

Maximizing Legacy and DOCSIS 3.0 Bandwidth in the Forward Path using M-CMTS and EdgeQAM Technologies

Jeff Finkelstein, Cox Communications
1400 Lake Hearn Drive NE, Atlanta, GA 30319
email: jeff.finkelstein@cox.com phone: +1.404.843.5955

John T. Chapman, Cisco
170 West Tasman Drive, San Jose, CA 92677
email: jchapman@cisco.com, phone: +1.408.526.7651
web: www.johntchapman.com

Abstract

In the traditional DOCSIS model, new downstream capacity required hardware capable of supporting both downstream and upstream channels. These MAC domains required a higher entry cost partially related to the complexity of the hardware components needed at the CMTS, head-end and HFC plant. With DOCSIS 3.0 we now have the ability to increase downstream capacity using modular components that reduces cost of the equipment providing the new bandwidth, simplifies future cabling and enables converged services.

The authors have been involved in DOCSIS since its inception, through actual field deployments, development of the specifications and architecting solutions for future growth including convergence of video and data services.

In this paper the authors hope to assist cable operators in finding ways to use DOCSIS 3.0 capabilities and maximize bandwidth availability for new and legacy services. We focus on identifying key selection criteria for deploying M-CMTS and EdgeQAM technologies.

1.0 DOCSIS 3.0 Evolution

As an introduction we would like to highlight the key changes from DOCSIS 1.0 to DOCSIS 3.0.

1.1 DOCSIS 1.0

The focus of DOCSIS 1.0 was to create a standard structure for providing high-speed internet services. The primary features of DOCSIS 1.0 included standards based provisioning, SNMP, standards based encryption using DES, and rate-limiting features per subscriber to control access speeds. [DOCS10]

1.2 DOCSIS 1.1

DOCSIS 1.1 included new features that allowed cable operators to provide services that relied on quality-of-service. The primary focus was on supporting voice services, but the new features were designed to be capable of supporting any service requiring guarantees for jitter, latency and packet loss.

The key features provided in DOCSIS 1.1 included quality-of-service controls via service flow level provisioning, dynamic services with admission control, concatenation, fragmentation and payload header suppression. [DOCS11]

1.3 DOCSIS 2.0

Where DOCSIS 1.1 enabled new services using MAC layer enhancements, DOCSIS 2.0 focused primarily on the physical layer to increase upstream bandwidth.

The key enhancements to the upstream physical layer in DOCSIS 2.0 include higher order modulations, increased symbol rates, new modulation types including A-TDMA and S-CDMA, logical upstreams and autonomous load balancing. Another key improvement as part of the new modulations was the ability to use a 24-tap FIR equalizer, which has the potential to improve MER in the presence of linear distortions. [DOCS20]

1.4 DOCSIS 3.0

DOCSIS 3.0 contains significant enhancements to the entire core specification and touches every component in the DOCSIS network. It is important to note that in the DOCSIS 3.0 specifications there is nothing that infers a modular architecture. [DOCS30PHY] [DOCS30MULPI]

1.4.1 MAC Layer

The feature most identified with DOCSIS 3.0 is channel bonding in both the downstream and upstream direction. This provides the ability to create a single or multiple channels capable of speeds equal to the sum total of channel bandwidth. For example, if four 40 Mbps downstream channels are bonded into one large channel, the potential aggregate bandwidth is 160 Mbps.

A key feature of channel bonding is the potential to use statistical multiplexing to gain more efficient bandwidth distribution across the channels associated with the bonding group. One important by-product of statistical multiplexing is that it not only provides for speeds greater than can be achieved by a single channel (e.g. > 25 Mbps) by bonding the separate channels into a single large logical channel, but for a typical data rate that would traditionally be provided in a single channel there are efficiencies gained by taking advantage of multiplexing to gain more better utilization of non-bonded capacity. In short, if you use bonded channels for our current data rates, you gain efficiency by self-averaging of peak to average ratios.

1.4.2 Network Layer

IPv6 is a key driver for many to deploy DOCSIS 3.0. With IPv4 there are 2^{32} possible addresses, but with the increasing number of devices this address space is being rapidly depleted. With IPv6 there is support for 2^{128} addresses, which equates to everyone on Earth having their own 5×10^{28} addresses. IPv6 allows for either provision of addresses to occur through DHCP or through stateless auto-configuration.

IP Multicast support adds features including multicast QoS, intelligent multicast routing and source specific multicast. These capabilities enable the operator to provide broadcast-like services more efficiently and with a guaranteed QoS than previous DOCSIS versions.

1.4.3 Security

Security enhancements in DOCSIS 3.0 include certificate revocation management, runtime validation, AES for enhanced traffic encryption and certificate convergence.

1.4.4 Network Management

Extensions to the IPDR SAMIS schemas defined in DOCSIS 2.0 provide better granularity that enhance the operators ability to manage the network. These enhancements include schemas to support network performance monitoring, capacity planning and security.

1.4.5 Physical Layer

Physical layer improvements include support for S-CDMA with active code selection and switchable upstream bands.

2.0 I-CMTS and M-CMTS Architectures

In this section we discuss the differences between the integrated CMTS (I-CMTS) and the modular CMTS (M-CMTS) architectures.

The two primary goals of the M-CMTS architecture are to provide a method of independently scaling the downstream PHY functionality by separating it from the core CMTS and to lower the cost of delivering video over DOCSIS at a price competitive with traditional MPEG VOD.

2.1 I-CMTS Architecture

An I-CMTS is defined as a single chassis with integrated MAC and PHY functionality supporting all upper layer protocols. In this architecture, a line card supports both downstream and upstream channels combining them into a MAC domain. Traditionally, these line cards support a 1x4, 2x8 or 5x20 configuration.

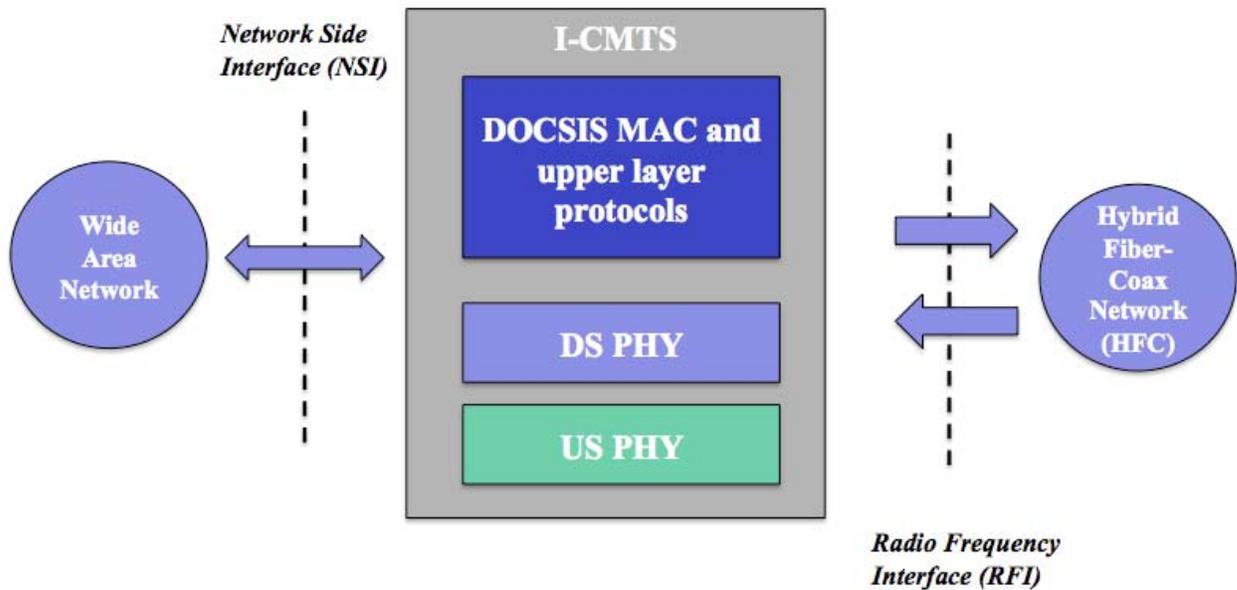


Figure 1 - I-CMTS Architecture

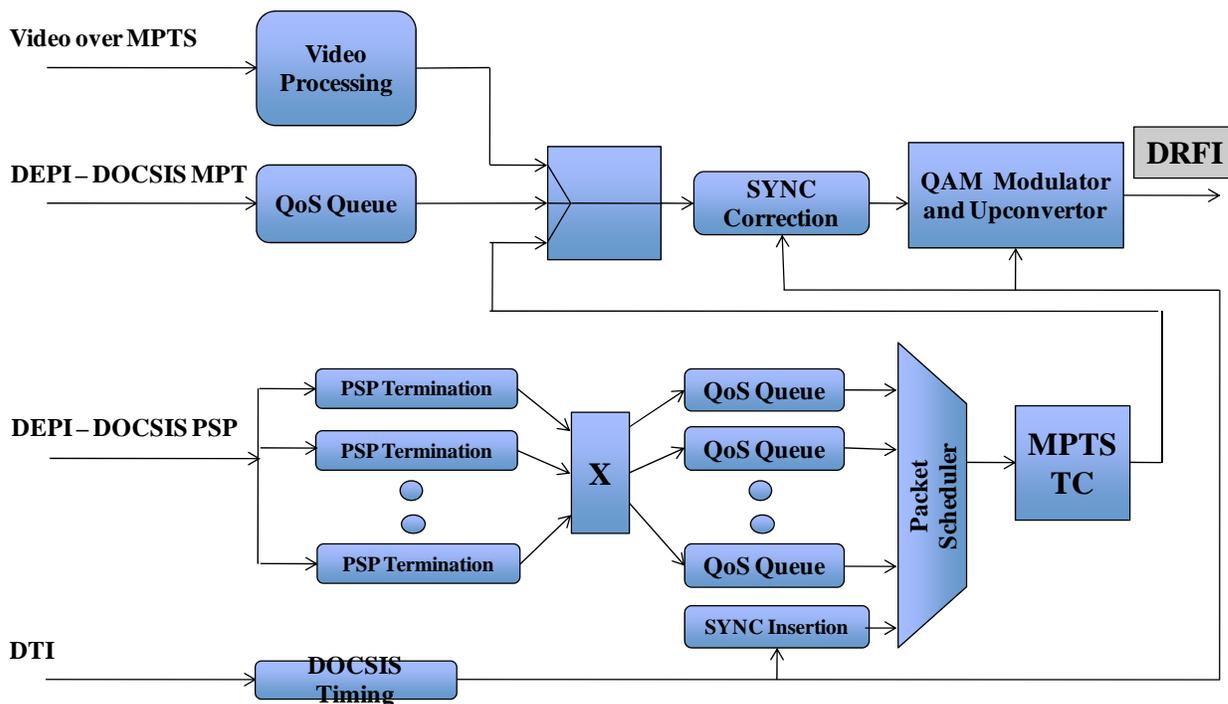
2.2 EdgeQAM Architecture

Before digging into the modular CMTS architecture, let's spend some time understanding the EdgeQAM itself.

EdgeQAM are edge devices that are capable of supporting a variety of applications including broadcast, video-on-demand (VOD), switched digital video (SDV), and M-CMTS. QAM densities

are increasing at a rapid pace with devices today ranging from 48 to 480 QAM's served per device, with futures looking to reach even higher.

The universal EdgeQAM supporting both video and data services helps reduce operational cost and enables evolution of services without stranding network capacity. Many have increasingly highly available services with redundant supervisor, RF switching, power supplies, and network components.



Courtesy of Greg White at CableLabs

Figure 2 - EdgeQAM Block Diagram

2.3 M-CMTS Architecture

The M-CMTS architecture separates the MAC and PHY functionality using standard interfaces between components. The M-CMTS core controls the MAC and upper-layer protocols. The EdgeQAM supports the downstream PHY using the DOCSIS External PHY Interface (DEPI) to communicate with the M-CMTS core. All devices must support the DOCSIS Timing Interface (DTI) specification to synchronize to a common timing source.

In the current release of the M-CMTS specification, neither the upstream PHY component nor protocol for communication between the M-CMTS core and external upstream PHY device is defined, instead the upstream MAC and PHY components are included in the M-CMTS core.

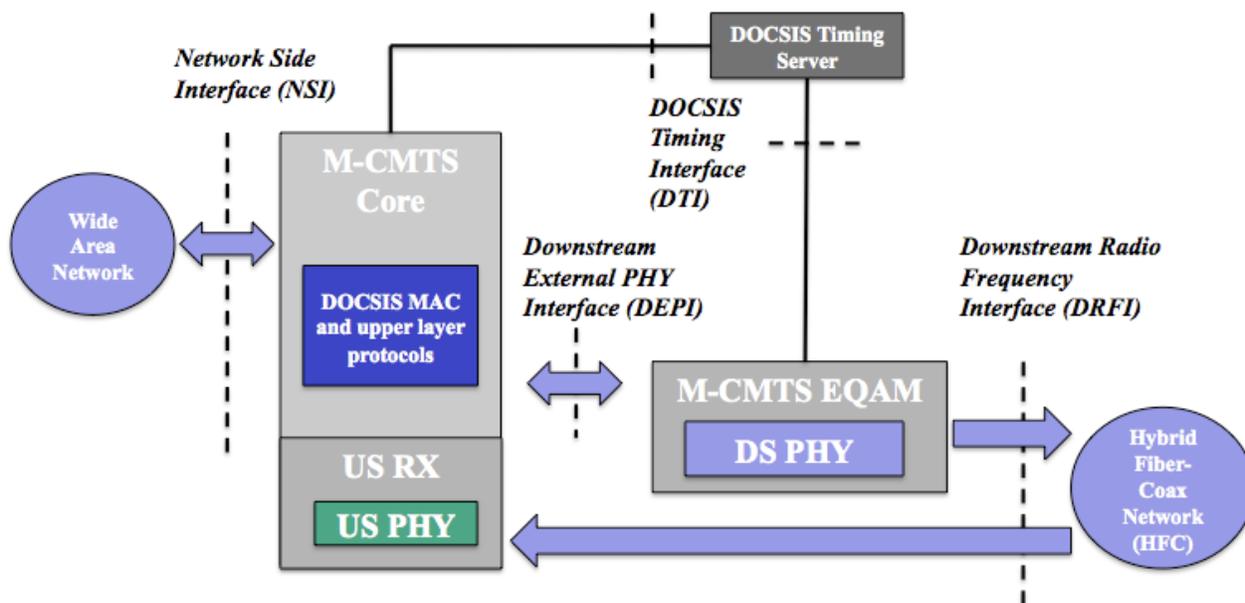


Figure 3 - M-CMTS Architecture

There are some mixed configurations as well. An I-CMTS may choose to add additional downstreams to an existing installed CMTS by adding DS MAC linecards that connect to external EdgeQAMs. These configurations thus contain both I-CMTS and M-CMTS downstreams.

2.3.1 Downstream External PHY Interface (DEPI)

This interface provides the data plane for all DOCSIS frames from the M-CMTS to the EdgeQAM device. Key features of this specification include L2TPv3 tunneling protocol as a pseudowire interface, MPT (MPEG Transport) and PSP (Packet Streaming Protocol) as pseudowire types, flow control, SYNC insertion, latency measurement and control connections/sessions/flows.

The difference between MPT and PSP is in how the DOCSIS frames are prepared for transmission to the EdgeQAM. In MPT mode the DOCSIS frames are generated and encapsulated in MPEG frames and in PSP mode raw DOCSIS frames are transmitted to the EdgeQAM. All EdgeQAM devices MUST support MPT mode and MAY support PSP mode.

Delay buffers are used to avoid large amounts of jitter in a multi-hop environment by requiring the EdgeQAM to maintain 20ms of buffer and insert MPEG nulls if the buffer runs dry. The M-CMTS core uses an open-loop flow control to try and avoid this situation.

SYNC insertion occurs in the M-CMTS to maintain the precise timing defined by the specifications. A 500ns max peak-to-peak jitter is required to meet the requirements.

The M-CMTS core performs latency measurements by using a time-stamped message echo and response. The M-CMTS core generates the message and inserts its timestamp, sends the message to the EdgeQAM, which then inserts its timestamp and returns the message to the core.

Control connections, sessions and flows are used to provide a single path from the M-CMTS core to an EdgeQAM for control, each control supports multiple DEPI sessions to that EdgeQAM and the flows carry the actual data from M-CMTS core to EdgeQAM.

2.3.2 DOCSIS Timing Interface (DTI)

The DOCSIS Timing Interface provides the mechanism to synchronize timing between the M-CMTS core and EdgeQAM.

2.4 CMTS Line Card Architectures

An I-CMTS may place upstream (US) and downstream (DS) components either on the same line card or on different line cards. When using separate line cards, the DS line cards contains the DS MAC and DS PHY and a separate US line card contains the US MAC and US PHY. A M-CMTS may place the US PHY, US MAC, and DS MAC on the same line card, or the M-CMTS may have a line card with only a DS MAC that gets combined with an I-CMTS line card.

There are many combinations, and they are usually all valid for one reason or another. The choice of approaches depends upon factors such as:

- **Investment Protection:** If existing I-CMTS line cards contain DS and US on the same line card, the least expensive way to increase DS capacity is to add additional DS line cards. These DS line cards could be I-CMTS with the PHY internal or M-CMTS with the PHY external.
- **Backplane Bandwidth Utilization:** Backplanes contain bi-directional bandwidth. Uni-directional line cards will only ever at the most use half of that bandwidth. Thus, the highest chassis density, that is when the backplane bandwidth 100% used in both directions, can only be achieved with bi-directional line cards. For example, if a backplane could support 50 DS and 50 US in each direction, a line card with 50 DS and 50 US on one line card would produce a denser chassis configuration than separate line cards with 50 DS on one line card and 50 US on the other line card. This is a real concern as downstream channel densities increase.
- **Configurability:** Separate DS and US line cards may offer more flexibility at making different combinations of DS and US ratios work. In practice, there are only a few ratios really needed. The ultimate measure of a system is cost and density of the overall system.
- **Redundancy:** Line card architectures must support RF Switch connectivity for I-CMTS and Ethernet switch connectivity for M-CMTS. Line cards with DS and US on one line card only require one extra line card. Separate DS and US line cards require two extra line cards. This also requires two RF Switches. Overall, this impacts the system cost.

- Head End requirements: Each head end is a different size and will have different requirements. Often, a very basic configuration will meet most requirements.
- Overall system cost and density: This is the true measure of a CMTS. There are many factors that determine this in addition to the ones mentioned here.

The choice between I-CMTS and M-CMTS should not be based upon philosophy or that one is right and the other is wrong. They are both valid and useful architectures. The choice of architecture should be the one that gives the best cost and density for your particular installation. Many times, it may actually end up being a blend of both I-CMTS and M-CMTS architectures in the same chassis.

3.0 Frequency-Space Diagrams

One of the challenges in visualizing the placement of channels in a multi-dimensional world is how the DOCSIS channels are related to physical entities. The authors have attempted to develop a simple method for showing these relationships, called the “Frequency-Space Diagram” (FSD).

In an FSD, the horizontal axis starts from low frequencies on the left and moves towards higher frequencies on the right. The vertical axis shows the spatial relationship between different fiber nodes. (Some alternative representations may have their axes reversed.)

Upstream and downstream channels are represented by building blocks that may span one or more fiber nodes on the vertical axis. Each building block has a designator. Typically, U_x is used for upstream channels, D_x for I-CMTS downstream channels, and M_x for M-CMTS downstream channels. Primary downstream channels are indicated with a bold outline while secondary downstream channels contain a thin outline. Alternatively, the designators “P” and “S” maybe included. A primary downstream contains the DOCSIS SYNC, MDD and MAP messages, and supports all pre-DOCSIS 3.0 CMs. A secondary downstream does not contain MDD, SYNC or MAP messages, and only supports DOCSIS 3.0 CMs.

Logical associations such as bonding groups (denoted with BG_x) and MAC Domains (denoted with MD_x) can be illustrated by drawing larger building blocks around the channel building blocks. A physical association of RF channels to line cards is done by including the name of the line card in an oval adjacent (above or below) to the channel blocks. If a building block needs to represent more than one RF channel, a designator of “x ch” can be included rather than drawing x building blocks.

The logical grouping of fiber nodes (denoted with FN_x) to one another is called a service group (denoted with SG_x) and is represented by dotted lines surrounding fiber nodes. A classic Service Group is determined by its downstream wiring in that the forward carrier path is common to the entire Service Group. Service Groups can be frequency dependant. For example, SDV (switched digital video) may group four fiber nodes to a service group at one set of frequencies while DOCSIS may group two fiber nodes to a different service group at another set of frequencies.

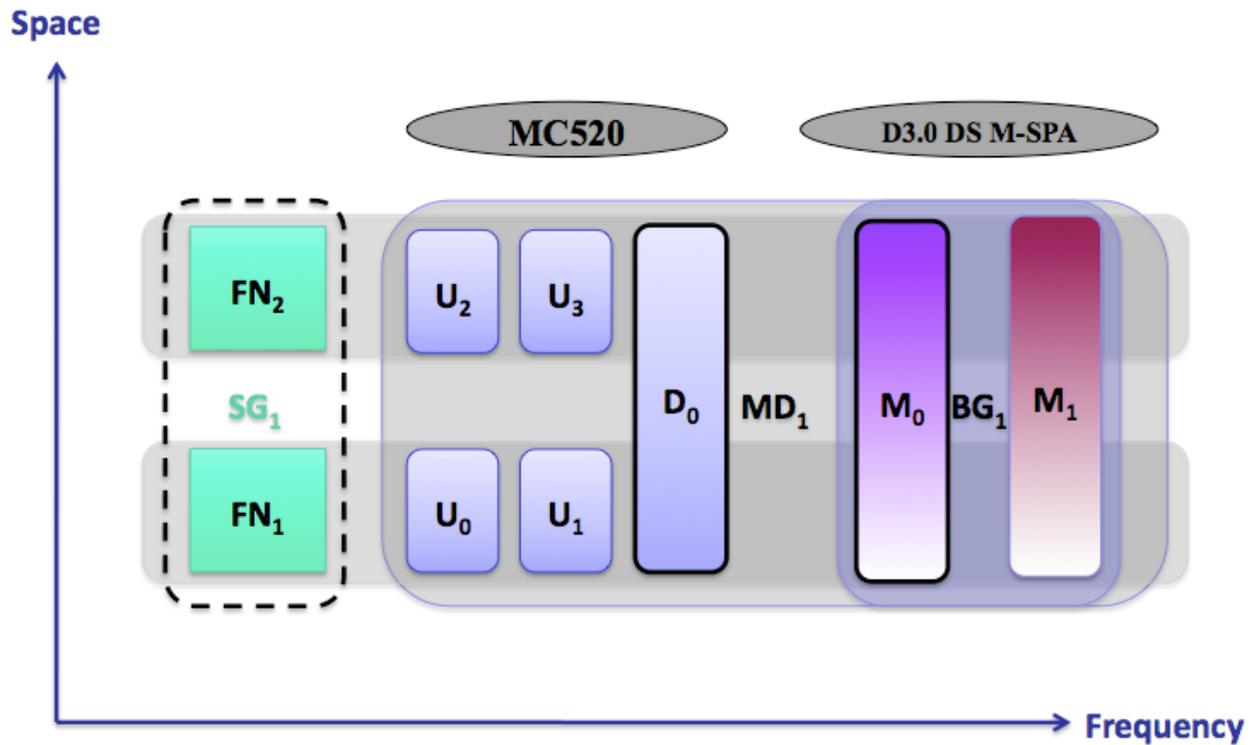


Figure 4 – Frequency-Space Diagram Example

We will use Frequency-Space Diagrams to show the relationship of channels to fiber nodes in the following examples.

4.0 Adding Downstream Capacity with I-CMTS

In the legacy MAC domain configuration, adding additional downstream capacity requires new upstream burst receivers as they are bound to the downstream transmitter forming the MAC domain. If the need is primarily for more downstream bandwidth, there is a potential to strand upstream capital as the receivers may not need to be used.

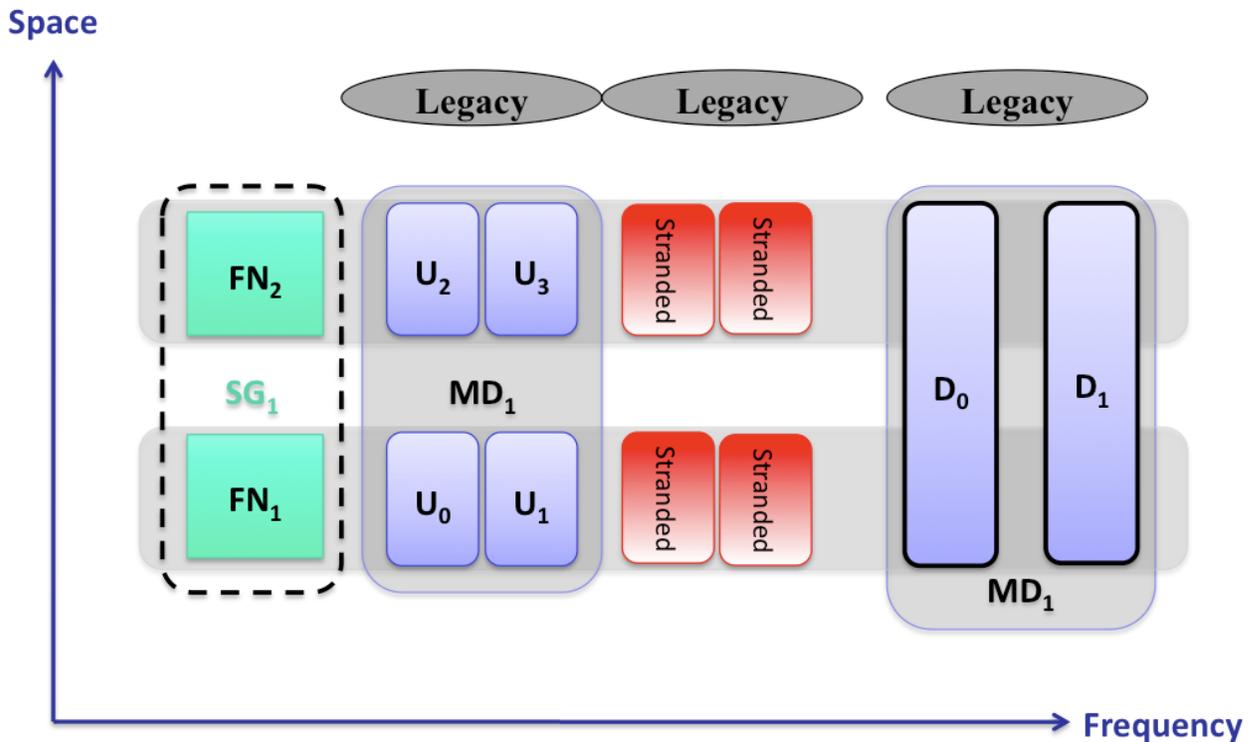


Figure 5 – FSD Example of Stranded Upstreams

As a result, there is a significant financial motivation in moving towards a modular CMTS architecture to lower the cost of downstream bandwidth. The flexibility of the M-CMTS design creates the ability to add new downstream capacity at a lower cost, as there is minimal stranded capital.

However, there are situations where additional legacy (pre-DOCSIS 3.0) capacity is needed. A primary use case example is the desire to keep voice calls on the integrated cards to maintain full redundancy on all components. Most CMTS vendors have years of experience providing a fully redundant carrier class product that has set a high standard for reliability. EdgeQAM vendors have made significant progress towards similar capabilities, but existing EdgeQAM equipment may not have the desired features for reliability. However this gap is being closed with each new release of EdgeQAM hardware. As CMTS vendors have implemented a mechanism for DQoS (PacketCable Dynamic Quality of Service [DOCS11]) flows to remain on the core CMTS, there is little effort required to maintain voice on the integrated cards.

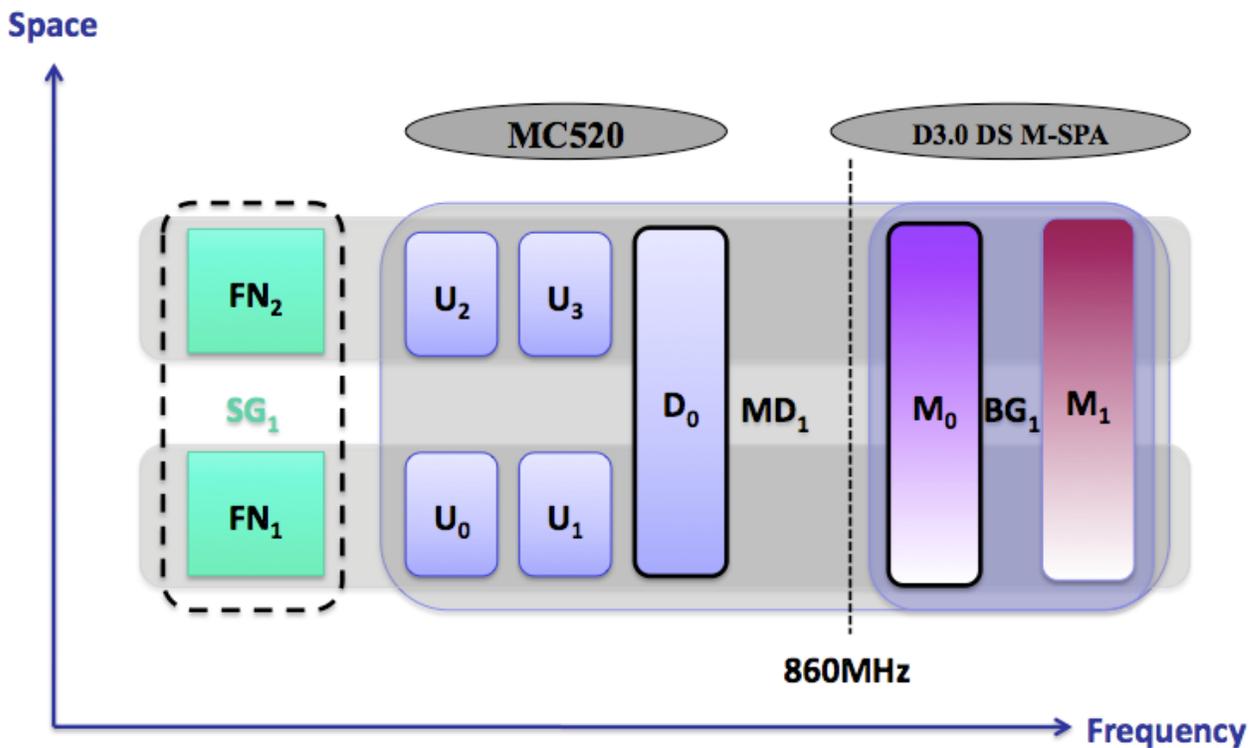


Figure 6 – FSD Example of Channels above 860 MHz

5.0 Adding Downstream Capacity with M-CMTS

Frequency placement of DOCSIS 3.0 capable downstreams is a key decision point when adding new channels. The typical EdgeQAM today in North America provides four 6 MHz wide downstreams per port with future generations capable of producing eight or more.

There are two limitations to the frequency agility of individual carriers within a port. The first limitation is based upon the capability of the EdgeQAM. The EdgeQAM builds a carrier group, then performs a block up-conversion. Typically, the carriers can be placed anywhere within a limited frequency range, and that limited frequency range can be located anywhere in frequency spectrum.

The receiving devices have additional constraints on the spacing of the received channels as well. Cable modem implementations use one or more tuners with one or more channels per tuner. Each tuner can be independently located within the 92 MHz to 1000 MHz frequency range, but has a limited maximum bandwidth typically of 60 MHz. A very common configuration is an eight downstream channel CM with two tuners and four channels per tuner.

Primary DOCSIS channels provide the required MAC Domain Descriptors (MDD) TLVs used by the cable modem at initialization. The primary channel must be located within the tuning range of the cable modems receivers, which in the case of a single tuner cable modem must be an approximately 60 MHz band. Many operators have placed the legacy downstream channels in the

lower portion of the forward path. Given the need to locate the DOCSIS 3.0 frequencies within the tuning range of a DOCSIS 3.0 CM, when placed above 860 MHz a least one primary channel has to be included in that block.

5.1 Frequency-Space Diagram Examples

The first example shows adding a second primary capable channel (M_0) for legacy devices using a modular downstream. The new downstream channel is added to the MAC domain (MD_1) and service group (SG_1) from the EdgeQAM, which communicates with the CMTS core using DEPI.

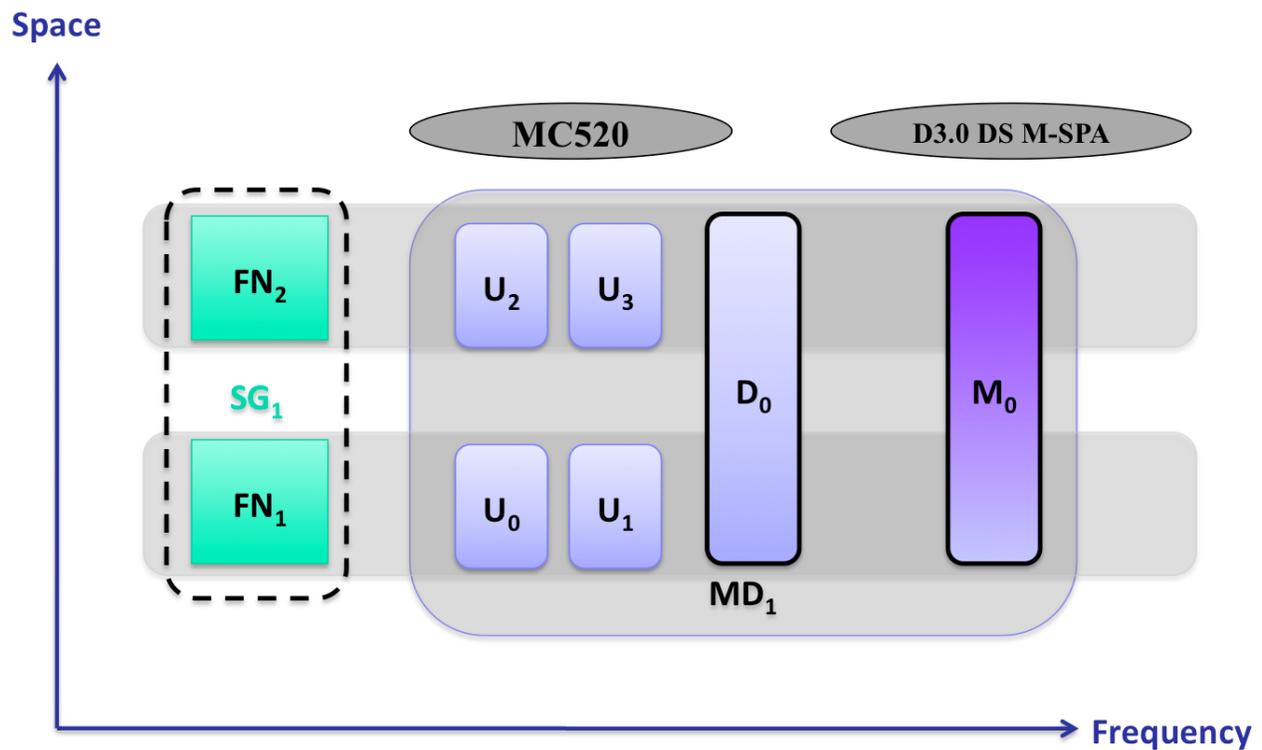


Figure 7 – Adding a Primary Capable Modular Downstream

The second example shows adding a second primary capable channel (M_0) and a secondary channel (M_1), both of which are bound in a bonding group (BG_1). The new primary capable downstream channel is used not only to provide additional legacy capacity, but also functions as the primary channel for the bonding group. The secondary only channel strictly provides capacity to the DOCSIS 3.0 capable cable modems and cannot be used by the legacy devices. Both are added to MAC domain (MD_1) and associated with the service group (SG_1).

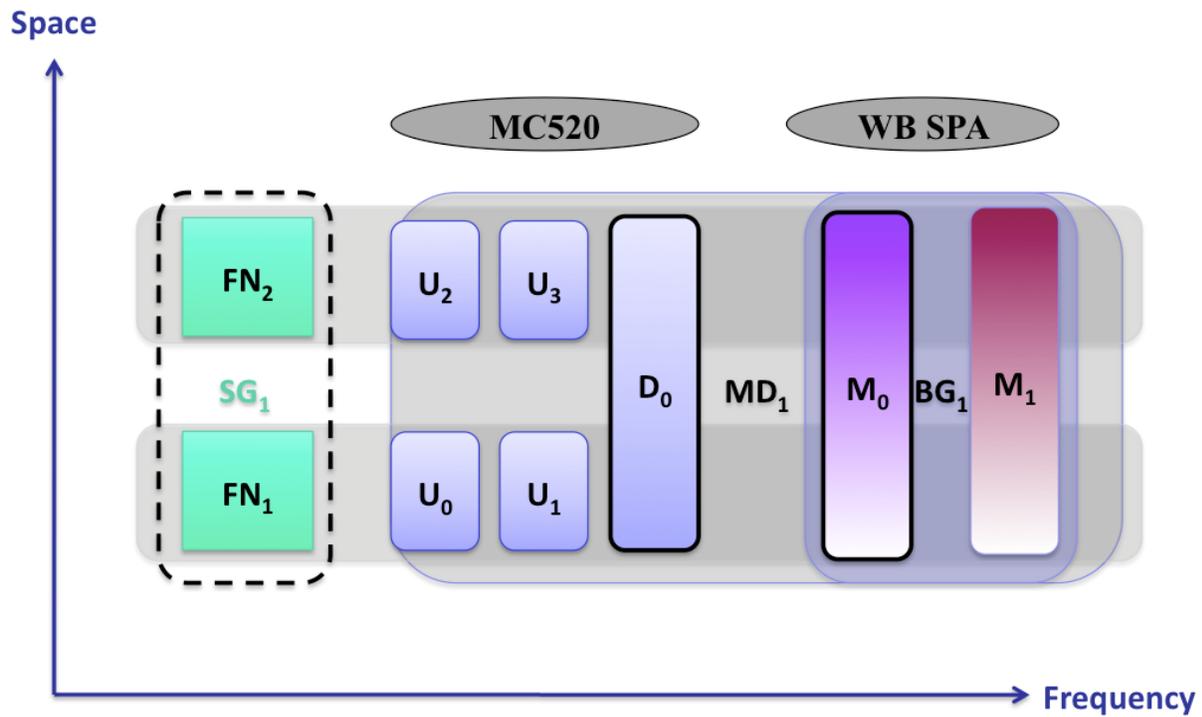


Figure 8 – Adding a Primary and Secondary Channel

The third example shows adding a second primary capable channel (M_0) and a secondary channel (M_1), both of which are bound in a bonding group (BG_1), then striped across 2 service groups (SG_{1-2}) and MAC domains (MD_{1-2}).

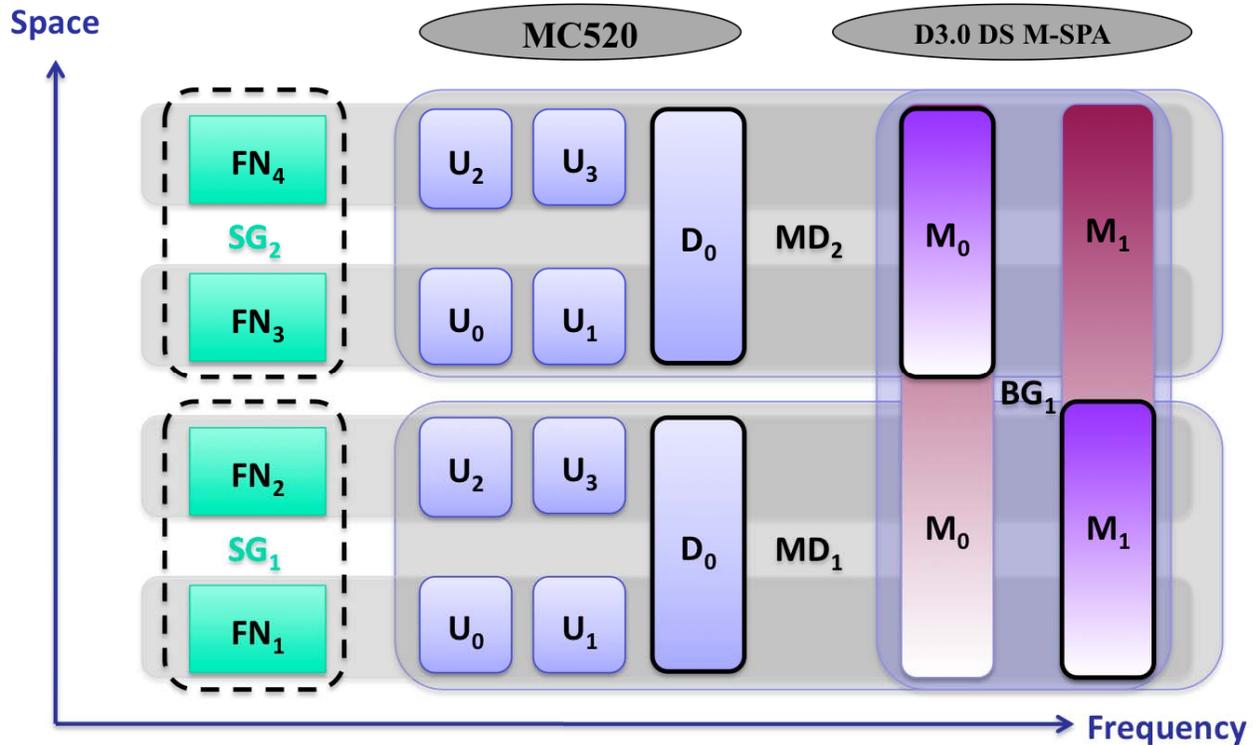


Figure 9 – Adding a Primary and Secondary Channel Striped across 2 MAC Domains

The next example shows adding a second primary capable channel (M_0) and a secondary channel (M_1), both of which are bound in a bonding group (BG_1), then added into MAC domain (MD_1) and associated with the service group (SG_1).

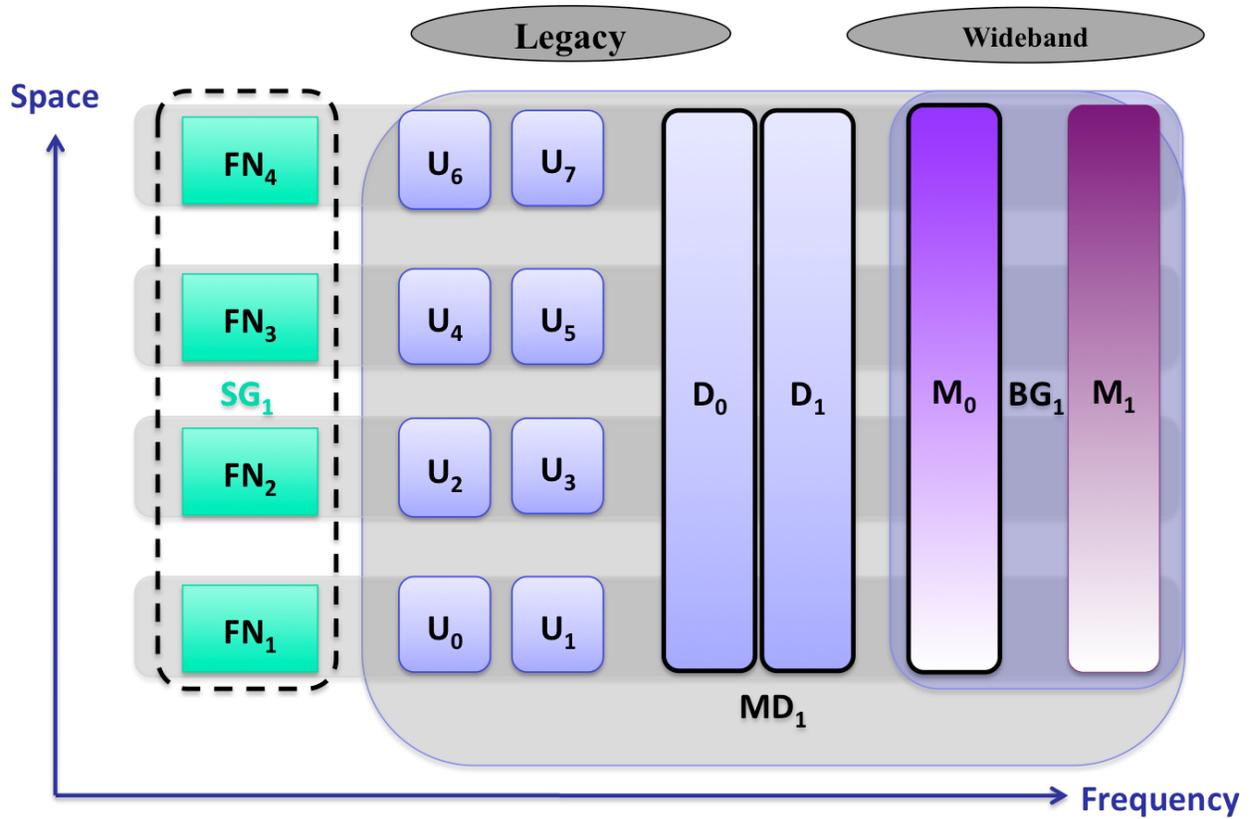


Figure 10 – Adding a Primary and Secondary Channel to a 2x8 MAC Domain

The following example shows converting a 4x2 MAC domain (4 downstreams, 2 upstreams) into an 8x2 MAC domain (8 downstreams, 2 upstreams) using an integrated modular downstream with 4 channels (D_{0-3}) in a bonding group (BG_1), adding EdgeQAM second primary capable channel (M_0) and multiple secondary channels (M_{1-3}) which are then bound in a bonding group (BG_2), then both bonding groups are added into MAC domain (MD_1) and associated with the service group (SG_1).

Note that the EdgeQAM downstreams are placed above 860MHz.

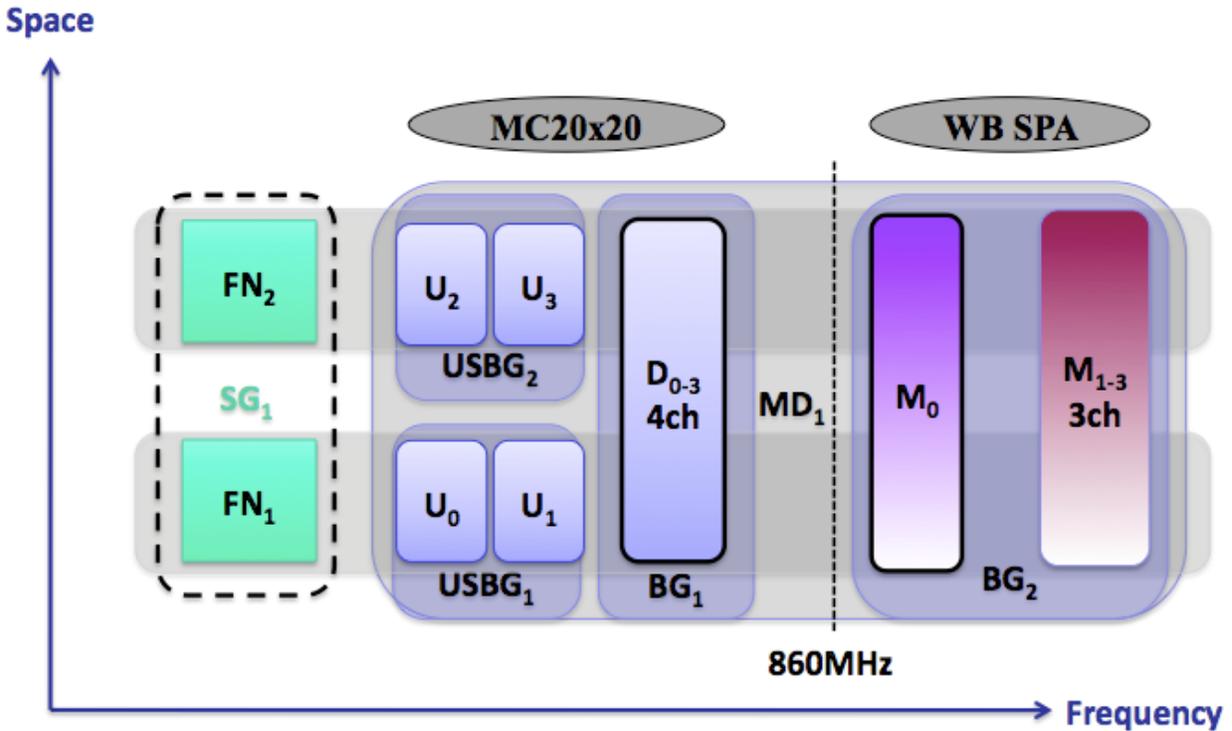


Figure 11 – Converting a 4x2 to an 8x2 MAC Domain

The final example shows an entire 5x20 line card (5 downstreams D_{0-4} and 20 upstreams U_{0-19}) providing and one primary channel downstream and four upstreams to each of five service groups comprised of two fiber nodes each. There are five wideband channels (M_0-M_4) configured as primary channel to a single service group and secondary to others, with an additional three wideband channels configured as secondary channels. All eight of the channels are striped across the five MAC domains. This advantage of this configuration is that it uses the minimum number of additional CMTS resources (one-third of a SPA card) to achieve maximum coverage of 8 channels of bonding. Three SPA cards can do an entire eight slot chassis and provide 8 channels of bonding to all service groups.

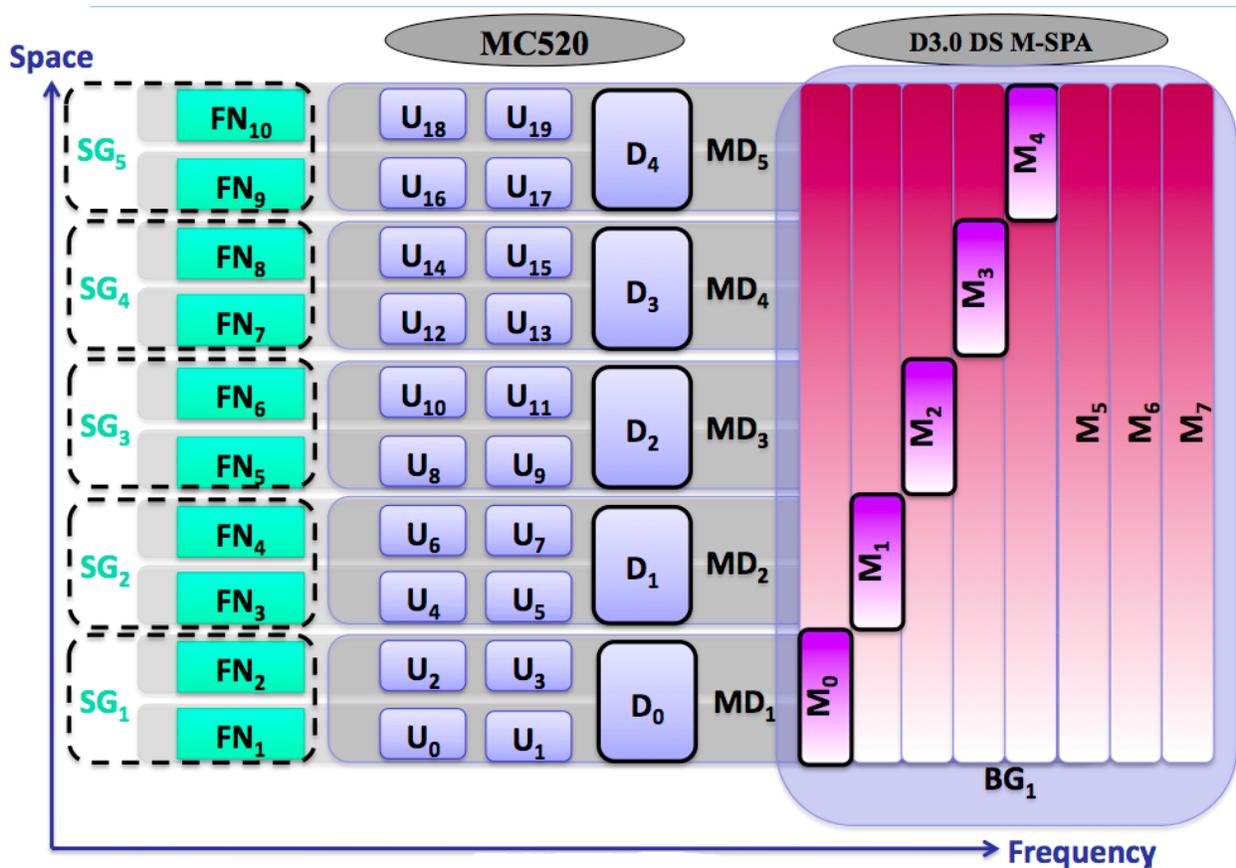


Figure 12 – Full 5x20 Card with Striped Wideband Channels

6.0 M-CMTS Futures

There are many ongoing discussions concerning what changes CMTS technology require to future-proof the access network. In this section we will explore the drivers that may impact what the next generation architecture may become. The use of the term “CMTS” is loosely referring to the edge devices that will provide data, voice and video services in some form, which should not be construed to mean the specific device, as we know them today.

6.1 IP Video

Using bonded channels to provide IP video results in improved performance based on the overall peak-to-average ratio being lower when more streams are combined in a wideband channel. This self-averaging of peaks allows more videos to be packed into less bandwidth using variable bit rate (VBR) encoding than can be seen with committed bit rate (CBR) streams. Additionally, by using wider channels, more VBR streams can be fit into these wideband channels than if they were placed into the separate channels that comprise the bonding group.

For example, if 12 total VBR streams need to be provided to a particular service group, and each 6 MHz channel could support three HD VBR streams, four channels would be required to send those streams downstream. But if a four-bonded channel group was created, there is a potential to send up to 16 streams to the same group. That is a gain of 33% using self-averaging to achieve a better peak-to-average ratio.

For that same four-bonded channel group, if some of the streams have packets delayed (no more than the jitter buffer of the receiving device), there is a potential for another 33% gain by fitting up to 20 VBR streams into the same wideband channel.

There is ongoing research into these gains, but initial results look promising.

6.2 Next Generation Access Architecture

A newer topic for future next generation access architectures currently being discussed centers around a complete separation of the upstream and downstream MAC-PHY from the CMTS core, along with integration of video processing for MPTS video streams.

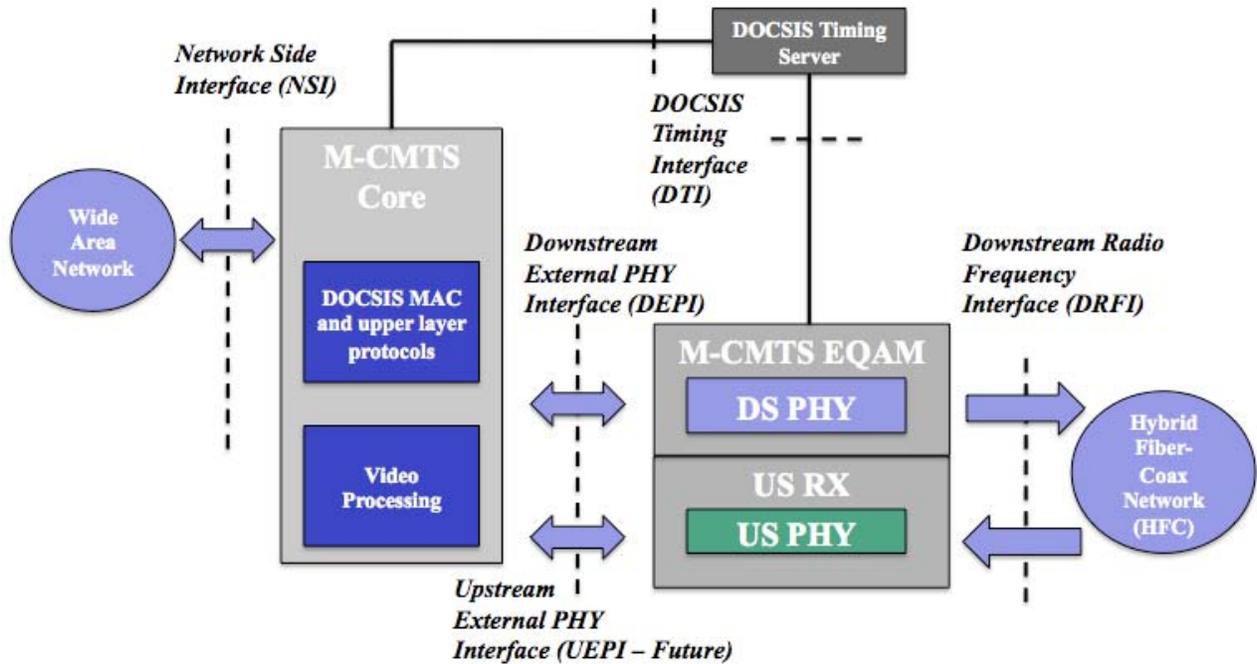


Figure 13 – Next Generation Access Architecture Example

7.0 Conclusion

The decision between modular and integrated CMTS is a difficult choice. The modular design simplifies addition of downstream capacity, but newer integrated cards provide similar functionality, which complicates the decision. The authors suggest that there does not have to be a decision of modular versus integrated, but that instead a mix of both technologies is stronger than either one can stand by itself.

By mixing integrated and modular we can achieve a higher level of investment protection than by replacing existing technologies. Both are capable of providing similar functionality, it is a matter of selecting the product packaging option that best fits your needs. By choosing a mix of technology the creativity potential is more than meets the eye.

It does not have to be one or the other; sometimes the best choice is a little bit of both.

8.0 Acronyms and Abbreviations

CBR	Constant Bit Rate
CM	Cable Modem
CMTS	Cable Modem Termination System
DEPI	Downstream External PHY Interface
DOCSIS	Data Over Cable System Interface Specification
DTI	DOCSIS Timing Interface
I-CMTS	Integrated CMTS
M-CMTS	Modular CMTS
SLA	Service Level Agreement
UEPI	Upstream External PHY Interface
VBR	Variable Bit Rate

9.0 References

[DOCS10] “Data-Over-Cable Service Interface Specifications, DOCSIS 1.0, Radio Frequency Interface Specification”, SP-RFI-CO1-011119, November 19 2001, CableLabs.

[DOCS11] “Data-Over-Cable Service Interface Specifications, Radio Frequency Interface Specification”, SP-RFIV1.1-IO8-020301, March 1 2002, CableLabs.

[DOCS20] “Data-Over-Cable Service Interface Specifications DOCSIS 2.0 Radio Frequency Interface Specification”, SP-RFIV2.0-IO4-030730, July 30 2003, CableLabs.

[DOCS30PHY] “Data-Over-Cable Service Interface Specifications DOCSIS 3.0 Physical Layer Specification”, CM-SP-PHYv3.0-IO1-060804, August 4 2006, CableLabs.

[DOCS30MULPI] “Data-Over-Cable Service Interface Specifications DOCSIS 3.0 MAC and Upper Layer Protocols Interface Specification”, CM-SP-MULPIv3.0-IO1-060804, August 4 2006, CableLabs.

[ERMI] “Data-Over-Cable-Service-Interface Specifications Modular-CMTS Edge Resource Manager Interface Specification”, CM-SP-ERMI-IO2-051209, May 12 2009, CableLabs.

[SNMP] R. Presuhn, B. Wijnen “An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks”, RFC-3411, December 2002, Internet Engineering Task Force.