CONVERGED MULTISERVICE ACCESS PLATFORM

Jorge D. Salinger
VP, Access Architecture
Comcast Cable

Abstract

A new headend equipment architecture option for implementation of the traditional CMTS and Edge QAM functions is presented. The equipment resulting from implementing this new architecture, called Converged Multiservice Access Platform (CMAP), incorporates all the CMTS and Edge QAM functions. Each CMAP downstream RF port implements all QAMs for digital narrowcast and broadcast services for a single service group. Similarly, CMAP upstream RF line cards implement multiple demodulators per port. This architecture, which can be implemented in a single, integrated chassis, or by separating the packet processing from the PHY and MAC in separate modules, enables unprecedented density in MSOs' headend facilities.

Starting by outlining the key goals and objectives of the architecture, this paper describes the various CMAP components, their key features, the density achieved by the architecture, multiple operational simplicity and efficiency improvements, and the transport agnosticism achieved.

A description of the specifications spelling out the full details of the CMAP product requirements and the timeline for their development is provided. The relationship between the CMAP product specifications and the various CableLabs® specifications, such as DOCSIS®, M-CMTS™, DRFI, MHA, PacketCable™, etc., is also explained.

Examples are presented to show how CMAP could be deployed in typical cable systems, including its deployment in MSO networks of varying sizes, capacities and composition of services. Space and power savings, the key benefits of CMAP, are depicted in comparative analysis.

<u>NOTE:</u> All examples presented in this paper are only for illustrative purposes, and do not reflect the actual deployment plans of Comcast or any other cable operator.

BACKGROUND AND RATIONALE

For a few years now, MSOs have been increasing the number of QAM channels used for narrowcast services. Most MSOs are deploying more and more QAM channels to support growth from the success of Video on Demand, especially as a result of the availability of more High-Definition TV (HDTV) content. Additionally, the use of Switched Digital Video (SDV) for an increasing number of multicast content offerings is driving the deployment of QAM channels even further. And, with the availability of channel bonding in DOCSIS 3.0, MSOs are deploying additional QAM channels for their CMTS equipment to support the newer, higher bandwidth data services.

At the same time, MSOs continue to reduce the size of service groups to make more efficient use of their networks. The drivers, for many years now, have been operational streamlining (smaller service groups result in improved service quality) and efficient use of spectrum (support narrowcast service growth by reusing spectrum).

These two trends result in a continuous increase in the number of QAM channels per service group. Moreover, the expectation from the current service projections is that such growth will continue and even expand, especially as MSOs reduce the number of analog channels available in the network.

However, the deployment of additional QAM channels in Edge QAM or CMTS equipment cannot easily be supported in the space available within existing typical headend and/or hub/OTN sites.

As a result, the cable industry needs ever denser QAM-channel-per-RF-port Edge QAMs to reduce both the resulting environmental requirements and the capital and operational costs of the equipment itself.

TECHNOLOGY EVOLUTION

Interestingly enough, cable industry suppliers identified the above trends quite some time ago. Because of such foresight, Edge QAM vendors have been developing denser QAM-channel-per-RF-port implementations for several years now, even approaching densities that allow implementations of unique QAM channels for every channel in every RF port.

However, this technology evolution has been difficult to incorporate in equipment available for purchase. It is not a simple operational and financial matter for MSOs to take the leap towards such higher densities for any one service, and consequently it is difficult for vendors to justify the investment to implement this technology. This is not only true for Edge QAM vendors, but particularly difficult for CMTS suppliers implementing integrated architectures.

With Modular CMTS and the Modular Headend Architecture, as defined by CableLabs, it should be possible to achieve such densities as Edge QAM development evolves towards higher densities and CMTS equipment is developed for these network architectures. However, most CMTS development has focused on an integrated architecture for a variety of technical reasons.

ENTER CMAP

A new equipment architecture option is emerging that enables the implementation of denser network architectures in yet another way, providing both MSOs and vendors an alternative approach for achieving the original goals of the Modular Headend Architecture – denser QAM-per-RF port implementations.

Such equipment architecture is described in work underway at Comcast, which is developing product specifications for a new class of equipment called Converged Multiservice Access Platform, or CMAP.

CMAP leverages existing technologies such as DOCSIS 3.0 and current HFC architectures, incorporates newer ones such as dense Edge QAM architectures and Ethernet optics (EPON, in particular). It also leverages the experience acquired over the many decades of technology evolution for cable networks.

The key goals of CMAP include:

- Enabling implementation of denser headend equipment targeting a much higher number of narrowcast services, reducing costs and environmental requirements in headend/hubs/OTNs.

- Developing an access technology-agnostic architecture, making it possible to deploy newer access technologies with the same services architecture.

- Leveraging new and/or broadly deployed technologies to unleash further capacity in the cable industry's HFC network, using overlay architectures to simplify deployment.

CMAP OBJECTIVES

The Converged Multiservice Access Platform is intended to provide a new equipment architecture approach for manufacturers to achieve the Edge QAM and CMTS densities that MSOs require to address the costs and environmental constraints resulting from the success of narrowcast services. In addition to the architecture described in the Modular Headend Architecture Technical Report from CableLabs (i.e., Modular CMTS with Universal Edge QAM), the CMAP provides an alternate approach to the implementation of headend equipment that delivers QAM channels for different services.

To achieve the functionality described above, a CMAP device implements the various Edge QAM and CMTS functions in a consolidated platform. The result, as shown in the figure below, is that a single CMAP downstream port will include all the QAM channels for all digital services. For example, a typical downstream RF port may be licensed to include 32 QAM channels for narrowcast and 96 QAM channels for broadcast services. If deployed in a 750 MHz system that maintains 30 analog channels, the CMAP RF port will provide 32 QAM channels for narrowcast video, data and voice services and approximately 50 additional QAM channels for broadcast services.

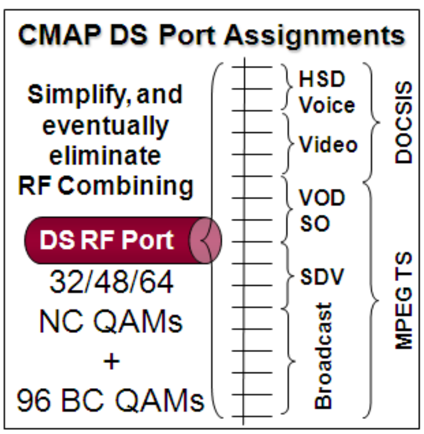

As with the existing CMTS architectures, a CMAP device can be implemented in an integrated or modular manner.

In the first case, all functions are implemented in a single chassis.

In the second, CMAP functions are divided between a Packet Shelf (PS) and an Access Shelf (AS), as follows:

- The PS implements the packet processing functions, such as subscriber management, service flow management, video program stream edge manipulation (e.g., multi-program transport stream creation, PCR restamping, etc.), layer-3 routing and higher layer protocol manipulation, and other such functions.
- The AS implements all the upstream and downstream PHY functions normally associated with the CMTS and the Edge QAM, and as much of the MAC as needed to support both upstream and downstream flows. A documented interface between the AS and the PS is defined to enable interoperability between AS and PS vendors.

The figure in the next page outlines one possible modular implementation of CMAP, where multiple types of Access Shelves are available for different access architectures. However, other implementations are also possible. For example, multiple access technologies could be incorporated into a

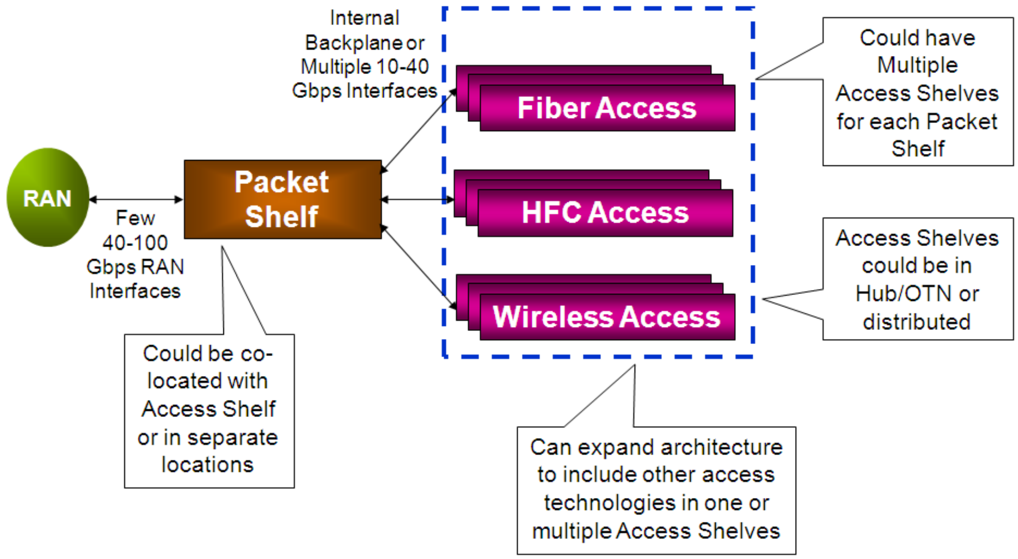

single AS, so that one centralized PS would interface with multiple AS devices possibly distributed across various sites.

The key functional goals for CMAP include:

- Flexible use of QAM channels for the various services offered by MSOs, enabling the configuration of the CMAP to provide a changing number of MPEG transport stream-based services (e.g., for VOD, SDV, etc.) versus DOCSIS-based services (e.g., HSD, voice, etc.).

- Individually configurable assignment of QAM channels to the various service groups, so that it would be possible to have service groups for HSD/voice, VOD, and/or SDV overlap in different ways without requiring the various service groups to provide homogeneous coverage.

- Efficient implementation of Edge QAM blocks by implementing separate sets of QAM channels for narrowcast and broadcast applications, such that QAM channels for narrowcast services are individually implemented for each RF port but QAM channels used for broadcast services are shared amongst all the RF ports in each downstream line card (DLC).

- Simplification of the RF combiner by providing all QAM channels for all digital services from a single RF port, leaving only certain legacy functions for the RF combiner network. The list

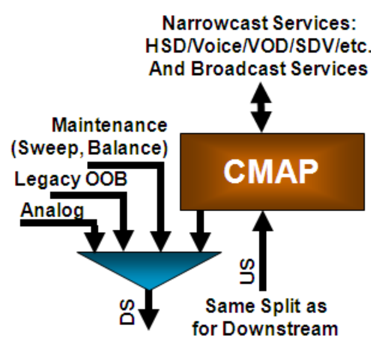

includes any remaining analog channels, the legacy out-of-band control channel, and any maintenance equipment not incorporated into the CMAP, as shown in the figure above.

- Implementation of sophisticated proprietary encryption systems (e.g., PowerKEY, DigiCipher, etc.) without requiring special-purpose hardware, so that a CMAP from any vendor can implement either encryption mechanism, or both mechanisms, with the same platform.

- A transport-agnostic network architecture, including implementation of PON and other access network technologies natively within the CMAP.

- Significant operational improvements, including significant environmental efficiencies (e.g., much less space, power consumption, and heat dissipation), implementation of functions such as upstream spectrum surveillance, continuous wave carriers for plant amplifier biasing, and many other operational enhancements.

SCOPE OF CMAP SPECIFICATIONS

The requirements included in the CMAP specifications currently under development at Comcast outline product requirements, including capacity, performance, network implementation functions, and other such targets and objectives. In doing so, the CMAP specifications reference industry standards, such as CableLabs specifications (e.g., DRFI, DOCSIS, VSI, PMI, PacketCable, etc.) and SCTE standards (SCTE-02, etc.), without duplicating the requirements detailed in those standards.

Please note that the CMAP specifications do not contradict or redefine any industry standards. Where necessary, changes are made in the industry standards, and not in the CMAP specifications. For example, certain changes are being made to the CableLabs DRFI specification and the SCTE-02 standard, which the CMAP specifications take into account or anticipate with appropriate descriptive language.

In some cases, the CMAP specifications detail requirements that exceed those of other industry standards. For example, an industry standard may indicate a preference with a SHOULD requirement while the CMAP specification might label it an absolute requirement with a MUST.

To develop the CMAP product specifications, Comcast is working with a broad group of industry leaders and technology experts from companies interested in the development of a CMAP, all of whom have volunteered to help Comcast develop these requirements. Staff members from CableLabs, Cable Europe Labs, and other advisers are assisting Comcast in this effort. Most importantly, a number of contributors from various North American and European MSOs are participating in the effort as well.

As detailed in the table above, the CMAP Team plans to complete three product requirements specifications in the next few months. The first of these product specifications, called the CMAP Hardware and Functions Specification, has been completed. The other two, Configuration and Management and PASI, are currently under development and should be completed by the middle of 2010. Additionally, following the completion of the product requirements specifications, the team plans to develop recommended test procedure specifications.

Specification	Objective
Hardware and Functions	Hardware components and requirements, and the various features and functions implemented by the CMAP.
Configuration and Management	Interfaces and requirements for configuring and managing the CMAP
Access Shelf-to-Packet Shelf Interface (PASI)	Functions performed by the PS vs. the AS and the characteristics of the interface between the two components.

In addition to the CMAP product requirements specifications, the team has contributed to additions and/or changes to existing or new industry specifications. The main body of work in that regard has been related to the CableLabs DRFI Specification, which has undergone several Engineering Change Requests (ECRs) to accommodate the

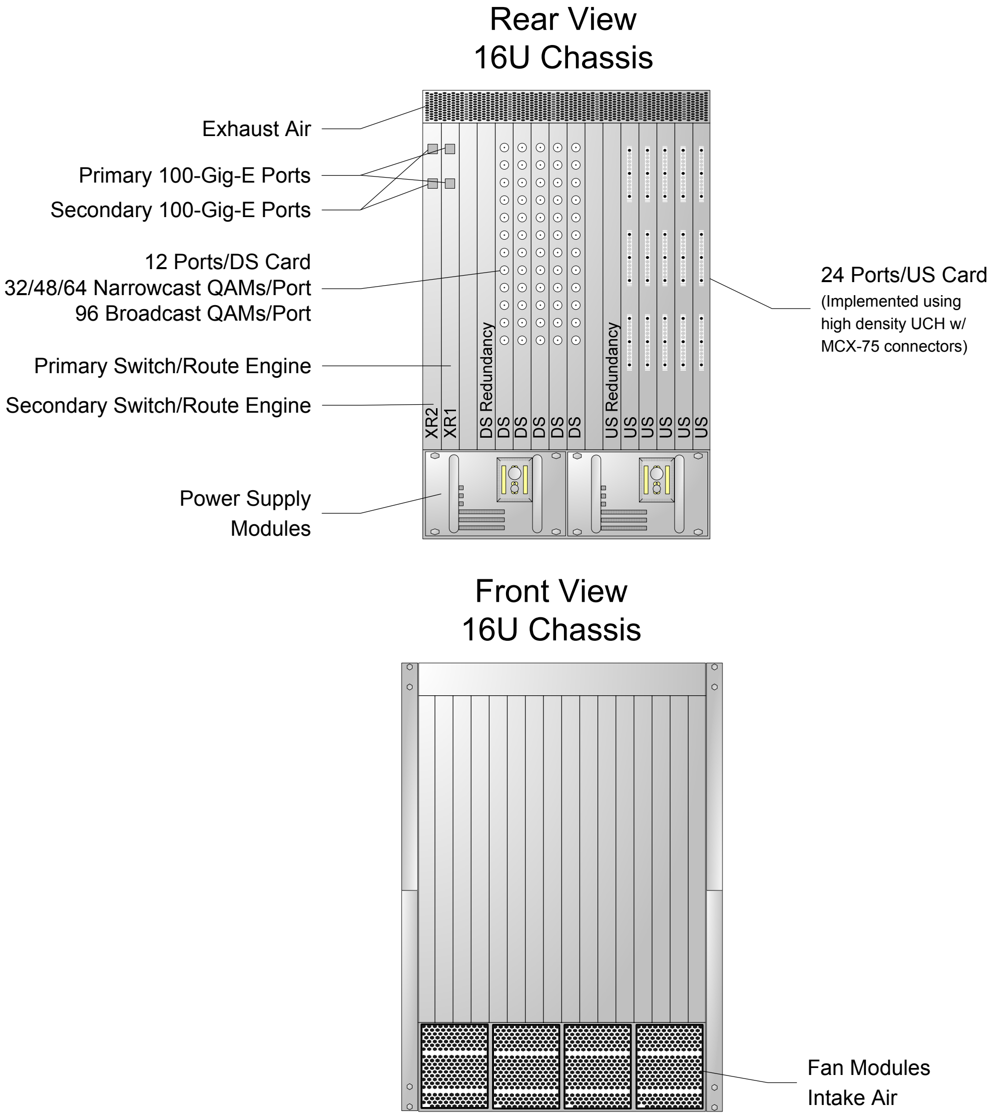

design and operation of implementations of large numbers of QAM channels-per-RF-port, which are applicable to the CMAP as well as other dense QAM channel-per-RF-port equipment implementations. Another area of industry specification work relates to SCTE-02 through the SCTE IPS Working Group, regarding enhancements to the "F" connector requirements and the addition of a 75 Ohm version of the MCX connector. The 75-Ohm MCX connector is commonly implemented in a gang holder known as UCH, which consists of a row of 10 connectors typically used with mini coaxial cable. Other industry specifications may be updated as deemed appropriate.

Unlike other similar efforts to date, such as RFIs, RFQs and/or prior requirements documents for CMTS and Edge QAM equipment, the CMAP specifications outline very specific chassis requirements. These requirements include comprehensive line card implementation details, such as the number of supported QAM channels for each function in terms of density and redundancy characteristics. They also spell out such physical interface objectives as type of connectors, detail preferences for power supply locations and orientations, set airflow direction and entry/exit requirements, and detail many other such implementation requirements. These requirements should limit vendor innovation because vendors will still be left with many opportunities for differentiation. At the same time, they recognize the need to simplify operations by creating standards for key operationally beneficial parameters.

The following figure shows a possible front and rear view of the CMAP chassis that would be compliant with the CMAP specifications. In the figure, the following details are depicted: rear-facing connectivity for all components; downstream line cards (DLCs); upstream line cards (ULCs) with twice as many upstream ports as downstream ports for 2:1 upstream-to-downstream ratio; redundant switch-route engines with primary and secondary 100 GigE ports; redundant power supplies; and vents for air flow. The diagram does not depict PON line cards for business services, which are not mandatory but are strongly preferred.

Given the scope of each RF port providing all services for a given service group, it is important that the operation be highly reliable. Therefore, the chassis is required to implement N+1 redundancy for upstream and downstream line cards and 1+1 redundancy of all common equipment. This line card redundancy is achieved by way of a mid-plane near-passive RF switch and the use of physical interface cards (PICs), which provide the separation between the active components with critical mean time between failure (MTBF) and the RF interfaces to the minimum remaining external combining and downstream/upstream lasers.

CAPACITY ESTIMATES

To help guide equipment and network design, the table included below depicts three deployment scenarios for a CMAP, as follows:

- Minimum, albeit unlikely, deployment in the left column,
- Maximum, also unlikely, deployment in the right column, and
- Estimated probable initial deployment in the middle column.

These scenarios show that a downstream NSI capacity of 15 Gbps is easily necessary, while an absolute maximum of 155 Gbps could possibly be required for the line cards envisioned in the specification. But a capacity of approximately 30 Gbps is most likely.

Similarly, for the upstream NSI direction, a capacity of about 7 Gbps might be required upon initial deployment, as depicted in the table.

Please note that the scenario considered probable in this example depicts 5 downstream line cards, with 32 active narrowcast QAM channels, for which 50% of the content is unique (e.g., 50 % of the content is replicated via multicast). It also includes 1 broadcast line-up for the entire chassis with 75 active QAM channels.

Approx Calculations	Minimum Chassis Capacity (Mbps)	Probable Capacity (Mbps) 5 DLC, 32 NC/RF (50% unique), and 1 BC/Chassis (75 QAMs)	Maximum Chassis Capacity (Mbps)
DS NC QAMs per RF Port	a1 = 22 QAMs * 36 Mbps	b1 = 32 QAMs * 36 Mbps	c1 = 64 QAMs * 36 Mbps
DS BC Lineups per LC	a2 = 40 QAMs * 36 Mbps	b2 = 75 QAMs * 36 Mbps	c2 = 96 QAMs * 36 Mbps
DS NC+BC Chassis	((a1 * 8 ports) * 50% unique) * 4 LCs + a2	((b1 * 8 ports) * 50% unique) * 5 LCs + (b2 * 2 LCs)	(c1 * 12 ports * 100%) * 5 LCs) + (c2 * 5 LCs)
US Port	a3 = (2 US * 26 Mbps) + (1 US * 8.8 Mbps)	b3 = (3 US * 26 Mbps) + (1 US * 8.8 Mbps)	c3 = (4 US * 26 Mbps) + (2 US * 8.8 Mbps)
US Chassis	a3 * 8 ports * 4 LCs	b3 * 16 ports * 5 LCs	c3 * 24 ports * 5 LCs
DS NC QAMs per RF Port	792	1,152	2,304
DS BC Lineups per LC	1,440	2,700	3,456
DS NC+BC Chassis	14,112	**28,440**	155,520
US Port	61	87	122
US Chassis	1946	**6,944**	14,592

For clarity, the top portion of the table shows the calculations and the bottom portion of the table shows the calculated results in Mbps.

QAM REPLICATION

In order to simplify integration of the CMAP into existing systems, the CMAP requires a QAM channel replication feature. With this feature, the CMAP can "copy" the contents of one QAM channel onto the same QAM channel in one or more other RF ports, thereby implementing electronically an analogous function as RF splitting an Edge QAM port and combining the splitter outputs into multiple service groups.

The figure below illustrates the QAM Replication feature. Note that each RF port has a unique HSD service group depicted with a circle.

The purpose of QAM Replication is to allow

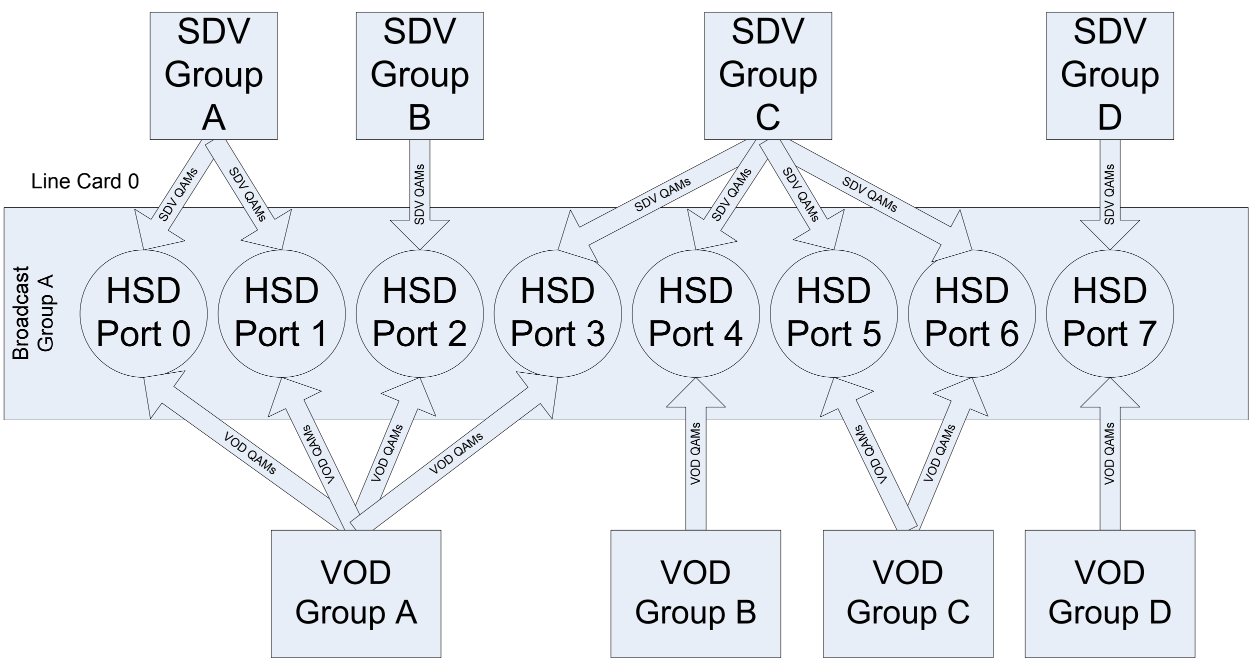

the creation of service groups on a decoupled (service-by-service) basis. With this feature, an MSO can replicate SDV and/or VOD QAM channels across multiple ports on a given line card.

The genesis for this feature can be found in the current deployment scenarios, where VOD service groups implemented by a separate set of Edge QAMs may span multiple HSD service groups. In other cases, a SDV service group may span a different number of HSD service groups, and perhaps more than one VOD service group.

Because HSD, VOD and SDV service groups are currently implemented by separate Edge QAM devices, today's combining is done at the RF level on an as-needed basis. But, as MSOs move to the CMAP, where each RF port has its own QAMs, it is not possible to combine service groups at the RF level any longer.

One way to deal with the new scenario is to align service groups, where a VOD service group would geographically coincide with an SDV service group and with an HSD service group. Another approach is to use this QAM Replication feature outlined herewith.
Neither approach offers a perfect solution. Service group alignment requires additional work, as well as extra equipment in many cases, and results in a loss of economies of scale such as those that benefit multicast for SDV. QAM channel replication, on the other hand, saves on the work and equipment required to align service groups and maintains the current deployed scenario, including its economies of scale. On the flip side, however, it does not save on QAM channel cost because the hardware for each individual QAM channel has to be in place for each port anyway.

ENCRYPTION CAPABILITY

Given the broad video services supported by the CMAP, it is imperative that the CMAP implement extensive encryption capabilities. Therefore, the CMAP Team invested a significant amount of time in developing a strategy for encryption within the CMAP so that a CMAP from any manufacturer, given the appropriate licensing, will support encryption to the fullest extent from any Conditional Access System (CAS).

To accomplish this, the CMAP implements a very clever scheme, previously envisioned for a cousin technology called Next Generation

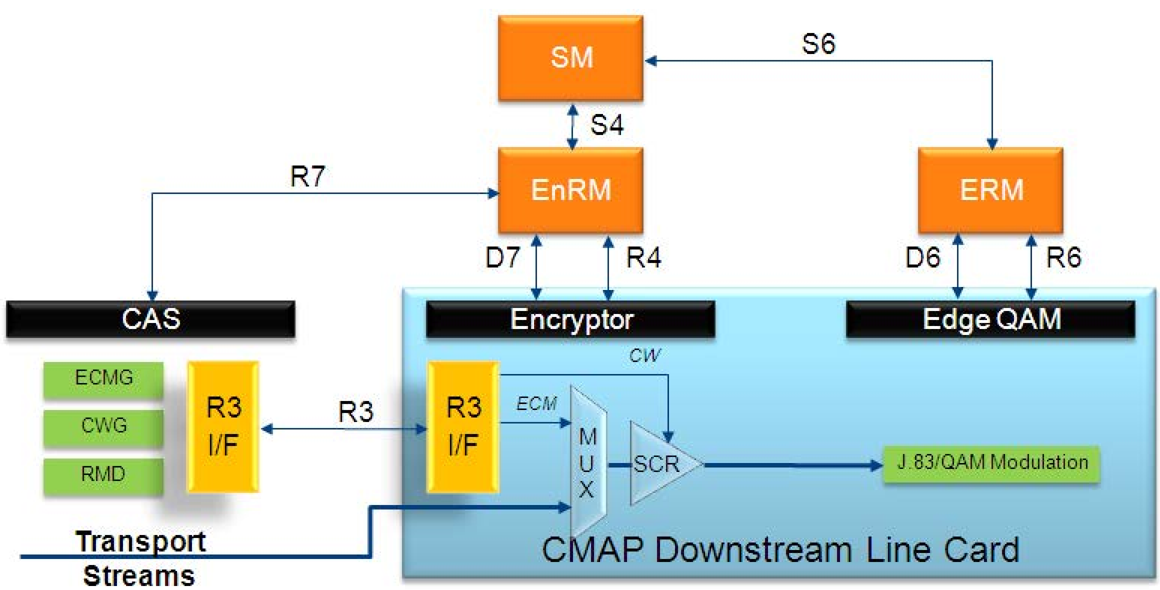

on Demand, or NGOD. As depicted in the figure below, the CMAP Downstream Line Card implements an Encryptor. This functional entity within the CMAP implements a superset of standard algorithms for scrambling content that work for most, if not all, CAS vendor systems. Additionally, the CMAP implements a set of interfaces, which are specific and defined by the CAS vendor for interfacing to an Encryption Control Message Generator (ECMG) and Control Word Generator (CWG). In turn, the CAS vendor would implement the ECMG and CWG according to the intricacies of their own CAS.

With this approach, the CMAP from any vendor that has entered the appropriate agreements with the CAS developer can implement full encryption system capabilities, including session-based scrambling. Not only is this not possible today, but today's technology allows only very minimal encryption functions by third party vendors, and requires very complex agreements to implement. The interfaces proposed in this approach are far more straightforward, revealing close to nothing regarding the CAS methods and procedures. Consequently, the agreements should be simpler to establish.

DEPLOYMENT EXAMPLES

Included in this section are examples of the possible deployment of a typical CMAP configuration in two types of systems, one implementing an HFC network with 750 MHz of spectrum capacity and another with 860 MHz of capacity. Both use cases are for typical systems, including a normal number of homes passed per node and per hub.

The two examples are for the deployment of a CMAP chassis consisting of the same configuration, as detailed in the following diagram.

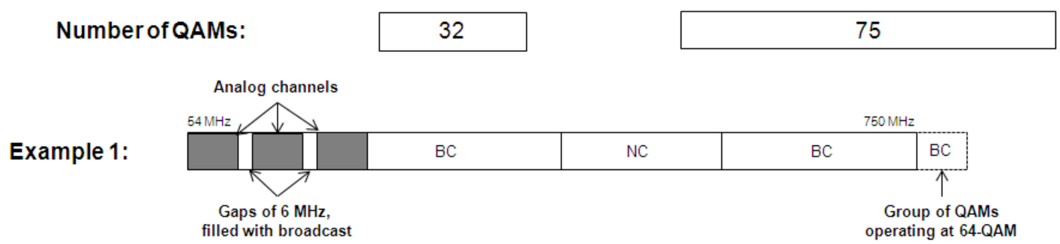

The CMAP chassis is capable of supporting a capacity of up to 64 QAM channels for narrowcast services and up to 96 QAM channels for broadcast services.
In the first example, detailed in the above figure, there is a group of analog channels (approximately 30) in the lower portion of the spectrum with a small number of gaps (2 as depicted in the example) consisting of a few 6 MHz channel slots. These gaps between analog channels are occupied by digital programs from the group of broadcast QAMs.

Additionally, there is a group of narrowcast QAM channels located towards the center of the spectrum. The remainder of the spectrum is occupied by broadcast QAMs -- a few of which are configured to operate in the roll-off portion of the spectrum and set to 64-QAM modulation, as opposed to all other QAM channels which will operate at 256-QAM modulation. Consequently, while all 32 narrowcast QAM channels would be used, approximately 75 of the broadcast QAM channels would be used in this example.

The second example, detailed in the figure at the bottom of this page, depicts the use of the same chassis.

In this example, the cable system is capable of supporting 860 MHz of spectrum. Similar assumptions for analog channels are made for this example, but additional narrowcast QAM channels are used instead and fewer broadcast QAM channels are needed to fill the available spectrum.

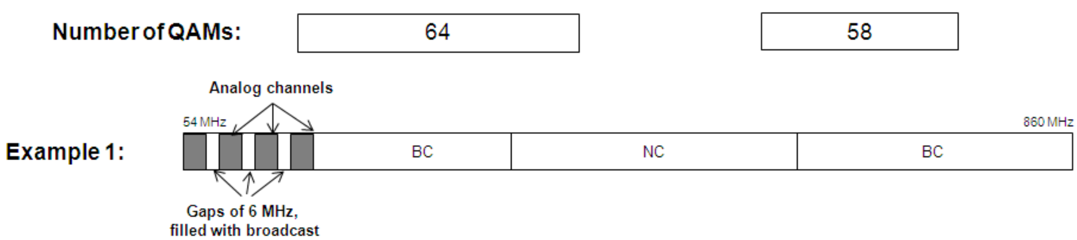

ENVIRONMENTAL EFFICIENCY

One of the key objectives of CMAP is to achieve significant environmental efficiencies. To that end, the following is an example of the space and power savings achieved by deployment of the CMAP in a typical system. The following figure depicts a typical installation in a headend consisting of the various digital services, including broadcast, SDV, VOD and HSD equipment, plus the corresponding combiner and lasers/receivers.

The example shown above is intended to serve a typical population of 200 nodes, combined in such a way as to result in 160 HSD service groups, and 120 VOD and matching SDV service groups.

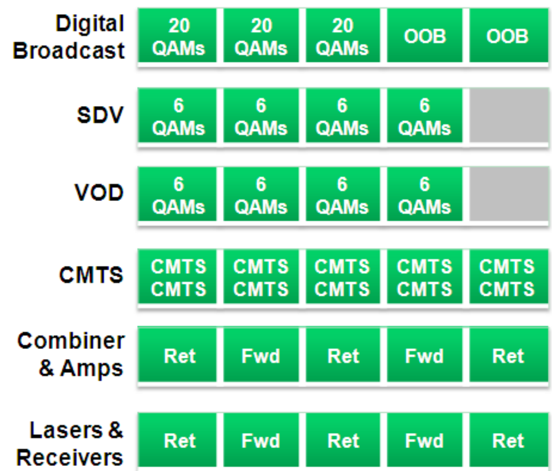

Considering typical CMTS and Edge QAM equipment as available today, the figure above depicts about 10 CMTS chassis, and about 4 racks for VOD and SDV, each containing 6 Edge QAM chassis configured for 64 QAM channels each at a density of 4 QAM channels per RF port. The digital broadcast lineup is composed of 60 individual QAM channels, plus the corresponding out-of-band equipment.

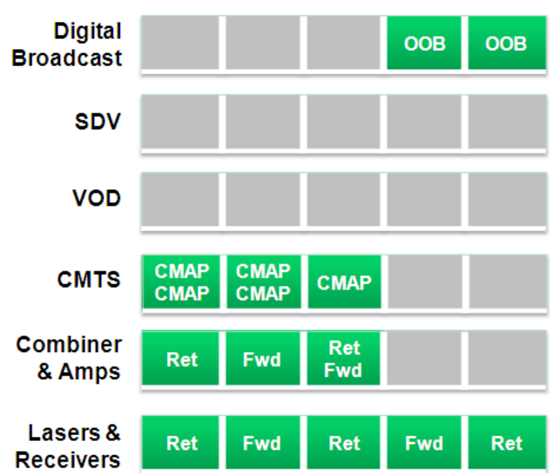

The following figure depicts the analogous installation when considering the deployment of equivalent CMAP equipment.

The above figure shows the following:

- Given that the CMAP chassis would have twice the density of a typical CMTS, only ½ the number of CMAP chassis are required, resulting in equivalent space savings.

- However, the CMAP chassis in its basic implementation includes all the necessary QAM channels for supporting the VOD and SDV services. Therefore, no additional equipment is needed to support these functions, resulting in significant additional space savings.

- Given that the CMAP also supports sufficient broadcast QAM channels, the space previously allocated to the broadcast QAMs channels is no longer needed, further contributing to space savings.

- Finally, it is estimated that ½ of the space allocated to the combiner network would be saved, resulting in even further space savings.

With all this taken into account, it is easy to see how as much as ½ of the space previously required is needed for deploying the CMAP. But, moreover, given that the CMAP can serve twice as many narrowcast QAM channels as the previous architecture could, the depicted CMAP scenario actually results in even greater space savings, providing twice as much capacity in ½ the space.

Clearly, the space savings are staggering!

In addition, it is worth considering the power savings.

A cursory analysis of the difference in power consumption, assuming typical power draw for existing equipment and the expected power consumption for the CMAP, yields an estimated power savings >50%. And, this is taking into account the use of 32 QAM channels in the CMAP, or 2x the capacity indicated in the original typical deployment. And, this does not even include the cooling savings from the great reduction of equipment and power consumption.

Without a doubt, the power savings are also very significant!

SILICON DEVELOPMENT

One important consideration is the evolution and availability of silicon components.

The functionality described by the CMAP specifications does not require of any new silicon. This is the case for both the upstream and downstream. CMTS vendors are already using and/or planning on making available line cards with existing and available high density burst demodulator silicon and corresponding MAC chips for upstream. For the downstream, vendors can utilize existing technology for Direct Digital Synthesis (DDS) consisting of readily available FPGAs and Digital to Analog Converters (DACs) from multiple vendors.
However, multiple silicon suppliers are in the process of implementing chips that provide very large QAM channel counts for downstream implementations. Some of these implementations are able to support QAM modulators for the entire RF spectrum from a single chip!

Even though these new silicon implementations are not required to develop a CMAP, they will certainly simplify designs, help reduce printed circuit board space and power/heat dissipation requirements, help reduce costs further, and accelerate development once the new silicon is available.

ANTICIPATED TIMELINE

As with any other technology evolution such as this one, things take longer than desired. On the flip side, their acceptance usually has farther reach than expected.

Clearly, from the many discussions with other MSOs, both within North America and throughout Europe and South America, interest for the deployment of this platform is very high. Almost without exception, MSOs at large are interested on the operational simplifications that the CMAP offers and the new functions it enables.

From preliminary discussions with vendors, and without revealing confidential information and plans, initial availability of equipment for laboratory and field testing is planned for late in 2011, early deployments are planned for 2012, and broad availability from multiple vendors by 2013.

CONCLUSIONS

The CMAP Platform is a viable alternative to the existing Modular Headend Architecture, and in fact may represent the Next Generation CMTS and Edge QAM.

CMAP implements all the QAM channels for each RF Port, supporting all digital services, including VOD, SDV, broadcast, HSD, and others in the future.

Given the environmental savings alone, field operations could benefit from CMAP immediately. Without CMAP, considering the expected growth in narrowcast services in the years to come, MSOs would likely have to resort to expensive headend/hub expansions.

Moreover, CMAP ports will be much more cost effective than current CMTS and Edge QAM ports, so the cost of expansion will be greatly reduced.

Finally, while CMAP can be implemented with existing technologies, it can greatly benefit from the natural technology evolution in chip development.

The challenge for MSOs is to know how and when to begin deploying CMAP. For some MSOs, the answer is as soon as the equipment can be manufactured.

ACKNOWLEDGEMENTS

Thanks to my colleagues John Bevilacqua, Peter Hutnick, Doug Jones, Saif Rahman, Arun Rajagopalan, Esteban Sandino, Joe Solomon, and Dave Urban from Comcast for their contributions to the CMAP program and specifications.

Also, many thanks to my colleagues from CableLabs®, Cable Europe Labs, the many vendors and advisors of the Comcast CMAP VE Team, and especially fellow MSOs for their many contributions to the development of the CMAP specifications.