simultaneous, Standard Definition Television (SDTV) and HDTV transmission will burden the HFC plants.

In all of the excitement of these new bigger and better TVs, the catalyst to this technology has been overlooked. The Federal Communications Commission (FCC) has mandated by 2007 that all over-the-air broadcast TV transmissions be digital. This requires the TV manufacturers to be delivering digital ready TVs well before 2007. The one-way communications specification has already been agreed upon, and the two-way communications specification to these digital TVs is therefore implied to arrive by 2007.

The digital TV mandate provides some relief and some consternation for the MSO. An opportunity to migrate to an all digital HFC plant is clearly opened with this mandate. This converts the bandwidth inefficient analog channels into bandwidth efficient digital channels. This bandwidth relief is short lived. The digital TVs will be

HDTV capable and the consumer will be demanding more HDTV programming. An HDTV video stream ranges from 4 to 5 times larger than an SDTV video stream.

The two-way digital TVs will have the opportunity to become more than just video playing devices. The advanced digital TV may offer program guides, web page capability, or text streaming services. These are low bandwidth services. The digital TV could evolve into a device with specialized gaming electronics that allow internet gaming. As the digital TV gets more complex, the bandwidth requirements of the HFC plant will continue to rise.

Next Generation Network Architectures

To begin to deal with all of the new technologies that either are destined to be implemented in a HE or simply loading the requirements for the HE, the architecture of the HE itself must be reassessed. As noted earlier in this paper, the HE has grown into a collection of systems sharing the HFC plant. To reduce both operating expenses and the expense of turning over equipment in the HE, new architectures have been proposed.

The Unified CMTS via the modular CMTS

The beginning of this new architecture is being formed with the CableLabs Modular CMTS (M-CMTS) specification. This specification is vital in the development of future CMTSs. The intent of future CMTS architectures is to allow it to become a central entity in more than just data. The next generation CMTS must be able to control, maintain, and organize all of the streams transmitted and received over the HFC plant, not just voice and data. The unification of service control is the target of the next generation CMTS.

To efficiently create access to every CPE device, the CMTS must have access to the entire HFC plant. The traditional CMTS contained QAM modulators and upconvertors and receivers embedded into the CMTS in hardware driven ratios. For the unified service providing CMTS, embedding an entire HFC plant worth of QAMs, upconvertors, and receivers would not be cost efficient. Two systems sharing the same downstream channel must be managed. Routing all traffic through a CMTS regardless of the need for the CMTS in the downstream path is not cost efficient.

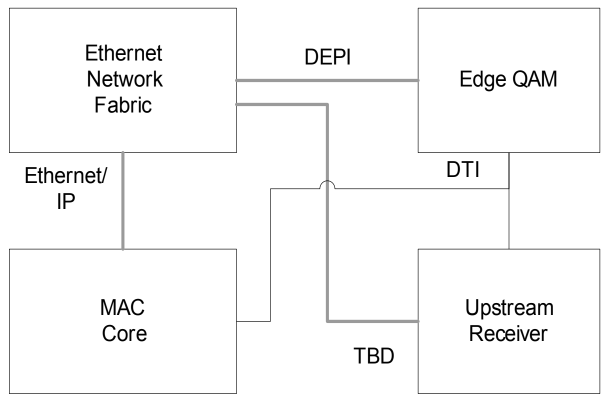

Diagram of M-CMTS

To deal with the cost and manageability of the downstream of the HFC plant the modular CMTS is required. The modular CMTS takes all of the functions of the traditional CMTS and subdivides these functions into three modules: CMTS core, Edge QAM, and upstream receiver. The CMTS Core provides all of the MAC functions required by DOCSIS. Classification, rate shaping, security, header suppression, routing, bridging, and other packet level functions. The Edge QAM performs all of the downstream PHY requirements of DOCSIS. The Edge QAM may also be accessed by other streaming devices other than a CMTS core. The upstream receiver performs the upstream PHY requirements for DOCSIS.

The Edge QAM

The key to this architecture is the Edge QAM. It is arguable whether the Edge QAM is part of the CMTS or really should be considered the stream-to-RF bridge to the HFC plant. Assuming an all digital HFC plant, the plant may have as many as 160+ QAMs on that plant. These QAMs may be transferring video, user data, voice, application information, and/or plant management information. All of this traffic does not have to flow through a traditional CMTS. However, the HFC plant may have CPE devices on any QAM that are in need of DOCSIS support. The CMTS must have access to the Edge QAMs.

The Edge QAM must be accessible by multiple devices that are providing a wide array of services over the QAM channel. The ideal Edge QAM solution would be to have an Ethernet port for receiving digital traffic and an RF connector to transmit the traffic onto to the HFC plant. An Edge QAM device may support a single QAM or multiple QAMs. The Edge QAM must be able to place the QAM signals on the appropriate frequencies, and there require some upconverting capabilities.

The Edge QAM device is really a sophisticated Ethernet Network Interface Card (NIC) for the HFC plant. This should allow the vendor community to increase the density of QAMs per RF connector, and reduce the cost per QAM. The MSO now can purchase the number of Edge QAM blades required to provide the necessary number of channels. If the plant extends the frequency range, then more Edge QAMs can be plugged in at that time.

The key to the Edge QAM versatility is the interface. The Edge QAM must provide some MPEG multiplexing capabilities. This will allow multiple devices to share an individual QAM or QAMs. The Edge QAM must have low latency Queues for video transport. The edge QAM must also have prioritized queues for data and management message transport. The QoS capabilities provided by DOCSIS must not be hindered by the Edge QAM. The M-CMTS specification provides two separate constructs for transmission.

The first construct is an MPEG Transport (MPT) mode. This mode is expecting MPEG elementary streams wrapped in a UDP datagram. Some video processing to de-jitter the incoming video and correct the PCR is required for these streams. This processing is required by the Edge QAM. The interface also provides a transparent mode that does not get processed by the video processor or DOCSIS timestamp function. The interface also provides an MPT operation for DOCSIS packets. The DOCSIS MPT operation would have some resyncronizing circuitry to maintain timing. Behind all the MPT processing is a MPEG multiplexer to merge the streams onto the QAM channel.

The second construct is a Packet Stream Protocol (PSP) mode. This mode preserves the DOCSIS QoS features. DOCSIS MAC management message are typically prioritized ahead of user data on the downstream channel. To preserve this prioritization multiple queues that can be prioritized must exist in the Edge QAM. The PSP mode of operation also differs from the MPT mode in that the PSP flows are terminated at the Edge QAM where the MPT flows are transparent and not terminated. The combination of modes provide the Edge QAM the versatility it needs for different applications.

The CMTS Core

The CMTS Core provides all the CMTS core functions in the traditional CMTS. The CMTS core also is burdened with the channel bonding process. The M-CMTS has placed the channel bonding in the CMTS core versus in the Edge QAM or a combination of both. The CMTS core will have an Ethernet interface to the HE network and another Ethernet interface to the Edge QAM network.

The CMTS core will bridge or route packets onto the HFC plant (via an Edge QAM). The CMTS core will classify, rate shape, police, bond, encrypt, and header suppress packets. The intent of the M-CMTS architecture is to have CMTS core implementations that are scalable. As more DOCSIS processing is require for a given HE, more CMTS cores could be added to the HE.

The Upstream Receiver

The definition of the interface between the CMTS core and the upstream receiver is will be considered in the future. The downstream division was fairly simple. Leave DOCSIS in the MAC, so the Edge QAM was simple and could be used by non DOCSIS entities. Applying the same philosophy to the upstream, the interface would lend itself to be some sort of FEC block wrapped in an Ethernet packet. Then the CMTS core could rebuild the packet. However, reusable upstream receivers with non DOCSIS upstream devices may not be preferred.

The intent of modularizing the CMTS is to allow it to have access to the entire plant and become the standard control mechanism inside the HFC plant. With this in mind, it may be preferred to have the upstream receiver return Ethernet encapsulated packets with DOCSIS information wrapped around the completed DOCSIS packet. This would allow the CMTS core to be a management tool versus a packet processor.

The concept of having the Edge QAM provide DOCSIS packet processing has been dismissed for the current revision of the M-CMTS specification. However, there is merit to revisiting a different definition of the Edge QAM. The current Edge QAM definition does not seem to take advantage of the capability of multiple Edge QAMs on a single blade. Concepts like channel bonding require the CMTS core to packet process the bonded stream and send the dispersed streams to individual Edge QAM devices. This prevents channel bonding, at the Edge QAM, from dynamically selecting the appropriate channel to send packet segments.

The Edge QAM device may actually support multiple QAMs per Ethernet port. The system would be far simpler if the CMTS core managed and routed packets, and the Edge QAM prepared them for transmission. A channel bonding mechanism in PSP mode would allow the Edge QAM to pull packets from a multi-QAM queue and put the channel bonded segments onto the bonded QAMs. This would allow the Edge QAM to optimally balance the QAMs versus the CMTS core making some crude estimation of where to send the channel bonded segments.

Putting the M-CMTS pieces together

The modularization of the M-CMTS provides architectural advantages. Another hidden advantage to this architecture is the connectivity of the parts. The CMTS Core, Edge QAMs, and Upstream receivers are all connected together with an Ethernet network fabric. Just as the Edge QAMs can be scaled and shared with other services in the HE, this network fabric can as well. By leveraging the commodity priced technology of the Ethernet network fabric, the MSO gets an inexpensive solution with excellent versatility. The technological and business success of the Ethernet network fabric clearly speaks for itself.

System Resource Manager

The M-CMTS is perfectly tuned to leverage shared Edge QAMs. In the traditional CMTS, the CMTS managed the load on the CTMS QAM or QAMs. Now, the M-CMTS cannot really manage all of the Edge QAMs. The load on the Edge QAMs may not be known by the M-CMTS. A new device must be put into place to manage the loading of the Edge QAMs. This device will be called a System Resource Manager (SRM).

The SRM will manage the video servers, audio servers, IP server, or any server or device streaming content over the Edge QAMs. The SRM must have the capability to dynamically start and stop services, and dynamically route services to the Edge QAMs. The SRM may also need to monitor congested links in the Ethernet network fabric to prevent over subscription of that fabric. Again the SRM will have to perform the rate shaping, load balancing and manage the routing of the streams in the same manner that the Traditional CMTS provides these functions for data traffic over DOCSIS.

New Technologies and the M-CMTS

The intent of modularizing the CMTS was to create a HE architecture that would leverage all of the current and new technologies seamlessly. To demonstrate the value of this M-CMTS architecture, the afore mentioned technologies will be applied to the SRM and M-CMTS architecture. As each technology is added, the SRM and M-CMTS will leverage the combination of the technologies. The result will be that the SRM and M-CMTS will have unified all of the services and technologies under one architecture.

DSG

In an ever competitive world, the cost reduction of the CPE device is inevitable. This will create CPE devices on the HFC plant with only one tuner. This tuner may be tuned to any QAM. For a HE to provide access for DSG messages to every QAM on the HFC plant, that HE would require a traditional CMTS for every QAM. The M-CMTS has several techniques to handle this issue without a traditional CMTS per QAM.

If the DSG CPE device is a one-way device, then the M-CMTS is required to replicate all DSG messages destined for that device on all the QAMs. This requires a single CMTS core receiving and forwarding the DSG message or messages. The messages get tunneled through the CMTS core and on to the existing Edge QAMS, with the message going to each Edge QAM. The load on the CMTS core would be minimal.

If the DSG CPE device is a two-way device that allowed the SRM and M-CMTS to track the QAM the device is receiving, then the DSG server could send the message to the appropriate QAM. A point to point message sent through a single CMTS core and routed to the proper Edge QAM is highly efficient.

PVR

A simple example of leveraging the PVR with DSG over the M-CMTS would be to download program guides to CPE devices. This could be as simple as the show listings per channel for video, or music listings for audio devices.

Using the techniques listed with DSG this would be extremely simple. This would allow operational content to be stored on the PVR. This content could be updated at the MSO convenience.

Channel Bonding

Channel Bonding provide extremely high speed service to a single device. Leveraging DSG, PVR and the M-CMTS, channel bonding will be used to download a movie as an impulse purchase.

A customer orders a movie from a service provided on the CPE device. The application and the data required for the application can be downloaded and stored using DSG and PVR. Now leveraging the DSG tunnel to make the transaction, the movie could be downloaded in around 3 minutes bonding 10 QAMs. The downloaded movie would be stored on the PVR on the CPE device. The trailer for the movie could be downloaded in seconds and playing while the movie itself downloaded. The customer would have the option of watching the movie right after the trailer or viewing the movie at a later date.

The movie could be stored in the HE as an MPEG stream on a MPEG server. The DSG tunnel is used to inform the movie ordering application of the address of the appropriate servers to complete the movie ordering transaction. Once the transaction has been completed, the SRM can get the appropriate MPEG server to stream video to the appropriate Edge QAMs. The SRM can then prevent any further activity from oversubscribing any given Edge QAM. This example again takes very little processing from a DOCSIS CMTS core and simply leverages the network connectivity in the M-CMTS HE to get the video to the appropriate Edge QAM.

1024-QAM

There is no need to glorify this, just take the last example and go 10/8 times faster. Now instead of 10 channels in the prior bonding example, only 8 are required.

HDTV and AVC

The example now grows into having multiple CPE devices on the HFC plant. Some devices require MPEG-2 video, some devices can receive MPEG-2 or AVC, and some devices are AVC only devices. The only real difference now is that the DSG tunnel is used to communicate multiple applications that in turn will make requests to the System Resource Manager (SRM). The SRM then configures the multiple video servers to stream the video to the appropriate groups of bonded channels for each service required. The number of MPEG-2 and AVC servers can dynamically change over time. The HE architecture and the CMTS tunneling commands to the CPE devices do not need to change. This leads to a cost efficient migration path from MPEG-2 to AVC in the HE. The same strategy will apply to MPEG-N in the future.

A second example of HDTV, AVC and the M-CMTS will mix HDTV and SDTV broadcast programming using AVC over the HFC. The assumption will be that the program that is being streamed is a broadcast program that can be viewed by several types of CPE devices. The Multiple Description feature of AVC can generate streams for HHD, SDTV STB and HDTV digital TVs. The SRM then directs the streams to the appropriate QAMs. If the CPE devices requiring more than 1 stream are channel bonding devices, then the streams can be dispersed across multiple QAMs. The HHD devices that require only a single stream could be instructed to find that stream on the appropriate QAM. The SDTV STB and HDTV digital TVs can be instructed to pull off the streams of interest off of the appropriate bonded QAMs. This is using bonding as a mechanism for picking up multiple streams as opposed to picking up a single stream striped across multiple channels. This system is using AVC for its compression and Multiple Description streaming benefits. The channel bonding can now be used as a load balancing tool rather than a burst performance tool. The DSG tunnel is used to direct the CPE devices to the desired streams. All of this is accomplished with the SRM and M-CMTS architecture.

IP Video

To make a long story short, integrating IP video takes nothing more than replacing MPEG video servers with IP video servers and providing the appropriate number of CMTS cores to support the IP Video through DOCSIS requirements. IP video requires more buffering, but the PVR or the never ending improvements in memory densities can both be implemented to deal with any IP video buffering issues.

The Big Ending

The big ending has already gotten to be a bit boring. As the examples continued to pile on new technologies and services, the system quite frankly easily handled them. The key to providing more services became a matter of having the source servers. The control and distribution mechanisms did not change. VoIP could be added to the mix. The boring solution is to add the required CMTS cores and upstream receivers to handle the load. If you already have idle CMTS cores or upstream receivers, then just use those. If VoIP declines and an IP gaming service picks up, then the system replaces the voice traffic with gaming servers and gaming traffic. Again, add the server for the service that is required and the problem is solved.

Boring, scalable solutions may not make for a dramatic ending of a technology paper, but they would be a dream to the current cable operator. Dynamically changing services are an operator nightmare. New proprietary equipment means lengthy bring up time, expensive installing and maintenance costs, and lengthy Return on Investments (ROI). Installing a server for a service that is in demand today is both time and cost efficient. The ROI could be measured in weeks not years. Somehow the boring M-CMTS has become the cable operators' Holy Grail.

Wow, all of this supplied by an M-CMTS architecture. When the M-CMTS architecture is viewed from the entire HE, it becomes the unifying piece to the multiple service puzzle. This architecture allows the unification of services through the M-CMTS. Adding a SRM to an M-CMTS puts the finishing touches on the Unified CMTS, a system for all services.

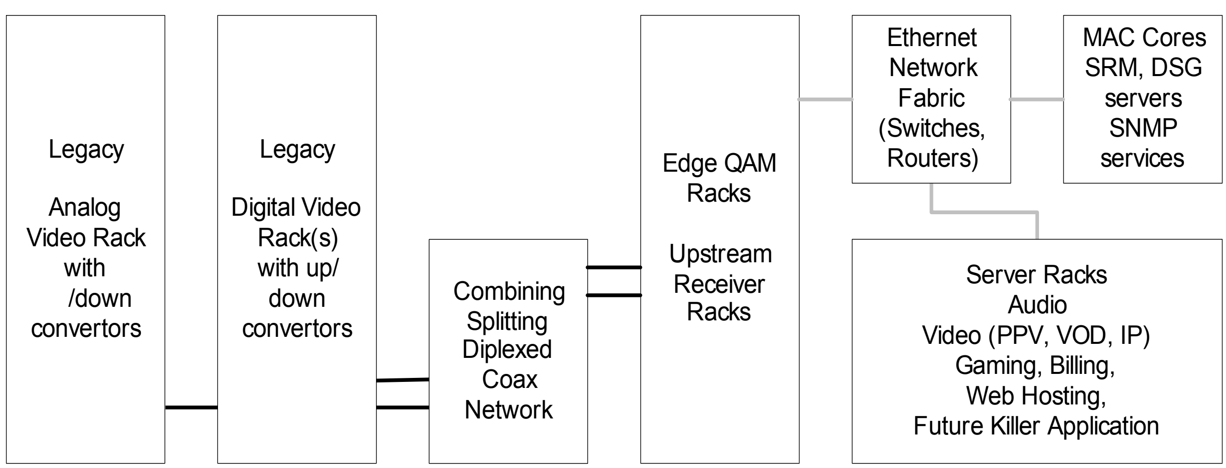

Diagram of the Unified CMTS, the system for all services