SCTE CABLE-TEC EXPO 2013 TECHNICAL PAPER

BEYOND CCAP v2 – NEXT, NEXT GENERATION ACCESS NETWORK ARCHITECTURE (N^2GAN)

Jeff Finkelstein, Executive Director of Strategic Architecture
Cox Communications

Boris Brun, Senior Manager of Technologies
Harmonic, Inc

Abstract

The Hybrid Fiber-Coax (HFC) plant has significant potential for providing the capacity our customers demand for many years. However this remains an untapped potential as we continue to exist in todays deployment scenarios of N+4 or greater. There are limitations to higher modulation orders resulting from today's architecture.

The Converged Cable Access Platform (CCAP) effort has defined a technology that provides the combined functionality of legacy edgeQAM, data processing of CMTS (Cable Modem Termination System), and IP (Internet Protocol) video processing. CCAP is able to provide full downstream spectrum through a single port that allows operators to have complete flexibility in deployment of services throughout the full range of channels. But what is next?

While emerging CCAP solutions leverage advanced technologies to solve many of today's infrastructure problems, there will be an inflection point where even these new devices will be challenged by the capacity needs of the future. These capacity needs not only create challenges from a spectrum perspective, they expose inherent head-end space and power limitations. These challenges, combined with the cost, capacity, and distance limitations of linear optics, will drive us to develop new ways of solving the bandwidth shortfall.

One approach is to increase the upstream and downstream bandwidth. While current HFC plants support of 5-42 or 5-65 MHz in the upstream and 88 MHz-750/860 MHz/1 GHz in the downstream, the number of available channels and the current modulation order limit the theoretical maximum bandwidth. Discussions are underway regarding alternative spectrum splits. A mid-split at 85 is being considered to increase upstream bandwidth, and possible use of the band above 1 GHz.

Another approach is to make the current bandwidth more efficient. There are evolving approaches to addressing the capacity problem. DOCSIS 3.0 allows bonding of many downstream DOCSIS channels using statistical multiplexing techniques to increase bandwidth for a single device or grouping of devices. DOCSIS 3.1 will provide the ability to reach even

higher speeds and capacity by using OFDM (Orthogonal Frequency Division Multiplexing) technologies to allow higher orders of modulation than was achievable with previous versions of DOCSIS. Discussions are underway regarding alternative spectrum splits such as 5-85 and 5-200 MHz for upstream, and possible use of the band above 1 GHz.

The purpose of this paper is to propose a complementary technology to monolithic CCAP. The proposal centralizes control and intelligence into the data center and distributes processing into the HFC network. We consider this a paradigm shift from previous technologies presented to such an extent that we call it the Next, Next Generation Access Network (N^2GAN) architecture.

HFC Network Capacity Today

The capacity of an HFC network is largely limited by the frequency ranges, or bandwidth, of actives and passives used to pass, direct, and amplify signals throughout the cable plant. The capacity of this bandwidth is further limited by the attenuation on the coaxial portion of the plant. The higher the frequency of the transmitted signal, the higher the attenuation of that signal, which limits the maximum distance that the signal may be effectively transmitted.

Capacity requirements are being driven by increased demand for high-speed data services, the proliferation of video products, teleconferencing, IP-based services replacing or augmenting existing legacy services, and other emerging applications. Operators have experienced annualized growth rates of 40% to 60% over the past few years with no sign of slowing down. Emerging applications and improvements – quality and functionality – in legacy applications are driving requirements for more bandwidth and capacity, and future applications will have even greater needs than previously thought possible. As an industry, we have learned that bandwidth growth across the backbone, metro, and access networks continues to increase at unexpectedly high rates. One immutable fact operators must come to terms with is that we have no control over how our customers use bandwidth.

Today, typical HFC upstreams use 5 to 42 MHz in the US and many other countries. Europe leverages EuroDOCSIS and uses 5 to 65 MHz. HFC forward networks have active components that normally support up to 750 MHz, 860 MHz, or 1 GHz, but in some deployments passive elements can support up to 3 GHz. We can consider future return network upgrades to 85 MHz or even 200 MHz in the return, which coupled with DOCSIS 3.1 will enable downstreams of up to 10 Gbps and 1.5 Gbps in the return path using higher order modulations from 10-12 bits per second per hertz.

How much speed/bandwidth/capacity do we really need?

The short answer is, we don't know. We will not make predictions but every indication is that the 50% growth rate will not change appreciably until acted upon by some external force. No, not the external force you remember from Newton's first law. The external forces we are talking about are competition and regulation.

Three primary forces drive increased demand for high-speed data capacity: application needs, competitive pressure, and regulatory requirements. The prevailing wind speed of the Internet, the speed that can be achieved over time end-to-end by individual home users, is currently approximately 30 Mbps (as of this writing in 2013). To measurably change this 30 Mbps limit, many interconnects between providers will need significant upgrades.

New video services, such as IP video, network DVR (nDVR), and Video on Demand (VOD), also have impact on bandwidth. These services will likely add significant traffic to the network, with nDVR requiring 4 to 5 times the bandwidth used by VOD.

Together, the increasing growth in high-speed data and video easily reaches the projected numbers of a 40% to 60% compound annual growth.

Eras of Consumer Speeds

Historical trends of user speed consumption show no sign of decreasing usage. When we project those consumption into the future we see that if the growth continues undiminished, we reach a point where DOCSIS 3.0 technologies can no longer meet the requirements of user demand. At that point, the cable industry will need to find a new technology to keep up with this growth.

When we review the evolution of technology, we see that new methods of increasing bandwidth were developed just in time to meet demands. From 300 bps dial-up in the early 1980's, to 56 kbps dial-up in the 1990's, to DOCSIS 1.0/1.1 in the early 2000's, to DOCSIS 3.0 today, technology has kept pace with consumer speed requirements. Or possibly the reverse is true - the consumer has pushed dial-up and cable networks into the future.

The following graph shows how over time new technologies were deployed to meet user capacity needs.

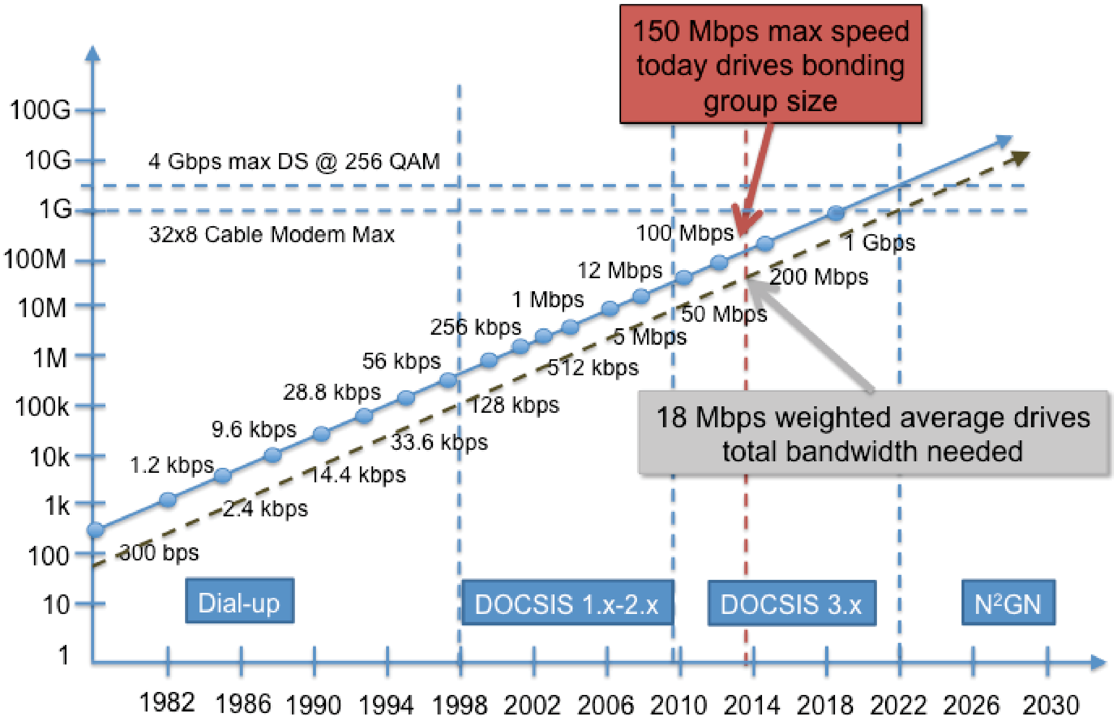

Figure 1 - Eras of Consumer Speeds

As the above chart illustrates, you can see that if the trend continues, DOCSIS 3.x capabilities will be exceeded within 10 years. As the grey box shows, we are introducing a concept called weighted average speed to help predict future bandwidth needs.. While the top speed (brown box) is defined by the largest bonding group's size (depending on whether we use a 2:1 or 3:1 oversubscription model), it is the weighted average consumption that drives the total bandwidth needed. The weighted average is important as it represents the speed offering that proportionally impacts consumption. For example, in a representative spread of 10% of users taking the lowest speed tier, 60% in the middle tier, 20% in the next highest, and 10% in the highest tier, even a minimal increase in the middle tier will cause a higher overall data consumption rate than significantly increasing the highest tier based on the number of customers it is impacting.

A speed increase in the highest tier of customers has little impact in the overall consumption number since we typically have less than 1% of customers who take the highest speed tier. The majority of customers are in the middle tiers (often less than 20% of the highest speed tier offering). In these tiers, a small increase has a significant impact on the overall bandwidth consumption. This impact is even greater in the upstream where resources are even more limited.

The graph predicts that a 1 Gbps maximum speed will be offered in the next 5 years. To those who have been in the industry 10 years or more this may seem an impossibly high number, but we need to plan for this as a probable future need.

One Gbps top speed on cable? That's crazy! How do we get more capacity?

Good question. BAU (Business As Usual) node segmentation and splits add more capacity but network demands are forcing a node splitting rate that is unsupportable with current solution. If we split nodes as required we may simply run out of head-end space, power, cooling, and cabling necessary to support the new infrastructure.

For the next section we will discuss some options. We will focus on how to solve the multi-variable problem of capacity, power, rack space, and cabling.

CCAP Architectural Options Explored

As the authors discussed alternative CCAP deployment architectures, we realized that we needed a common language and way to describe the architectures.

Let's start with the basics. What does our outside plant look like today? From left to right in the graphic below, there is a metro ring interconnecting hub sites. The hub sites house routers, switches, video insertion equipment, edge QAMs, CMTS, CCAP, combiners and splitters, business equipment (e.g. metroE),metro Ethernet, and laser transmitters and receivers. On the right side of the hub, the upper connection shows a typical HFC deployment with fiber nodes, trunk amplifiers, feeder amplifiers, and line extenders. The lower connection shows a fiber based deployment for PON (Passive Optical Network), switched Ethernet or business services.

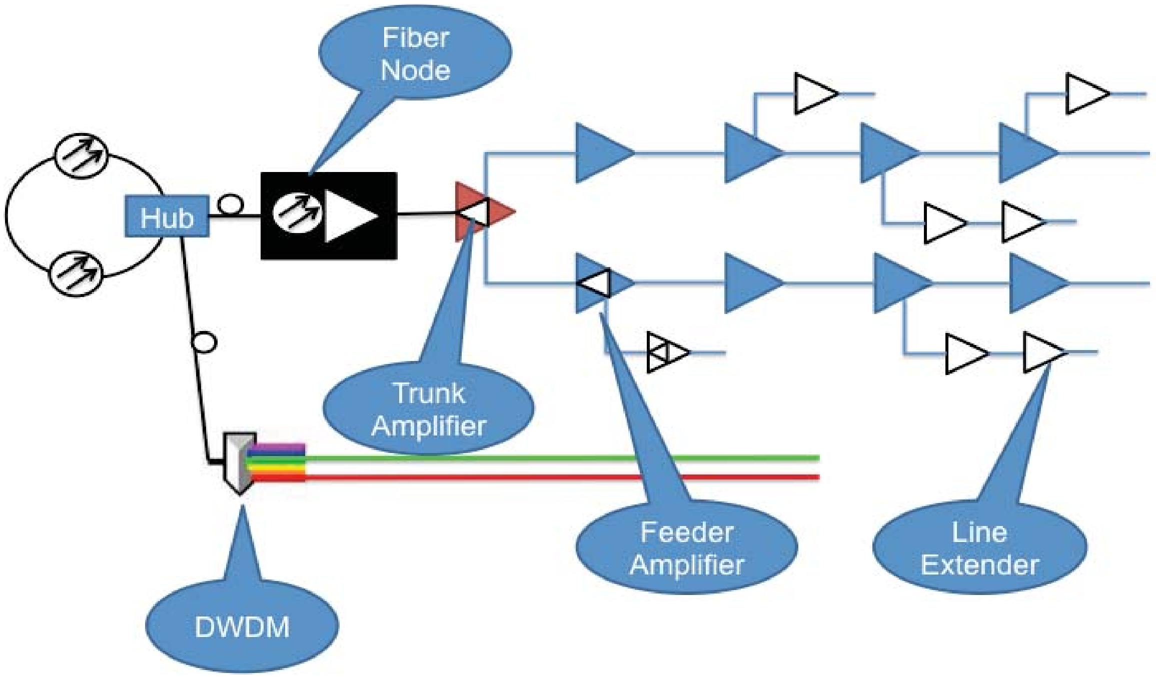

Figure 2 - HFC Architecture

Starting with this network, what happens as we drive to a fiber deep architecture in the HFC plant, but only use the fiber for greenfield deployments? As figure 3 shows, we need to add more fiber nodes in the plant to support smaller service group sizes, which requires more fiber be pulled to the new nodes.

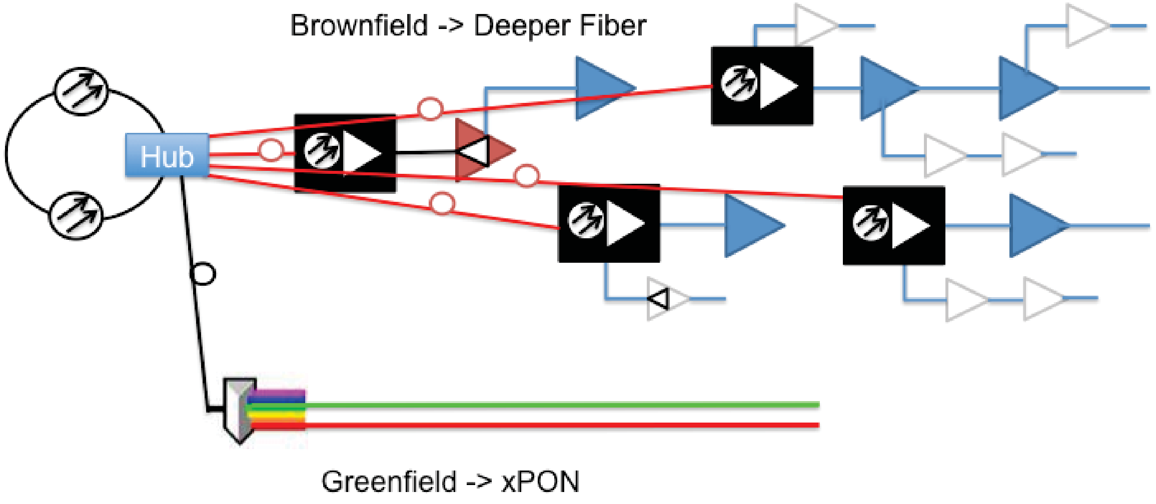

Figure 3 - Brownfield and Greenfield Deployments

The below chart discusses the good and bad news of the architecture changes with a fiber deep deployment topology.

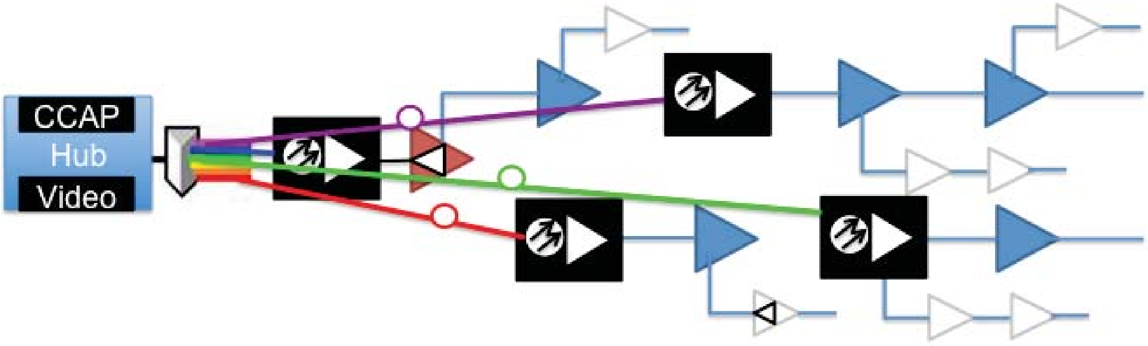

Good news
- Less equipment in the Head-End
- Multi-service single box
- Backward compatible
- Easy migration
- Simplify Head-End

Bad news
- Complex to build
- Continued node splits needed
- Timeframe requires transitional approach
- Head-End requirements will be difficult scaling at N+0 and possibly N+1

Figure 4 - Business as usual node splits with CCAP

We have added a few comments on good and bad news with this and the following topologies. The key take away is that when scaling from a N+6 to a N+0 architecture, we will need to continue splitting some percentage of nodes to passive coax, but that creates problems in the head-end from a space perspective.

We have discussed space numerous times, and it important to put it into perspective. The graphic below compares the equipment and space required in a legacy modular CMTS head-end versus a CCAP head-end.

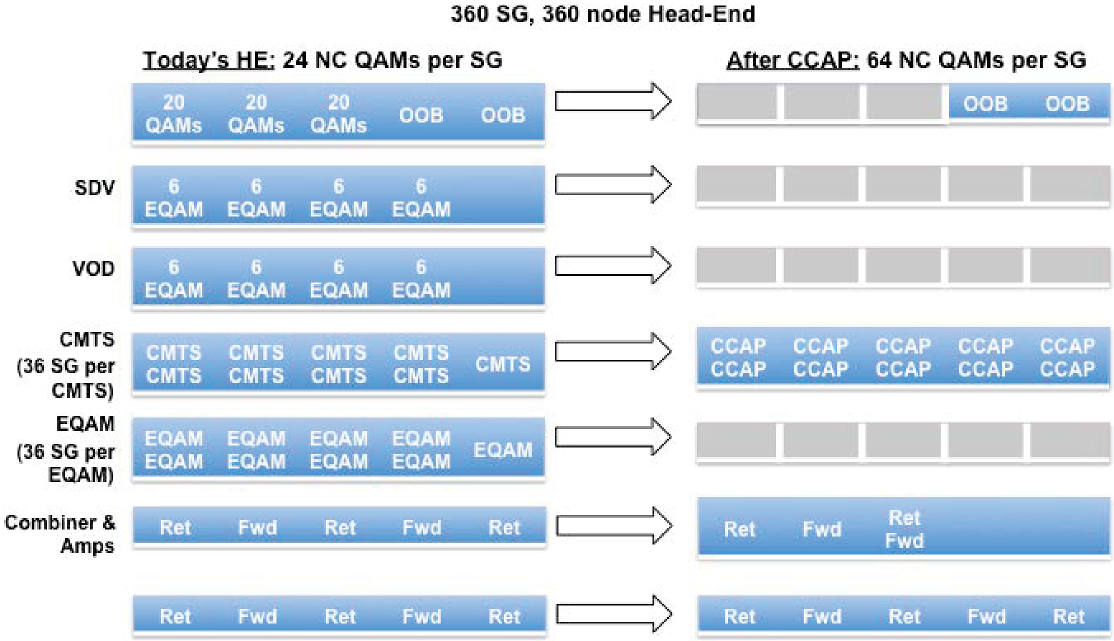

Figure 5 - Before/After of CCAP deployment

The key take-away here is we gain back 50% of the floor space and can even add 16 more QAMs per service group. This is the magic of CCAP as it integrates the edge QAM and CMTS into a single chassis. Note that this example requires integrating all data and video, including broadcast, into one CCAP platform.

Now, lets split the nodes connected to this CCAP hub from N+6 to N+0…

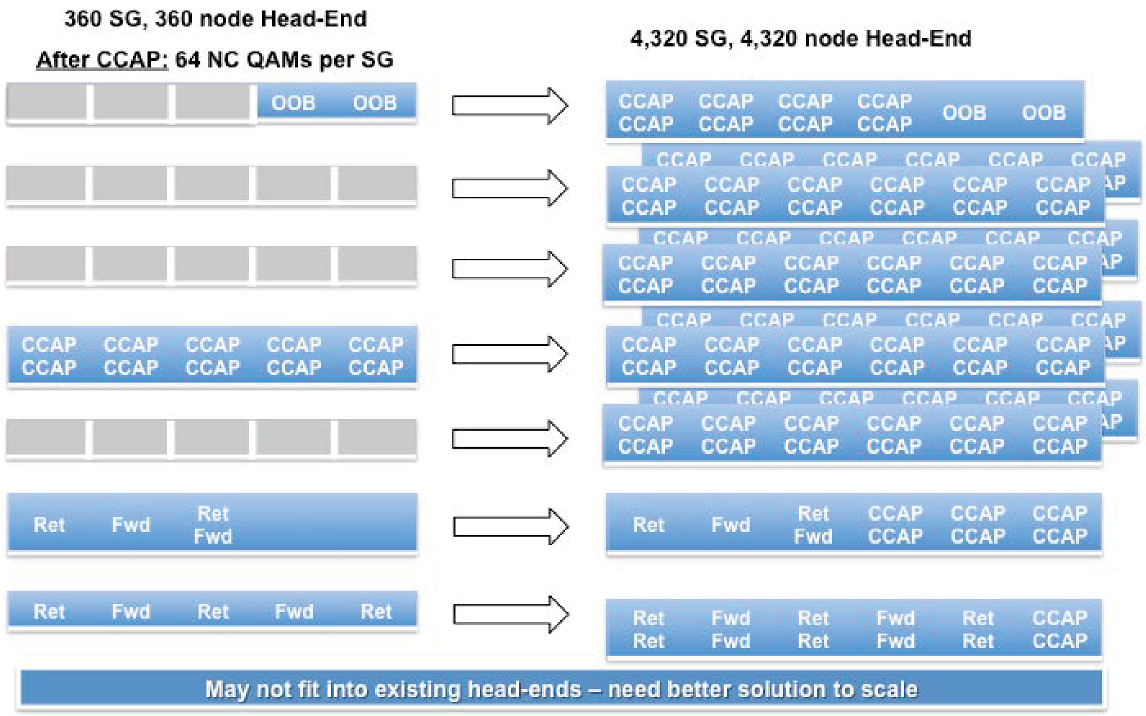

Figure 6 - From N+6 to N+0

As the illustration shows, you can see not only do we fill up the space we cleared, but we also need to add an additional 4 full rows with racks of CCAP equipment (assuming we are providing 40 service groups capacity from a single CCAP chassis and that we can put 2 CCAP devices in each rack). It is possible that there could be DWDM (Dense Wave Division Multiplexing) optics installed into the CCAP chassis, but that adds an additional level of complexity and may result in challenges in dealing with legacy equipment e.g. analog carriers, set-top out-of-band telemetry, and other equipment.

Given constraints in current facilities, we have a problem. It may not be there today and will not likely exist for years after CCAP has been deployed. The time to think about it is now if we want to have a solution when we need it.

Where are we going to put the necessary equipment required to solve our capacity and bandwidth problems of the future?

Remote PHY (R-PHY) and Remote CCAP (R-CCAP)

At this point let us re-introduce the concept for a remote PHY and also discuss a new concept for the remote CCAP…

Remote PHY (R-PHY)

From a simplistic view, here is a high-level view of remote PHY…

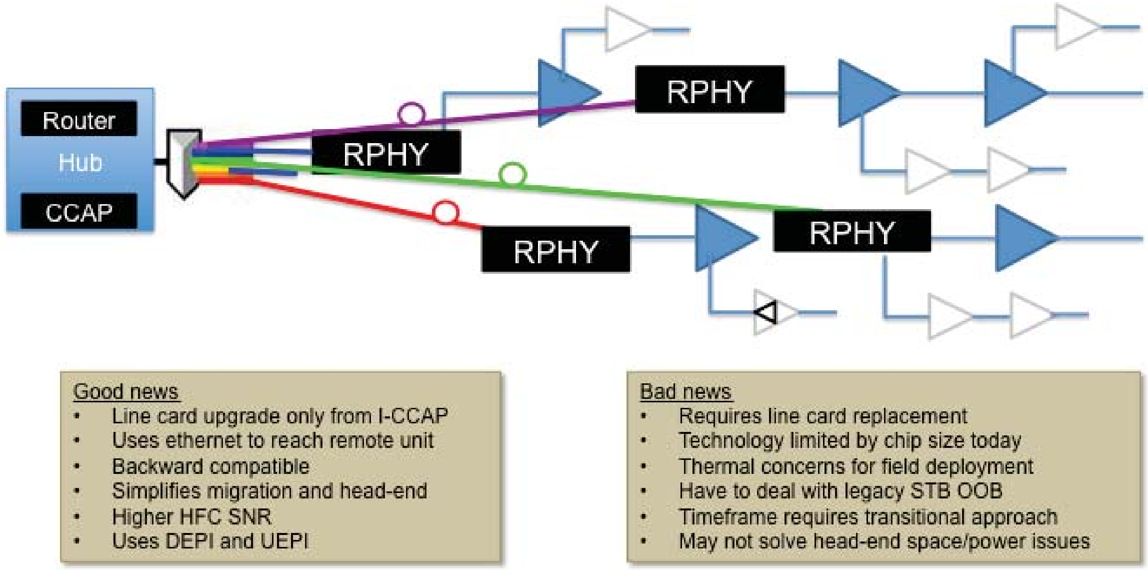

Good news
- Line card upgrade only from I-CCAP
- Uses ethernet to reach remote unit
- Backward compatible
- Simplifies migration and head-end
- Higher HFC SNR
- Uses DEPI and UEPI

Bad news
- Requires line card replacement
- Technology limited by chip size today
- Thermal concerns for field deployment
- Have to deal with legacy STB OOB
- Timeframe requires transitional approach
- May not solve head-end space/power issues

Figure 7 - Remote PHY Architecture

The PHY has been decoupled from the rest of the CCAP functional blocks enables the employment of advanced features. The idea is to locate the PHY (L1) deeper into the network, at the edge of fiber and coax plants, while PHY-MAC connection works through the digital optical link. This architecture creates conditions for delivery of higher-capacity QAM signals and improves overall network efficiency. Modulation of RF waves beyond the fiber prevents its damage by optical noises, which in normal conditions cause the SNR degradation in a network. This type of network architecture is named Remote PHY.

Remote PHY (R-PHY) CCAP represents a complete digital platform at the head-end with modulators/demodulators located on the other side of fiber at the node. The flow of data traffic and video traffic here is the same as for integrated CCAP (I-CCAP), the difference being the physical separation of the CCAP MAC and its associated PHY. The multicast and unicast streams go through the CCAP core, but instead of being modulated to QAM at the head-end, they are converted to an optical signal by standard digital optical transceivers and delivered through fiber to the remote PHY modulators, which returns the signal to its original electrical binary form and performs modulation. The output of Remote PHY is connected to a coaxial plant and combined with analog broadcast legacy services. In the return path, the same technique is used to demodulate the received signals and digitally transmit to the CCAP core over fiber using separate wavelength or physical fiber.

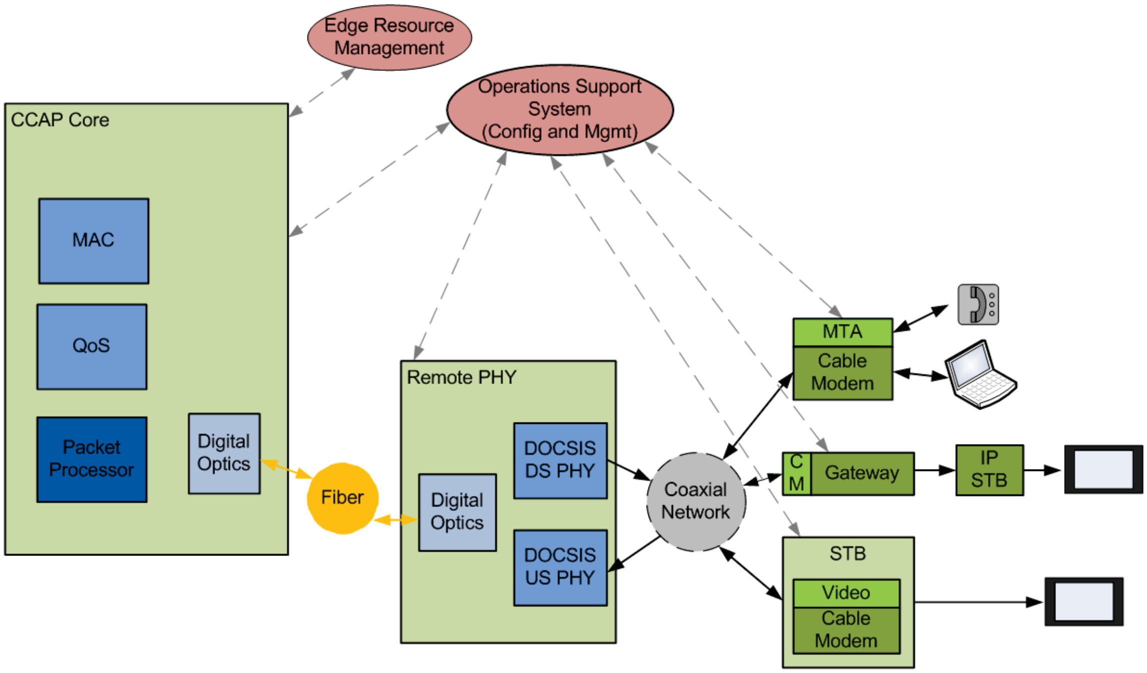

Figure 8 - Remote PHY Functional Architecture

Challenge 1 – Timing Synchronization

Synchronization of all elements is critical for the proper working of DOCSIS system. The modulator and demodulator must be synchronized to a Master 10.24 MHz clock, which is located in the R-PHY CCAP core system. In order to perform it, the fiber data clock has to be synchronized to the Master 10.24 MHz at R-PHY CCAP's optical transmitter. Clock Recovery circuit will derive the original 10.24 MHz clock from data input stream at PHY's receiver side. Usually a clock recovery scheme such as a Phase Locked Loop is used, which locks an input data stream. For clock recovery to work properly, the data stream has to change bits polarity consistently that can be guaranteed by scrambler, 8b/10b or 64b/66b encoders. Figure 9 demonstrates a classical scheme for clock recovery based on the PLL (Phase Locked-Loop) approach. If the clock jitter on an output of a standard optical receiver will not satisfy the tough requirements of DOCSIS, then an additional Clock Filter circuit should be added for jitter mitigation.

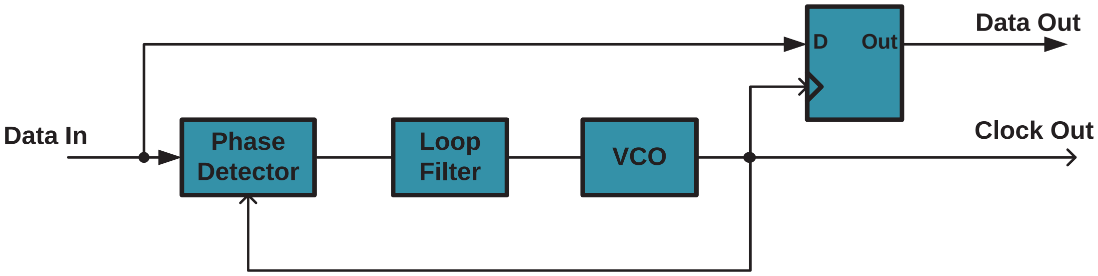

Figure 9 - Clock Recovery Scheme

Additionally, the PHY and MAC blocks must work in the same time domain. Therefore, Timestamp Counters on both sides of fiber has to be synchronized and must compensate for the fiber delay at the decoupled PHY. Figure 10 shows the block diagram for Timing Synchronization of R-PHY concept.

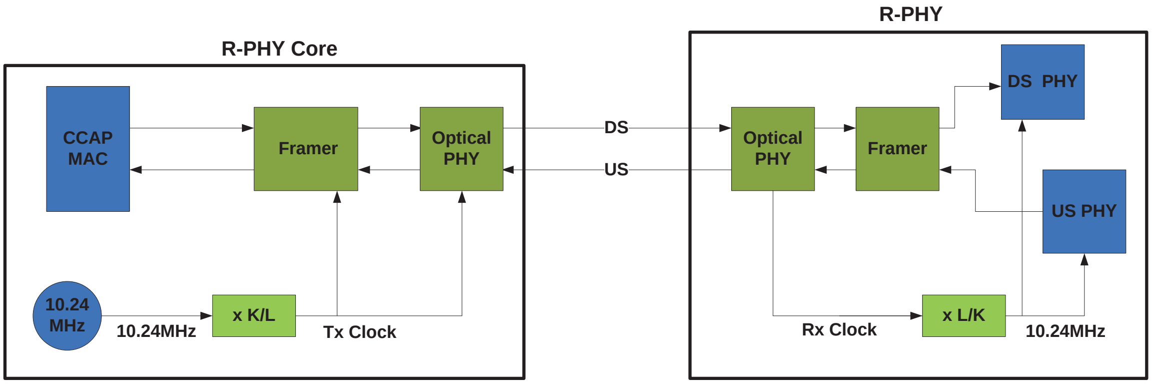

Figure 10 - Time Domain Synchronization

<u>Challenge 2 – MAC / PHY Protocol</u>

The separation between MAC and PHY functionalities will require a definition of a special protocol for communication between two logical blocks. DOCSIS 3.0 introduced the Downstream External PHY Interface (DEPI) specification for Modular CMTS architecture where DOCSIS downstream PHY was located externally to the core CMTS at edgeQAM device. However, additional discussions are required in order to conclude if DEPI can work for Remote PHY concept.

The Upstream External PHY Interface (UEPI) specification for separation of upstream PHY has not standardized yet. Initial work was done during the DOCSIS 3.0 effort and has the potential to dramatically shorten the amount of time needed to draft a spec.

Challenge 3 – Configuration, Monitoring and Firmware Update

The Remote PHY concept declared decoupling of Layer 1 to a remote location. Layer 1 or Physical Layer defines electrical and physical specifications commonly implemented by hardware elements that should not request any firmware updates. But in the reality DOCSIS PHY (DS, US or both) can (and does) use an FPGA (Field Programmable Gate Array) for implementation of FEC and modulation/demodulation. There will be circumstances when the FPGA will need to receive a new code version from the Central Office to ensure proper functionality.

The dynamic PHY configuration and network monitoring information has to be delivered to/from the DOCSIS MAC at head-end for proper network management. As part of the CCAP spec, the YANG model has been defined as the management interface to the system, which will scale well with remote architectures.

The MAC/PHY protocol has to take in account the delivery of management traffic between Central Office, DOCSIS MAC and DOCSIS PHY. As the management plane is now part of the data plane, this needs to be accounted for in traffic planning and possible use of a differentiated service to assure delivery of configuration and reporting data.

Lots of complexity involved in something that seems so simple. It also does not directly solve our original problem with CCAP, which is we are running out of space in the head-end as we reduce our cascades of amplifiers.

Now lets look at another idea for distributing CCAP functionality into the plant…

Remote CCAP

Again, from a simplistic view remote CCAP looks similar to remote PHY, except that all MAC and PHY is removed from the head-end…

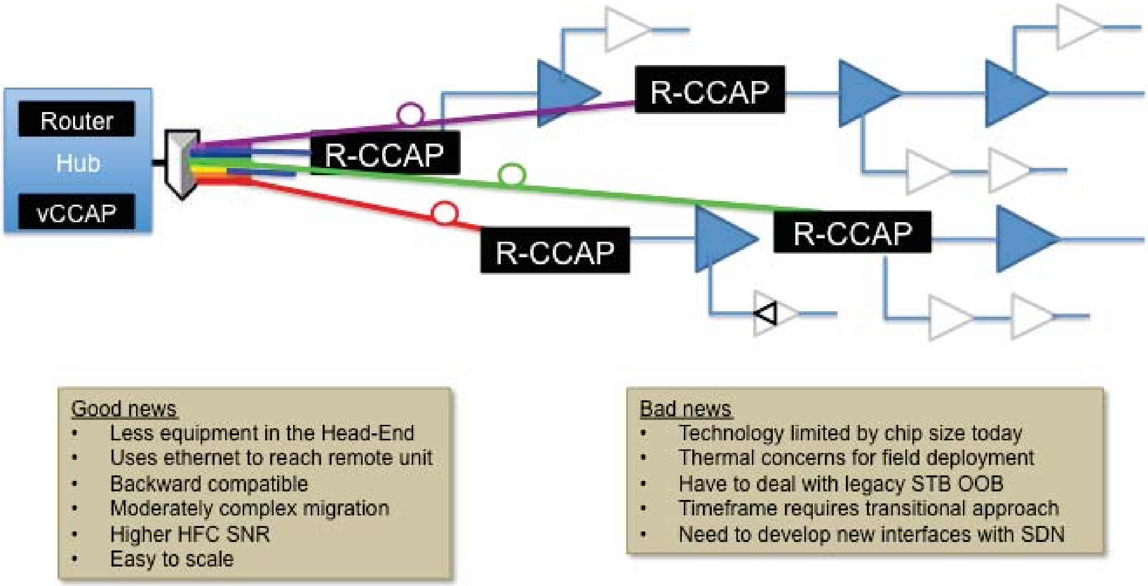

Good news
- Less equipment in the Head-End
- Uses ethernet to reach remote unit
- Backward compatible
- Moderately complex migration
- Higher HFC SNR
- Easy to scale

Bad news
- Technology limited by chip size today
- Thermal concerns for field deployment
- Have to deal with legacy STB OOB
- Timeframe requires transitional approach
- Need to develop new interfaces with SDN

Figure 11 - Remote CCAP Architecture

The concept of remote CCAP is to relocate all layer 1 (PHY) and layer 2 (MAC) functionality into the HFC plant, which enables further network and head-end transformation to a single-protocol IP system using Ethernet for all communications. This includes data, control and management plane traffic.

The remote CCAP concept converts a node to a standalone device, located deep in the network. It takes a role of media converter between fiber and coax for provisioning IP and MPEG services over HFC. From the fiber side, this is another element of Ethernet network, while from the cable side this is a network manager for provisioning of high quality video and data services. Figure 12 shows the remote CCAP architecture.

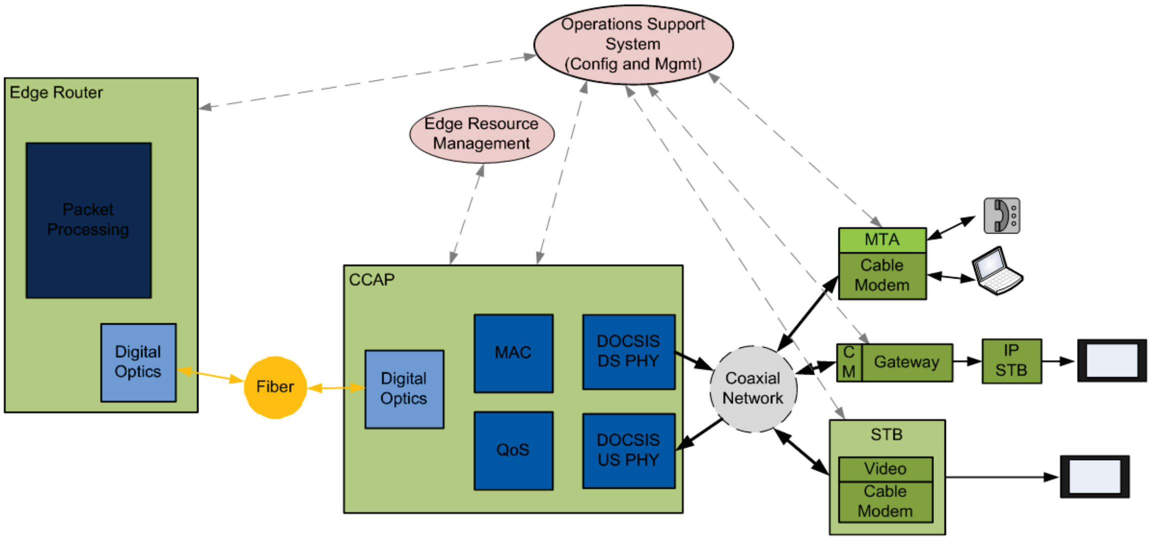

Figure 12 - Functional Architecture of remote CCAP

<u>Complex or Simple?</u>

This was a significant question the authors had as we discussed this technology. As we went through various implementation details we realized that if we would take the approach of virtualizing the majority of cognitive upper layer functions and used SDN to manage them through virtual hosts, the technology became easier to manage.

This is not meant to be a tutorial on SDN (Software Defined Networking), but as a brief introduction it is composed of open standards including OpenFlow and OpenStack, along with a number of IETF and IEEE evolving standards such as PCE (Path Computation Element) and others. For more information on these standards there are a number of excellent web sites or feel free to contact the authors.

The following are a number of questions we had as we thought through the process of deployment and management of these devices. The following diagram shows the virtualization of functions with a remote PHY, followed by one for a remote CCAP.

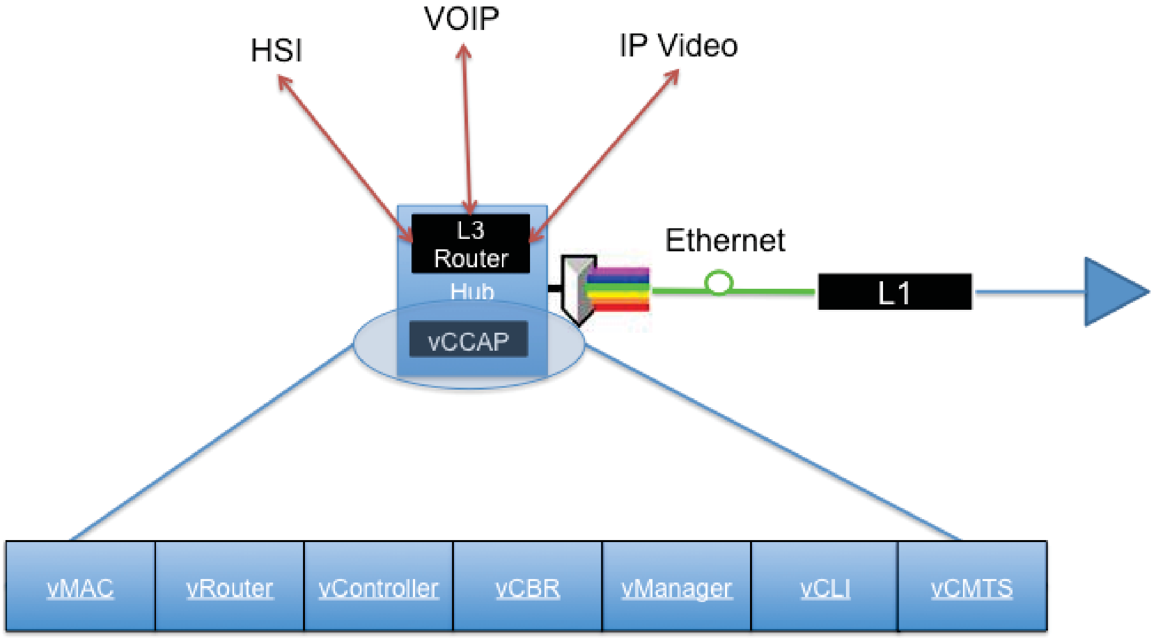

Figure 13 - Virtual CCAP with Remote PHY

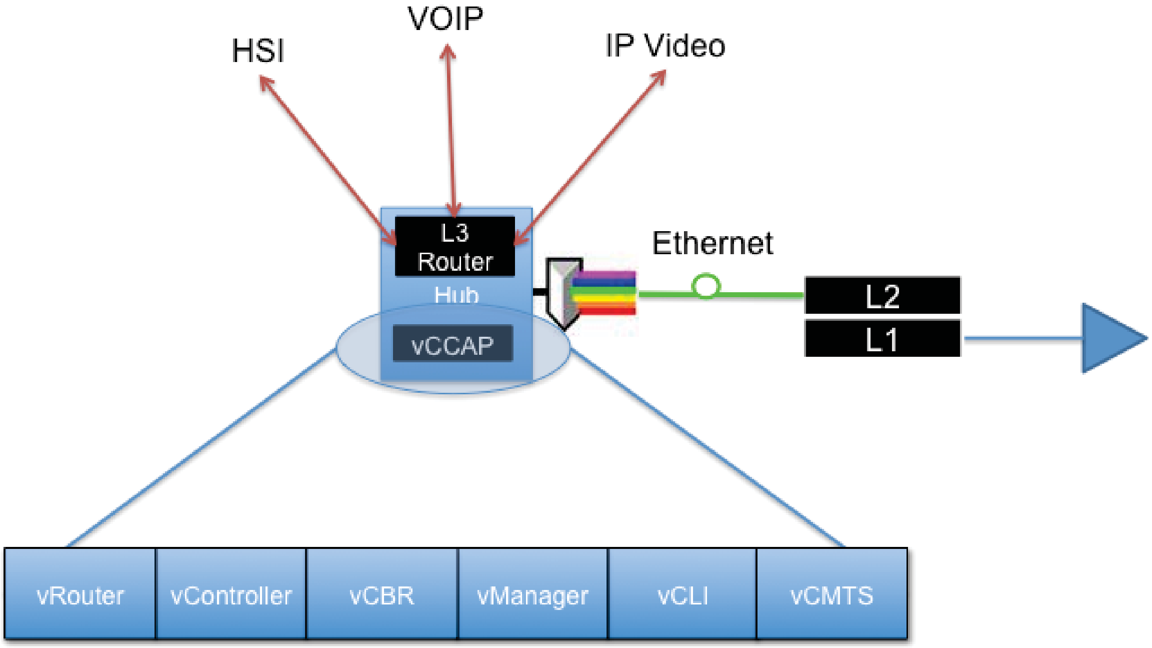

Figure 14 - Virtual CCAP with Remote CCAP

As we show the only difference between the virtualized functions is that with the Remote CCAP the MAC and PHY is decoupled from the physical components in the hub where with the remote PHY the MAC is still part of the CMTS core functionality in the hub and only the PHY has been decoupled.

Challenge 1 – Routing

Is the device going to be a router or switch? Today these devices are routers, which makes sense to us as you have larger systems with tens of thousands of users. We realized that managing routing resources for 500 or few users at the device was overly complex. Our basic premise with remote CCAP is that these devices are simple media convertors. Placing routing in the existing routing subsystem simplifies the remote device and helps reduce its power requirements. The connection from the head-end to the remote device is simple. It is Ethernet, nothing more, and nothing less. As one of our working mottos states, “Keep It Simple, Keep It Ethernet.”

Challenge 2 – Video Controller

As we worked through existing functionality the one we had to think through most is to consider video delivery. Today we have complex edgeQAM devices, many with embedded encryption. Our question was do we need to maintain this long-term or is it primarily for our existing device technologies?

We recognize that we will have legacy set-top box technologies deployed for many years. IP delivered video to the home will use DRM for protecting content, which removes complexity from the edge devices. How do we deal with this during the interim years?

Our decision was to push encryption out to bulk encryption, which simplifies the edge but does require pushing some complexity into the core. We felt that the trade-off in this regard dramatically simplifies the remote device making it worth the shift in how some may be currently encrypting legacy content.

Challenge 3 – Content Based Routing

Content based routing is a new concept that defines a content-based network as a communication network that features a new advanced communication model where messages are not given explicit destination addresses, and where the destinations of a message are determined by matching the content of the message against selection predicates declared by nodes. Routing in a content-based network amounts to propagating predicates and the necessary topological information in order to maintain loop-free and possibly minimal forwarding paths for messages. The routing scheme we propose uses a combination of a traditional broadcast protocol and a content- based routing protocol.

By virtualizing this functionality we further simplify the remote device by pushing the requests to the layer 3 devices, enabling us to keep the communication path to the remote device as a simple layer 2 connection.

Challenge 4 – Management

When going from hundreds to tens of thousands devices, how do we continue to manage them? Today we use CLI to manage CMTS and CCAP, but how will that scale?

One idea we have been discussing is virtualizing management functions using standard protocols. The concept is that a device maintains no state information; everything is kept in a distributed database that is shared between all the virtual management nodes in the network. It is similar to the stateful database kept on every device today in terms of data, but instead of being stored locally it is maintained in the network itself.

When the device boots it issues a boot request into the SDN network, which triggers a download of the latest configuration items to the device. As changes are made in virtual management hosts and sent to the device when committed, no persistent state is necessary on the remote CCAP device.

The same goes for the virtual CLI. As there is only a minimal layer 3 configuration at the remote device, the interface into the remote CCAP is provided through a virtual connection that mimics a CLI or HTTP interface as needed.

One big question that we had was the impact of these remote technologies on node power. Both solutions must meet the power requirements of the existing fiber nodes.

Challenge 5 – Virtual MAC

The only difference between a deployment for the remote PHY and remote CCAP is where the MAC components exist in the network. With the remote PHY the MAC (packet scheduling, registration, service flow creation, etc) remains in the head-end or other remote location. The communication between the MAC and PHY is standard Ethernet, so the interconnect to the remote device does not change.

In a remote CCAP deployment the MAC is co-located with the PHY.

Challenge 6 - Power Consumption

For the estimation of power consumption of R-PHY and R-CCAP, the system components can be reviewed and appraised individually. Not all of components are parts of both concepts, but differentiation can help to compare both systems.

- RF

The power consumption of RF circuits is a function of requirements to output power level, SNR and insertion loss of passive elements. Higher requirements will cause higher power consumption. Any system located in the node will have to include RF circuit for downstream and upstream signals.

- PHY

PHY is a heavy consumer of power because it performs a mass of DSP and FEC processing. The digital filters and FFT have to be a high resolution in order to generate a high performance signal. Digital up-converters have to convert signal from low to a very high frequency and locate it everywhere on a wide DS spectrum. The channel estimation algorithms have to perform powerful mathematical calculations. And future spectrum and complexity upgrades will cause to a higher power demand.

- MAC

MAC provides a management layer for the cable network. Despite the importance of this layer and heavy combinatorial logic, the resolution of processing is not higher than packet length and so the power consumption of this layer is not a dominant factor.

- QoS

Quality of Service guarantees a bandwidth allocation and latency for DOCSIS Service Flows, packet classification and dynamic service establishment. It has a heavy hardware for queuing, but the processing resolution is equivalent to a packet length and the power consumption is close to the MAC layer.

- Network Processing (NP)

Packet Processor also works with packets, but has to provide big memory arrays to build MAC/IP tables for data forwarding. The heavy search algorithms must be implemented to provide efficient results.

- Digital Optical Transceivers (DO)

Digital Optical Transceiver connects a node to a digital fiber network. The power consumption of this component depends mainly on optical speed, but the protocol is also important. For example, in the case of PON, ONU’s MAC will slightly increase the power demand.

An existing DOCSIS system based on previous generations of ASIC technology was analyzed to provide estimation of power consumption for every node module. The Figure 15 shows the power spread between the different system blocks.

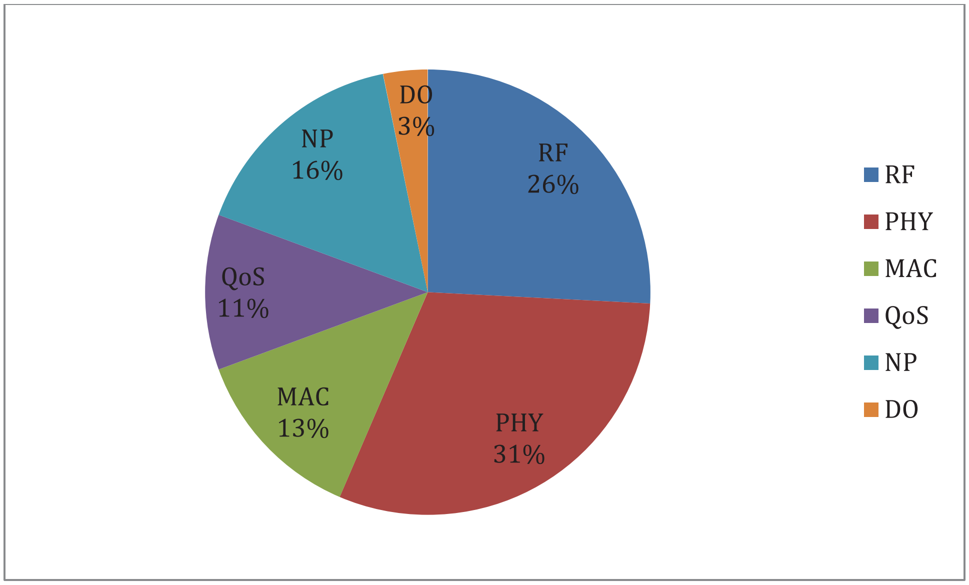

Figure 15 - Existing Technology Power Consumption Spread

As we know from Moore’s law, transistor density is doubling approximately every two years, which results in a significant power reduction per logic element. However, the power reduction is not always directly related to the size of a transistor’s area and can take lower proportions. With a two-fold increase in density, power will go down by approximately 30%. Based on that we can try to predict the future power consumption of the same modules by taking in account the above approximation. Figure 16 shows reduced power consumption for all digital parts, while the weighting of power for RF and Optics components increased.

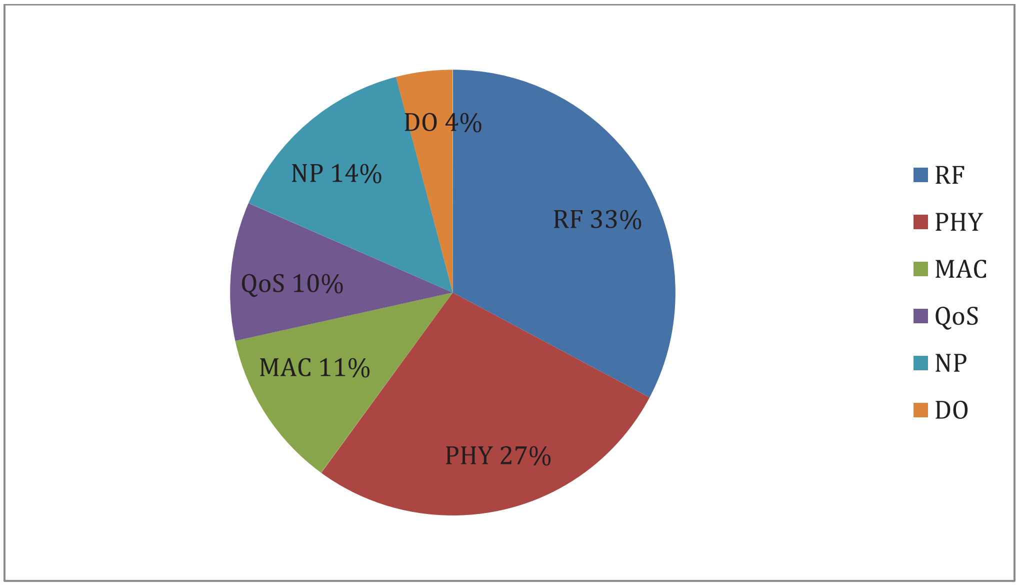

Figure 16 - An Approximation of DOCSIS System Power Spread in future years

The Figure 17 shows the tendency of power consumption for the next few years.

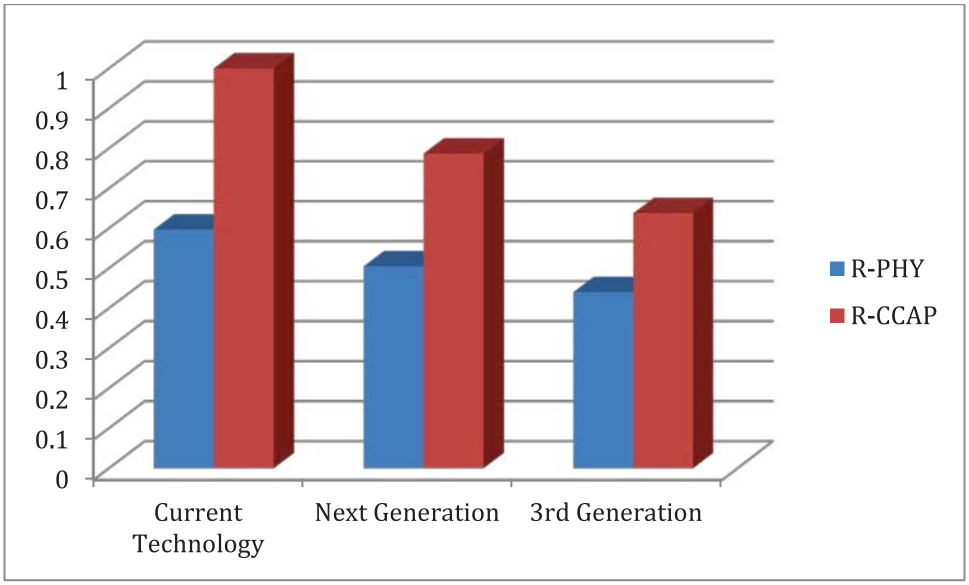

Figure 17 - Relative Power Consumption of R-PHY and R-CCAP

The other part of the power equation is cooling. The R-PHY CCAP architecture still requires a CCAP device in the hub, which dissipates significant amount of heat. Air-conditioning requirements must be adjusted to accommodate this, which will increase the overall power requirements for the hub. The Remote CCAP architecture removes the physical CCAP device from the hub, thereby relaxing overall air-conditioning requirements in the hub and lowering hub power.

Challenge 7 – Coexistence with Legacy Services

R-CCAP and R-PHY can replace a traditional optical node in future network splits, but what can be done to support legacy broadcast services and SCTE-55 (OOB) signaling?

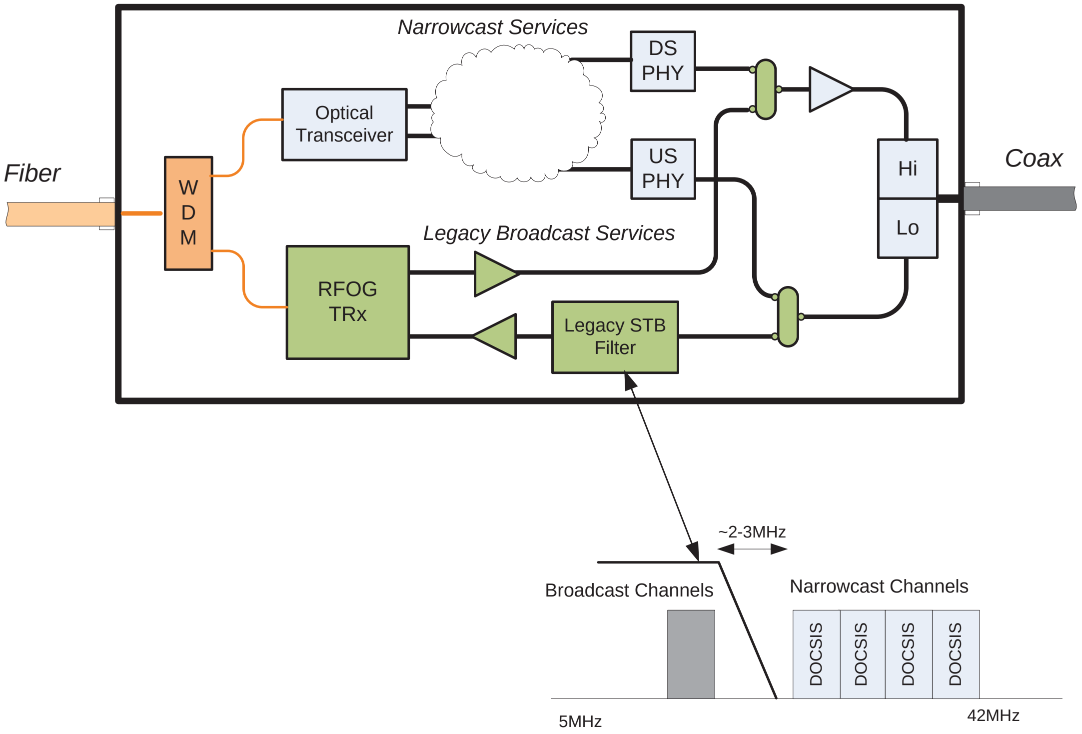

Figure 18 - Possible Node Architecture for Legacy Services Support

Existing coax cable can be used for delivery of legacy US OOB signal to the headend. A low pass filter can isolate between US service groups. On the DS side, a simple filter can't be efficient enough for many cases, in which case a fiber for legacy should be extended up to the new node.

Another method can be applied to simplify the system architecture. The digital return (spectrum digitization) technique can make possible using of the same fiber and protocol for traffic delivery to/from the node (including legacy). The digitized legacy spectrum can be framed and transmitted over fiber and recovered by the match devices.

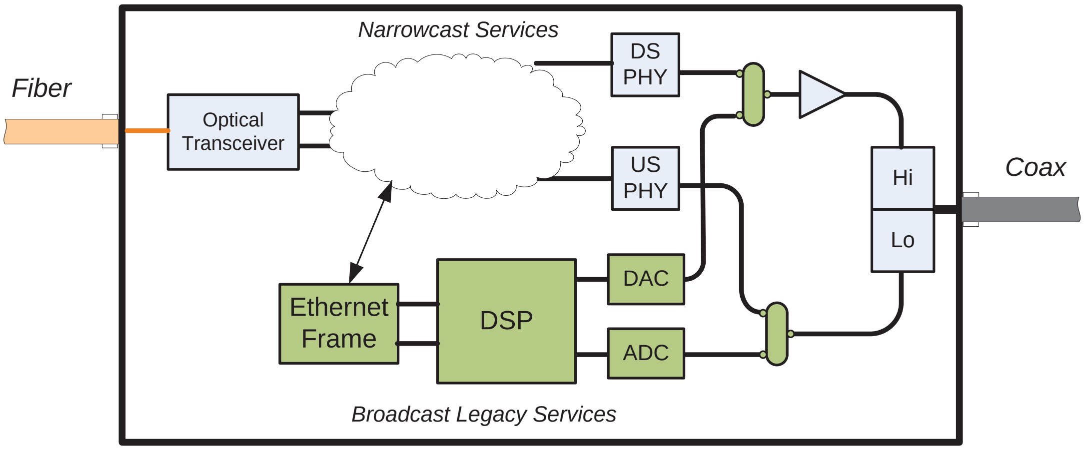

Figure 19 - Legacy Support Using in DSP Technique

Conclusion

HFC plants have a significant latent capacity upwards of 10 Gbps, but the usable bandwidth is limited by amplifiers in cascade, linear and non-linear distortions, diplex filter roll-off, and many other factors including the drop cable and in-home network. Given enough transmit power many of the impairments related to attenuation can be overcome, but there is a finite limit to how much channel capacity may be transmitted at any time based on the Shannon limit.

By distributing the PHY closer to the customer many fiber, coaxial, and amplifier impairments are lessened as the factors that impact the upper limit of transmitting or receiving speed. While there is nothing stopping a remote PHY or remote CCAP technology from working through multiple amplifiers, the most effective use of it is to drive us closer towards a passive coaxial architecture (node+0). By removing amplifiers from the network we are able to take advantage of all an HFC plant has to offer from a bandwidth, capacity, and transmission speed perspective.

Pushing the MAC to the plant adds some additional complexity, but also frees up the head-end from having scalability issues as related to rack space, power, cooling, and cabling. A question we need to answer is do the head-end savings outweigh the concerns in adding this additional complexity into the field.

We may deploy remote PHY and remote CCAP in a traditional HFC network with multiple amplifiers in cascade and migrate over time to a passive coaxial configuration, which allows us to take advantage of driving fiber deeper with all the benefits of having the modulators/demodulators closer to the customer.

From the analysis and discussions as we worked through our options, it was recognized that our current and near-term architectures would not easily scale to what was needed in order to compete in the future. We needed a new network paradigm allowing us growth to the scale necessary for our future needs and that the technologies proposed get us to an architecture that will meet that requirement.

Abbreviations

ASIC	Application Specific Integrated Circuitry
CCAP	Converged Cable Access Platform
CM	Cable Modem
CMTS	Cable Modem Termination System
CWDM	Coarse Wave Division Multiplexing
DF	Deep Fiber

DOCSIS	Data Over Cable System Interface Specification
DS	Downstream
DVR	Digital Video Recorder
DWDM	Dense Wave Division Multiplexing
EPON	Ethernet Passive Optical Network
FPGA	Field Programmable Gate Array
GbE	Gigabit Ethernet
Gbps	Gigabits per second
GPON	Gigabit Passive Optical Network
GEPON	Gigabit Ethernet Passive Optical Network
XGPON	10 Gigabit Per Second Passive Optical Network
HD	High Definition
HFC	Hybrid Fiber-Coax
IP	Internet Protocol
I-CCAP	Integrated Converged Cable Access Platform
I-CMTS	Integrated Cable Modem Termination System
MAC	Media Access Control
Mbps	Megabits per second
MHz	Megahertz
nDVR	Network Digital Video Recorder
OFDM	Orthogonal Frequency Division Multiplexing
OLT	Optical Line Terminal
OOB	Out Of Band
PCE	Path Computation Element
PHY	Physical Layer
PLL	Phase Locked-Loop
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
R-CCAP	Remote Converged Cable Access Platform
R-PHY	Remote PHY
SD	Standard Definition
SDN	Software Defined Networking
SDV	Switched Digital Video
SNR	Signal to Noise Ratio
STB	Set-top Box
US	Upstream
vCCAP	Virtual Converged Cable Access Platform
VOD	Video on Demand
WDM	Wave Division Multiplexing