# PUSHING IP CLOSER TO THE EDGE

Rei Brockett, Oleh Sniezko, Michael Field, Dave Baran
Aurora Networks

*Abstract*

*The ongoing evolution of cable services from broadcast video to narrowcast digital content (both data and video) has fuelled corresponding technical innovations to solve and support operators' operational and capital requirements. One area of particular interest is the QAM modulator. Accelerating subscriber demand for data and narrowcast video services will require a surge of new QAM deployments over the next several years, giving rise to a host of operational difficulties.*

*In this paper, we present the case for distributed headend architecture for HFC networks and discuss architectural and operational benefits of the Node QAM form factor, where the conversion of digital payload into QAM-RF signals is pushed from the headend to the cable TV optical node. In addition, we analyze the Node QAM in the context of the CableLabs® Converged Cable Access Platform (CCAP) architecture.*

## BACKGROUND

### Distributed Architecture Drivers

A key topic when discussing next-generation cable infrastructure is the balance between analog optical transmission, including the transmission of multicarrier QAM-RF signals, and baseband digital transmission of signals such as native Internet Protocol (IP) signals. Cable operators have gone through several transitions already, with the introduction of digital television; the growth of high-speed data; the use of IP-based distribution in the headend; and the use of native baseband IP-based communication between headends and hubs. The driving force has always been efficiency and cost.

The imperative to meet subscriber demands results in certain bottlenecks: physical space and power within the headend, bandwidth capacity in the deployed HFC, distance between headend and subscriber, limitations of hard-wired infrastructure.

For each of these areas, there are solutions, but a distributed headend architecture that extends the boundary point where content enters the RF domain addresses all of these:

- Headend space and power consumption can be mitigated by consolidating functionality and increasing port densities in next-generation CMTSs and Edge QAMs. Alternatively, functionality can be distributed to the hubs and nodes, leaving only the IP network and MPEG2-TS processing in the headend. Direct generation of RF output at the edge of the network eliminates the need for an RF combining network at the headend. This reduces headend space and power requirements and simplifies network operations by avoiding the need to mix signals in the RF domain.

- Distance limitations can be relaxed by pushing deeper the conversion of digital signals to RF. Analog optical transmitters and amplifiers are at the limits of their capabilities, and add expense and design complexity. However, by extending the headend IP domain to the node, not only is optical transmission distance extended, but RF signal loss budgets are mitigated and higher loss budget at higher frequencies can be accommodated, thus increasing bandwidth capacity of the subsequent coaxial section of the HFC network. For example, baseband optical links to the node would eliminate analog link contributors to signal degradation, thus allowing for higher modulation levels and hence better spectral efficiency in the available coaxial bandwidth. This can be especially effective and fruitful in passive coaxial networks (PCN), also known as Fiber Deep, Fiber to the Curb (FTTC), or Node-plus-zero (N+0) HFC networks.

- In addition to the effect of explicit signal impairments due to analog optical transmission, bandwidth capacity in the HFC network is further constrained by the complexity of carrying analog (RF) signals over distance. In the optical links to the nodes, the use of multiwavelength systems, while justified by fiber scarcity and revenue opportunities, introduces severe constraints on the usable number of wavelengths and their link performance. Impairments from analog (RF) modulated optical transmitters and erbium-doped fiber amplifiers (EDFAs) further limit the capacity of individual wavelengths. Converting from RF modulated transmitters to baseband digital optics would eliminate these impairments and increase the number of cost-effective wavelengths to 88 (yielding 880 Gbps of capacity to each node) using current technology, with room for growth in the number of wavelengths and the wavelength capacity of next-generation optics.

- The challenge of managing bandwidth allocation between unicast, multicast, broadcast, and data QAM signals is eliminated by mixing content dynamically in the headend IP network. This allows bandwidth to be allocated as-needed in response to market requirements without requiring "hands on" labor.

### Accelerating Demand for Narrowcast Services

Rapidly evolving subscriber behavior surrounding the consumption of multimedia is driving cable operators to confront two challenges. The first is the need to significantly accelerate the deployment of narrowcast services while also accommodating bandwidth-intensive services such as HDTV and 3DTV. These narrowcast services typically include high-speed data and packet voice, video on demand (VoD), and switched digital video (SDV), but also encompass other unicast and multicast services such as cable IPTV, network-based digital video recording (nDVR), and other services that leverage the IP cloud at the headend. The second challenge is the difficulty of planning a graceful and cost-effective migration from inefficient and obsolete service silos to new, dynamic methods of flexibly allocating capacity to different services in the face of constantly shifting customer demands.

The need to deploy an unprecedented volume of new QAM modulators is common to both challenges, and this raises concerns over issues including headend environmental constraints, flexibility of service allocation, RF combining issues, HFC transmission considerations, and the need to accommodate legacy equipment.

In these circumstances, one viable solution that achieves the benefits listed above is to relocate the QAM modulators to the HFC node, pushing the native baseband IP domain even further to the edge (closer to the user — the ultimate edge of the HFC network).

## DESIGNING A NODE QAM

### A Confluence of Technology and Need

Quadrature Amplitude Modulation (QAM) is a spectrally efficient way of using both amplitude and phase modulation to transmit a digital payload on an analog carrier. Cable QAM modulators[1] operate on packets in the MPEG2-TS format, and modern QAM modulators include integrated upconverters as well.

In the decades since the first baseband QAM modulators were assembled out of discrete components, silicon technology has increased a thousand-fold in processing price performance, and decreased a hundred-fold in size, giving rise to a surprisingly rich selection of special-purpose, general-purpose, and programmable chips, based on which we can re-design our modulators.

These advances can finally be used to their full advantage now that demand for modulators has swelled from tens and twenties per headend to hundreds and even thousands. Part of the advantage is in the availability of brute-force processing power, but a companion advantage is in algorithmic efficiencies derived from being able to perform certain steps in bulk. One result is that existing headend Edge QAMs can be made much denser, with thousands of QAM channels in a chassis. Another result is that it is now operationally feasible to put a full gigahertz' worth of QAM channels (or more) in the node.

## Node QAM Requirements

The node is a hostile environment for advanced electronics. Power budget and space are limited; cooling is passive; operating temperatures can be extreme; and accessibility is limited. In order for a Node QAM to be operationally neutral when compared to a headend Edge QAM, it must meet the following criteria:

- Low power. In order to avoid the need for non-standard node powering, a full-spectrum Node QAM must be able to generate at least 158 (6 MHz) QAM channels using the same amount of power as a traditional optical receiver. This eliminates the need for active cooling.

- Compact. The Node QAM should be designed to fit within the existing, field-proven node housings.

- Industrial grade operating temperature range (–40°C to +85°C). Unlike climate-controlled headends, or even cabinet-based hubs, components in the node must be able to withstand large fluctuations in temperature.

- Reliable. Servicing a node is logistically cumbersome and operationally expensive. A Node QAM must be robust and uncomplicated. Additionally, remote monitoring is critical. Ideally, cost, space, and power consumption profiles can be kept low enough to enable the deployment of spare modules, which would allow operators high levels of redundancy, even at the node level.

- Simple to install — "Set it and forget it". Installing a Node QAM must be as simple as plugging in a module and verifying the output with a field meter. Complex procedures such as configuration and management should be done centrally, to simplify operations.

- Low cost. Per-channel equipment costs need to keep pace with the cost of headend Edge QAMs.

- Future-proofed. Given the logistical difficulties of servicing nodes, the distributed Node QAM modules should have a margin for upgradability so future technological changes and additions can be accommodated by re-programming the existing modules. This not only simplifies architectural evolution, but also extends the operational lifetime of each module.

This is also applicable to the interfaces between the node modules and the headend/hub infrastructure; new modules can be introduced in a very scalable manner if they leverage the standard data networking interfaces used in the IP network in the headend.

These requirements, while difficult to achieve, are attainable given modern silicon capabilities and careful design, opening up the option to move to a more distributed architecture, with many of the benefits.

## ARCHITECTURAL BENEFITS

Generating some or all QAM signals at the node results in a number of advantages.

### Exploiting Digital Optics

A major advantage of moving the QAM modulator to the node is the ability to shift to digital optics between the headend and the node. In traditional usage, electrical RF signals are amplitude-modulated onto an optical signal. These signals are extremely sensitive to various fiber nonlinear distortions like cross phase modulation (XPM), stimulated Raman scattering (SRS) and optical beat interference (OBI) caused by the four-wave mixing (4WM) products that come into play depending on power, distance, wavelength count, and other factors. Together with other nonlinear and linear fiber impairments, they limit the capacity of the links and significantly impair transported signals. Designing and "balancing" optical links to the nodes in an HFC system is a delicate art. Furthermore, the lasers modulated with analog (RF) multicarrier signals have limited Optical Modulation Index (OMI) capacity due to the fact of high sensitivity of these signals to clipping. The limits reach up to 30% for directly modulated lasers and approximately 20% for externally modulated lasers. These limits, with the operational back-off of 2-3 dB, severely limit the capacity of every single wavelength in any multiwavelength system of practical distance.

Using baseband digital transmission is much simpler. Because data is not as sensitive to nonlinearities and other impairments, not only can distances be extended, but more wavelengths within a single fiber can be employed, resulting in higher bandwidth capacity to the node. Simpler and more economical optics and amplifiers can be used, as well. With their OMI approaching 100%, digital optics enable significant increases in the capacity and distance of each fiber optic link. Using existing technologies, they can support cost-effective transmission of 88 wavelengths with 10 Gbps/wavelength over distances in excess of 100 km from the IP headend/hub infrastructure. This opens significant opportunity to provide unparalleled bandwidth to the nodes for residential services as well as significant opportunity for additional revenue.

More cost optimizations for capacity and distance can be achieved by leveraging lower-cost optical amplifiers, simplified optical filters, and symmetric and asymmetric SFP, SFP+ and XFP transceivers. Furthermore, deploying distributed architecture and transmitting native baseband IP signals to the node finally enables HFC to take advantage of the high-volume economies of scale in modern digital (data) networking infrastructure, which outperforms the economies of scale for analog cable TV optics a thousand-fold.

Another benefit of using baseband digital optics between the headend and the node is the elimination of an HFC weakness: the analog link contribution to end-of-line noise budgets. The analog (RF) optical links to the nodes with analog (NTSC or PAL) video signals are designed for 47 to 50 dB carrier to noise ratio (CNR) for occupied bandwidths ranging from 700 to 950 MHz. For QAM signals placed on the same link, it translates to modulation error ratio (MER) between 39 and 42 dB. For links with QAM-only load, this limit is usually lowered by designers to 37 dB MER to take advantage of cost tradeoffs and increase fiber utilization efficiency and reach. This is sufficient to support a modulation order of 256-QAM, but it limits the capacity of the HFC link to between 5 and 6.4 Gbps. Improved noise budgets by using the Node QAM would allow the support of 1024-QAM modulation, over a bandwidth range up to 1800 MHz, resulting in throughput capacity of 15 Gbps, nearly triple the current capacity.

Digital baseband transmission would unlock practically unlimited capacity in the fiber links to the node. With the proximity of the node to the furthest service user, especially in PCN networks, distributed fiber to the home (FTTH) solutions like Next-Gen RFoG and xPON can be extended from the nodes selectively, based on the demand and opportunities.

Simplification of RF Combining Network

Generating QAM signals in the node allows those QAM signals to bypass the RF combining network. Node QAM output signals can be combined *at the node* with traditionally carried HFC signals in a single stage. New narrowcast QAM signals can be added at the node as needed, with no impact on either the existing RF combining network or the HFC plant alignment.

Besides removing the complexity of recalculating the headend combining plant each time new RF ports are added, it avoids both the signal and power losses associated with combining, splitting, and directional coupling, as well as the power, cooling burden, and significant space inefficiencies. Many of these advantages are delivered by the CableLabs Converged Cable Access Platform (CCAP)[2] architecture, as described later. A distributed architecture goes a step further by allowing legacy signals to be combined in a single passive combining stage at the node.

RF Signal Quality and Node Alignment

When QAM signals are generated in the node (Figure 1), with given output levels and the same or better output signal quality as headend-generated QAM signals, the resulting RF signal in the node is much cleaner. This is because it bypasses the signal losses, noise, attenuation, and distortions that are typically introduced in the RF combining network and the amplitude-modulated optical links to the node.

Operationally, it reduces the amount of RF aligning needed at the node; output power and tilt are generated exactly according to configured specifications, defined by the operator. The signal is not subject to any of the traditional distortions. The impairment contribution of combining network and analog (RF) optical links to nodes is eliminated, with the benefit of unlocking coaxial plant capacity as described above. This allows a 43+ dB MER (see Figures 1 and 2) at the node and gives the operator more options in the coaxial portion in terms of loss budget/coverage and, most importantly, bandwidth. In certain conditions, it makes higher-order modulation rates possible as well, resulting in better spectral efficiency.

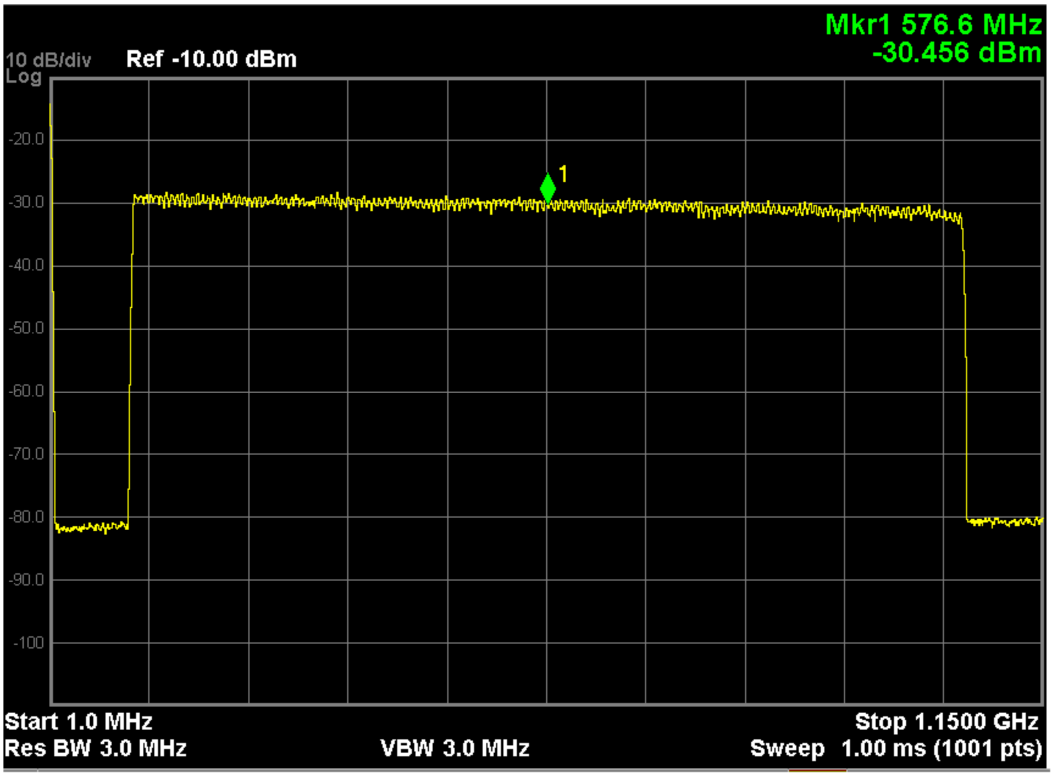


**Figure 1: 158 Node QAM channels**

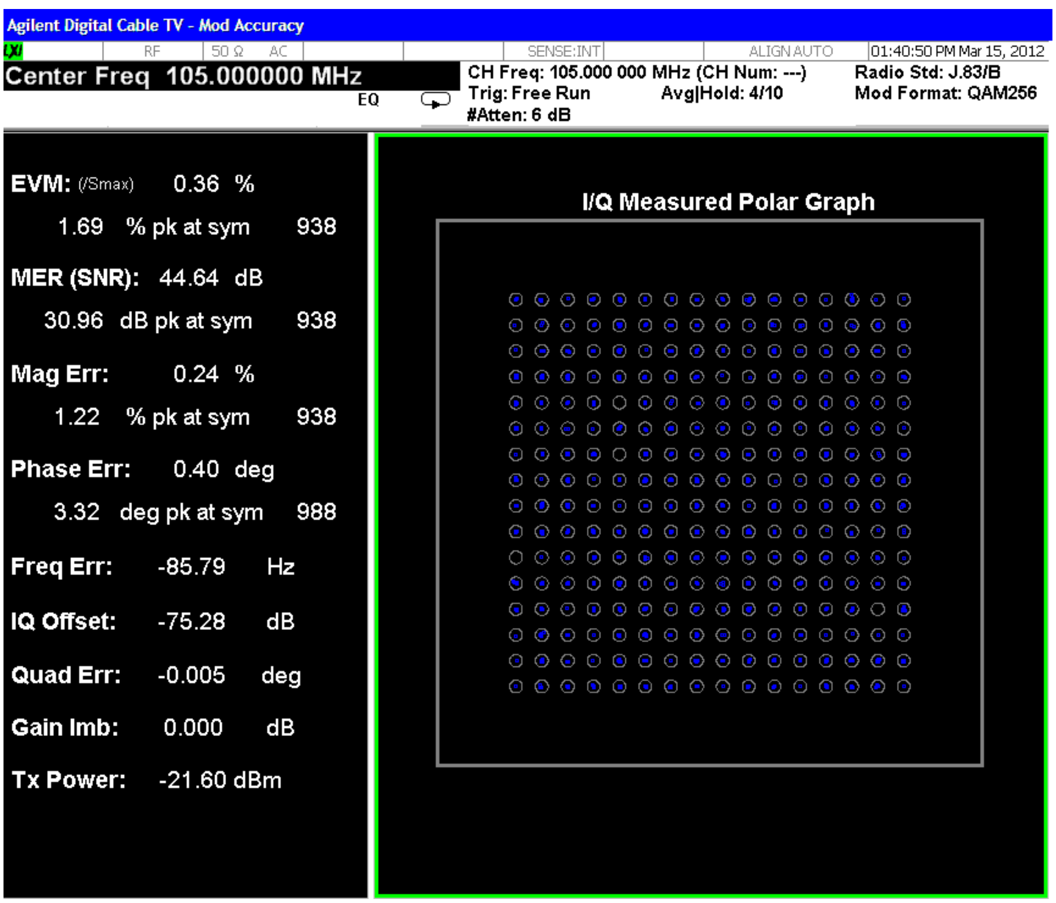


**Figure 2: Node QAM increases RF loss budget or bandwidth capacity.**

Service Flexibility

An important side effect of the Node QAM is that the optical network feeding it is a de-facto extension of the headend IP network, with access to all of the system's digital content — broadcast, narrowcast, unicast, and data.

The Node QAM itself is agnostic to the digital payload; it simply modulates the MPEG2 formatted transport streams that are delivered over the optical interface. The payload carried within the transport stream could be a groomed and re-quantized statistical multiplex; it could be an encrypted variable bit rate broadcast multiplex; it could be a simple multiplex of fixed-rate VoD streams; or it could be a DOCSIS M-CMTS-compliant data stream.

The contents of the transport streams are dependent only on the capabilities and sophistication of the headend service manager(s) and resource manager(s), and switched IP connectivity. Artificial service group constraints imposed by the hard-wired RF combining network are removed, leaving only a general-purpose pool of QAM signals to feed the population of subscribers attached to each node or node segment.

An enhancement enabled by the generation of QAM signals in the node from native IP input is the ability to selectively reserve local bands of frequency for other modulation and encoding schemes as well. See Figure 3.

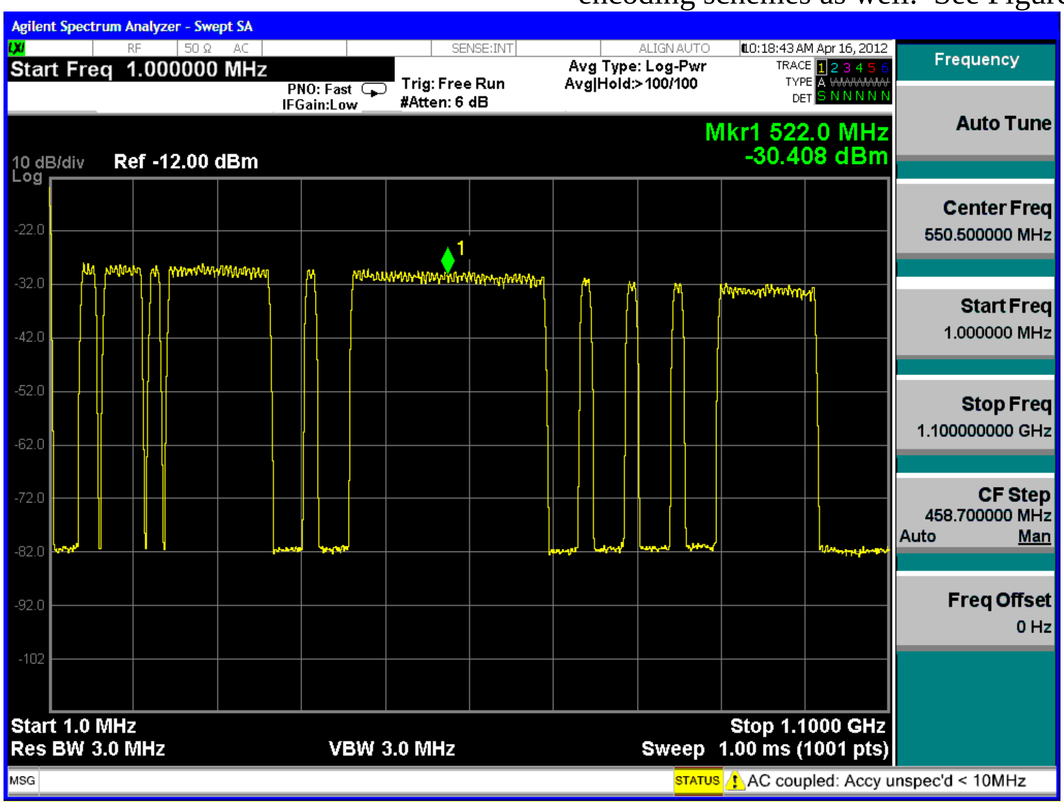


**Figure 3: Spectrum Allocation Agility. Individual QAM signals can be turned on or off.**

Some examples of practical applications include:

- Customized broadcast lineups. Certain niche customers, such as hotels, apartment complexes, hospitals, and campuses can receive their own broadcast lineups, created on the fly, without affecting the existing RF combining network.

- Uneven service usage. Usage of individual types of narrowcast and unicast services may vary unpredictably from node to node. Node QAMs with headend service switching allows each node to have a different service mix, without having to pre-allocate resources.

- Dynamic service allocation. Service usage may also vary within a single node, based on time of day or season. For example, a suburban node might experience heavy VoD usage during the day due to toddler addictions to children's programming, but switch to heavy internet usage late at night when parents use Netflix. With the Node QAM, a single pool of QAM signals can feed all services, without having to provision under-utilized service silos.

- Mixed services within a single channel. With sufficient sophistication from the headend multiplexers and resource managers, the Node QAM can deliver any mix of QAM services — broadcast, narrowcast, CMTS, VBR, CBR in a single channel, giving the operator complete flexibility.

Environmental

While modern headend QAM modulators are an order of magnitude more energy-efficient than earlier incarnations, and two orders of magnitude more compact, the addition of large quantities of new QAM channels via traditional methods creates a significant impact on the headend, in two ways. Headend Edge QAMs create a direct impact by their intrinsic consumption of power, rack space, and cooling mechanisms. They also have an indirect impact, due to the rack space occupied by the combining network; the power loss due to combining, splitting, and directional coupling of service groups, as well as the power consumption of intermediate amplification stages; and the power burden of heating, ventilation, and air conditioning (HVAC).

By moving QAM modulation to the node, not only are power and rack space requirements distributed, but overall per-QAM power and space consumption are reduced due to the fact that lower output levels are needed to drive the existing node RF amplification modules. This helps the Node QAM to live within the design constraints imposed by the node housing, including the use of passive cooling instead of fans. Node QAMs also eliminate the Edge QAMs' impact on the headend HVAC system.

In addition, by bypassing the RF combiner network at the headend, Node QAMs avoid wasting the signal power maintained by the RF combiner network's amplification stages, which end up being discarded when the signal is carried in its baseband digital format. Furthermore, power and space requirements are reduced when optical analog (RF) transmitters are replaced by low-power optical digital baseband transceivers.

These Node QAM benefits mesh well with the fundamental goals of the CCAP architecture, with the added advantages that Node QAM leverages digital optics, and that these benefits accrue on a node-by-node basis, allowing both small and large operators to migrate gracefully to CCAP.

## CCAP

CableLabs' CCAP architecture is a bold step in addressing many of the challenges related to the growth of narrowcast services. It leverages heavily the existing body of Data-Over-Cable Service Interface Specifications (DOCSIS) with the goals of increasing the flexibility of QAM usage and configuration; simplifying the RF combiner network; possibly adding content scrambling; creating a transport-agnostic management paradigm to accommodate native support of Ethernet Passive Optical Network (EPON) and other access technologies; improving environmental and operational efficiencies; and unifying headend configuration and management capabilities. CCAP includes a new Operations Support System Interface (OSSI)[3] specification and also takes particular care to ensure compatibility with existing DOCSIS resource management and service management and configuration specifications, in order to facilitate the migration from current CMTS/Edge QAM infrastructure.

CCAP Reference Architectures

CCAP unifies digital video and high-speed internet delivery infrastructures under a common functional umbrella, allowing a CCAP device to be operated as a digital video solution, a data delivery solution (both CMTS and M-CMTS), a Universal Edge QAM, or any combination. Each of the CCAP reference architectures (Video, Data, and Modular Headend) describe physical and functional interfaces to content on the "network" side, operational and support systems within the headend, and the HFC/PON delivery network terminating in various devices at the subscriber premises. Ancillary service and resource managers are allowed to exist both within and externally to a CCAP device.

CCAP OSSI

The lynchpin of the CCAP architecture is the CCAP OSSI, which defines a converged object model for dynamic configuration, management, and monitoring of both video and data/CMTS functions, but also makes provision for vendors to innovate within the framework. By creating a unified standards-based operational front-end to the video and data delivery infrastructure, CCAP OSSI provides a solid foundation for the headend's metamorphosis from a collection of separately managed service silos into an efficient service delivery "cloud".

CCAP and Node QAM

In the CCAP video and data reference architectures, the CCAP interface on the subscriber side is the HFC network. Traditionally, that interface exists within the headend. However, there is nothing inherent about the provisioning and management of QAM signals that *requires* the QAM modulators to be in the headend. Extending the logical boundary of the headend out to the node and minimizing the analog portion of the HFC remains consistent with the goals and specifications of CCAP.

## NODE QAM EVOLUTION

Initial Architecture

The initial configuration of the Node QAM topology can be envisioned as one presented in Figure 4. In this configuration, analog and operator selected QAM broadcast channels (*e.g.*, from a different location than the remaining QAM channels) are transported to the node in a traditional fashion but without the burden of combining with the remaining QAM channels in the headend/hub. The number of QAM channels originating in the Node QAM can be adjusted dynamically by the operator.

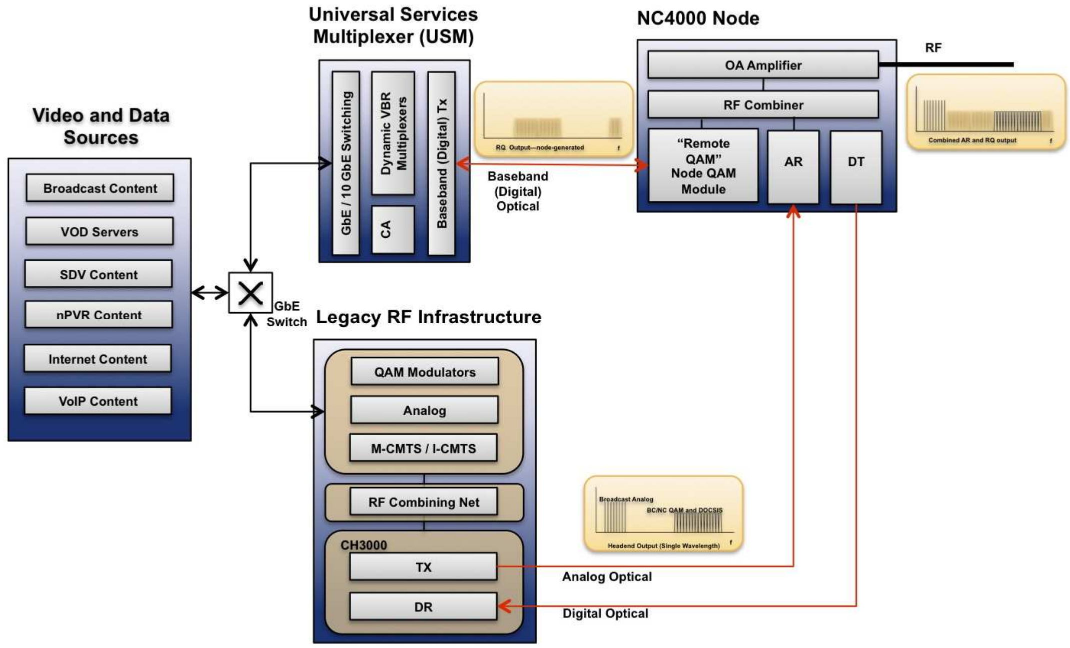

Figure 4 Node QAM Initial Implementation

<u>Conversion to Complete Digital Baseband Node Transport</u>

The next incarnation of the distributed architecture is presented in Figure 5. All analog channels and maintenance carriers are digitized in the headend and transported over the same transport (capacity allowing) to the node where they are frequency-processed and converted back to analog channels at their respective frequencies on coaxial plant. Some additional carriers (*e.g.*, ALC pilot signals) are synthesized in the Node QAM module.

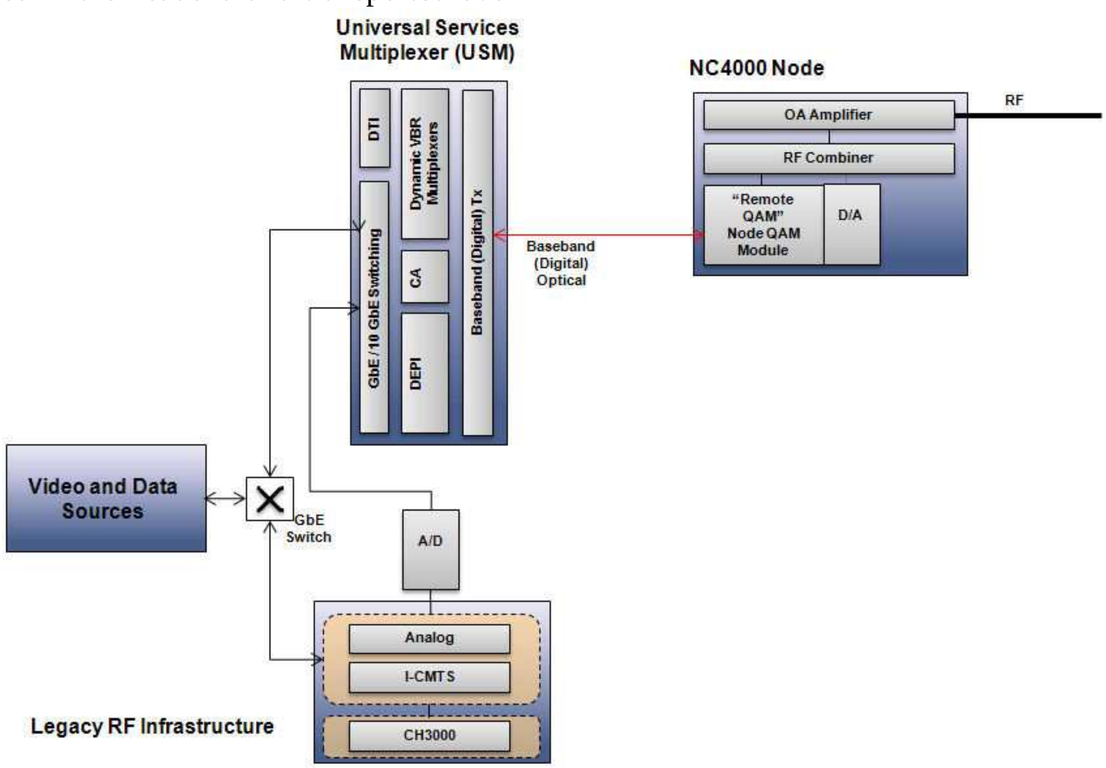

Figure 5: Node QAM Next-Generation

The reverse channel(s) from the node to the headend can also be converted to baseband digital optics, resulting in similar benefits. Options include traditional digital return (digitization of the return spectrum at the node), developing a node-based CMTS (or node-based DOCSIS burst receivers), or even next-generation native IP-over-coax technologies.

A related enhancement arising from the Node QAM's dynamic frequency agility is the ability to support flexible, remotely configurable frequency splits or capacity allocation between downstream and upstream communication, either using frequency division duplex (FDD) or time division duplex (TDD) transmission. This would enable full flexibility and adaptability to downstream and upstream traffic patterns and capacity/service demands.

### Future Enhancements

The Node QAM is an ideal platform to be modified to support other modulation schemes for next-generation transport mechanisms, such as EPON Protocol over Coax (EPoC). Implementing EPoC in the node allows significant reach expansion, preserving and facilitating headend and hub consolidation without deploying additional signal conditioners or RF-baseband-RF repeaters with their additional cost, power consumption, added operational complexity of provisioning and additional space/housing requirements in the field or hubs.

## OTHER ELEMENTS OF DISTRIBUTED ARCHITECTURE

### Node PON

A distributed node-based EPON architecture shares the Node QAM architectural advantages. Node PON modules allow for selective fiber placement from the node for commercial services in node areas where construction costs and effort are limited to fiber extension from the node. In PCN architecture, this is usually below 1 km, and mostly below 300 m if the node is placed strategically. In conjunction with DOCSIS Provisioning of EPON (DPoE) and CCAP, Node PON can address the needs of fast deployment of dedicated fiber links to selected high capacity demand users.

### Next-Generation RFoG[4]

In situations where fiber exists all the way to the subscriber, RF over Glass (RFoG) in a distributed architecture has the potential, with minor changes, to exceed the throughput of 10G PON/EPON, without the complexity of adding a PON overlay. This allows for seamless expansion of fiber from RF optical nodes to residences without replacing the distributed architecture node modules. Taking fiber from the node all the way to the subscriber with a FTTH network would allow for additional capacity enhancement beyond 15 Gbps downstream and 1 Gbps upstream facilitated by distributed coaxial architecture, especially with PCN and residential gateways deployed. With RFoG in a distributed architecture, 20+ Gbps downstream and 3 Gbps upstream is achievable today without PON overlay.

## SUMMARY

No-one knows precisely what the future will bring but it is clear that subscriber-side demand for IP-delivered multimedia continues to grow as “smart” home and mobile electronic devices proliferate. The cable industry is blessed with the most extensive and highest bandwidth conduit to that last-mile “IP cloud”. At the same time, cable headends have largely already made the transition to IP-based distribution. Moving the native baseband IP-to-RF transition point from the headend to the node brings the convergence of IP headend and IP home one step closer.

As discussed in this paper, there are many advantages to extending the digital headend domain as far into the network as possible, in terms of performance, resource utilization, operational simplicity, and service flexibility. There are many paths for the evolution to digital HFC: the Institute of Electrical and Electronics Engineers (IEEE) is proposing a new physical layer standard called EPON-Protocol-over-Coax (EPoC) to deliver IP traffic natively at 10 Gbps over last-mile HFC; fiber vendors continue to innovate on bringing fiber to the home; new silicon may enable conversion of large bands of RF spectrum at the headend into digital bitstreams that can be converted back to analog at the node. By bringing IP closer to the edge, the Node QAM helps pave the way to a distributed headend and digital HFC.

## ABBREVIATIONS AND ACRONYMS

| | |
|---|---|
| 10G-EPON | IEEE 802.3 Ethernet PON standard with 10 Gbps throughput |
| 3DTV | 3D Television |
| 4WM | Four Wave Mixing |
| ALC | Automatic Level Control |
| BER | Bit Error Rate |
| CBR | Constant Bit Rate |
| CCAP | CableLabs® Converged Cable Access Platform |
| CMTS | Cable Modem Termination System |
| CNR | Carrier-to-Noise Ratio |
| DOCSIS® | Data over Cable Service Interface Specification |
| DPoE™ | DOCSIS Provisioning of EPON |
| EDFA | Erbium-doped Fiber Amplifier |
| FDD | Frequency Division Duplex |
| FTTC | Fiber to the Curb |
| FTTH | Fiber to the Home |
| EPoC | EPON Protocol over Coax |
| EPON | IEEE 802.3 Ethernet PON standard with 1 Gbps throughput, a.k.a. 1G-EPON, G-EPON or GEPON |
| Gbps | Gigabits per second |
| HDTV | High Definition Television |
| HFC | Hybrid Fiber Coaxial |
| HVAC | Heating, Ventilation and Air Conditioning |
| IEEE | Institute of Electrical and Electronics Engineers |
| IP | Internet Protocol |
| IPTV | IP Television |
| M-CMTS | Modular Cable Modem Termination System |
| Mbps | Megabits per second |
| MER | Modulation Error Ratio |
| MPEG2 | Motion Picture Experts Group 2 standard |
| MPEG2-TS | MPEG2-Transport Stream |
| nDVR | Network-based Digital Video Recording |
| NTSC | National Television System Committee |
| OBI | Optical Beat Interference |
| OMI | Optical Modulation Index |
| OSSI | Operations Support System Interface |
| PAL | Phase Alternating Line |
| PCN | Passive Coaxial Networks |
| PON | Passive Optical Network |
| QAM | Quadrature Amplitude Modulation |
| RF | Radio Frequency |
| RFoG | Radio Frequency over Glass |
| SDV | Switched Digital Video |
| SFP | Small Form-factor Pluggable |
| SRS | Stimulated Raman Scattering |
| TDD | Time Division Duplex |
| VBR | Variable Bit Rate |
| VoD | Video on Demand |
| XFP | 10 Gigabit Small Form-factor Pluggable |
| XG-PON | ITU-T's broadband transmission standard with 10 Gbps throughput |
| XPM | Cross Phase Modulation |
| *x*PON | any of a family of passive optical network standards (*e.g.*, GPON, GEPON, 10G PON (BPON, GEPON or GPON) |

[1] ITU-T J.83 Digital multi-programme systems for television, sound and data services for cable distribution. April 1997.
[2] TR-CCAP-V02-110614. CCAP Architecture Technical Report. June 2011.
[3] CM-SP-CCAP-OSSI-I02-120329 . Converged Cable Access Platform Operations Support System Interface Specification. March 2012.
[4] O. Sniezko. *RFoG: Overcoming the Forward and Reverse Capacity Constraints.* NCTA Spring Technical Forum 2011.