HFC NETWORK CAPACITY EXPANSION OPTIONS

Jorge D. Salinger
VP, Access Architecture
Comcast Cable

Abstract

MSOs are deploying more narrowcast capacity than ever before, and there is no evidence of a change in this trend.

- *DOCSIS® 3.0 is widely deployed, with 4 and 8 downstream channel bonding groups becoming the norm. A continual annual growth of 40-60%, observed industry-wide for over 10 years, would require many more channels over time*
- *8-channel service groups for video on-demand (VOD) are commonplace. Growth rate increasing due to both higher usage and higher bitrate (high definition)*
- *10, 20 or even more channels for switched digital video (SDV) are frequently used for longer tail content*
- *Growth in business service applications requires additional increased capacity*
- *The advent of IP video services and network-based digital video recorder (DVR), which are anticipated to be very popular amongst current and potential subscribers, will compound the need for additional narrowcast capacity.*

The effect of the above trends, combined with the need to simultaneously support a full set of legacy broadcast services, including digital, analog and/or both, would likely require additional hybrid fiber-coax (HFC) network capacity.

While it seems conceivable that a transition from legacy and broadcast services to an all-narrowcast/IP services infrastructure could be established, the industry as a whole is looking for options that would provide additional capacity to support simultaneous uses, and increased capacity beyond such transition. These options include:

A. *Traditional service group segmentation*
B. *Move quadrature amplitude modulation (QAM) generation downstream into the network*
C. *Implement higher modulation physical layer (PHY) and/or more efficient media access control (MAC) protocols*
D. *Increase HFC downstream capacity beyond currently deployed, and/or move split to higher spectrum for increased upstream capacity*
E. *Develop technology that would operate in unused portions of the spectrum, and even unleash spectrum above current top range (e.g., above 1 GHz)*

Each of the above options has benefits and drawbacks. Each approach offers different network engineering and operational simplifications and complexities. And, the relative improvements in offered capacity versus cost and customer impact can be significantly different.

This paper will provide, from an operator's perspective:

1. *A technical overview of each of the options outlined above, describing how each would be deployed and evolved over time,*
2. *The key benefits/drawbacks for each of the options, including engineering and operational pros and cons for each option,*
3. *Possible implementation approaches for various applications, including residential and commercial services.*

TYPICAL HFC NETWORKS TODAY

Most MSO's hybrid fiber-coax (HFC) networks have been designed to either 750 or 860 MHz of spectrum capacity. If not fully utilized, it is expected that use of their capacity will be increased to the point of exhaustion as the use of DOCSIS® increases for the higher high-speed data (HSD) service tiers, additional high-definition (HD) programs for both broadcast (BC) and especially narrowcast (NC) services such as video on demand (VOD) and switched digital video (SDV) are deployed, or new services such as internet protocol (IP) video and network-based digital video recorder (n-DVR) are added.

Proportionally few HFC networks have been deployed to operate up to 1 GHz, although all equipment available today can support the use of spectrum up to 1 GHz and even 3 GHz for some components.

In recent years the growth in, and demand for, HD programming has resulted in the need for allocation of large numbers of EIA channels for HD services, both for BC and NC, which has filled every available portion of the spectrum. This is especially true for BC, where large numbers of programs are offered in HD format, while simultaneously the need for distributing the standard definition (SD) version has persisted. This has resulted in the need for use of 3x to 5x the number of EIA channels than previously required. For example, a typical digital multiplex including 10 to 15 programs would require an additional 3 to 5 EIA channels for the HD equivalent streams, even assuming the newer, more sophisticated multiplexing schemes available in the market. Of course not every program is available, or still sought by subscribers, in HD format. But very large numbers of them are, including 100 to 150 BC programs.

The above is also applicable to a great extent in systems utilizing SDV technology for distribution of its content. The difference is that the SD version of the program is not distributed unless a subscriber is requesting it, which reduces the marginal increase in capacity. Assuming that all programs are distributed in only one format, which is certainly a valid expectation for programs of low viewership, then the increase in capacity for a conversion from SD to HD would just be the increase in capacity required for the transmission of the HD program without requiring the simultaneous use of bandwidth for both formats.

Additionally, considerable spectrum is needed to deploy high-capacity narrowcast legacy video services, especially n-DVR, and a full-array of HD video-on-demand services. For the former, initial observations suggest that network requirements for n-DVR may be as high as 4x to 5x that of VOD, and that peak utilization overlaps, at least partially, with that of peak use for other narrowcast services.

Finally, the growth in HSD services shows no sign of letting up. Network operators have observed an increase use of HSD service capacity for well over a decade now, which amounts to a year-over-year compounded growth of 40% to 60%. The applications have changed throughout this time, but the demand has continued to increase at the same relentless rate.

In fact, such increase in demand for HSD capacity shows no evidence of decreasing. Should that trend continue, MSOs would be in a position to increase access network capacity through either one or more of the existing capacity tools and/or through one or more of the new capacity tools outlined in this paper.

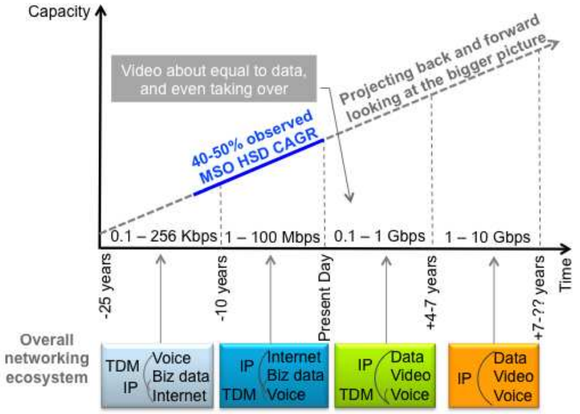

Figure 1: Example of HFC capacity utilization over time

How does this compare to other operator's data services and a longer period? As shown in Figure 1, projecting the MSO's HSD service growth back in time to when Internet services started as shown in the diagram, 25 years ago services should have been about 100 bps. This coincides with the history of telephone modems from 110 and 300 baud modems from the mid-80s, to 56 Kbps/V.42, into ISDN services.

This demonstrates that the growth seen in MSO's HSD services is typical over a much longer period of time, rather than an exception observed by MSOs in recent years.

GROWTH PROJECTIONS

From all of the above, it then follows that, should the usage growth pattern continue at the past experienced pace, networks will be required to provide HSD services in the range approximating 1 Gbps within the next few years. This growth, coupled with the surge in HD video formats, and more personalized narrowcast services, will result in a significant growth in NC capacity, as shown in Figure 2 below.

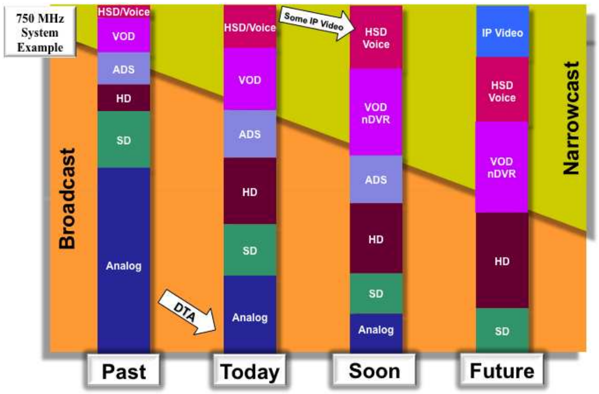

Figure 2: Example of narrowcast service growth over time

To support this growth, MSOs have deployed, or are considering deployment of, bandwidth reclamation tools such as SDV for digital broadcast, digital terminal adapters (DTAs) for analog services, or a combination of both. These tools have been extremely valuable to MSOs, which have seen their operational complexity and cost to be well justified.

In the case of SDV, early predictions several years back from industry analysts projected that the efficiency of SDV would reach 40% (e.g., programs requiring 10 EIA channels could be carried in 6). This has proven to be understated, since it was based on the use of SDV for reduction in bandwidth required for existing services. As SDV's role in the network grew, the efficiencies have been even greater, especially as SDV has been used to introduce niche services that have low viewership and would have otherwise been difficult to deploy.

The benefit of DTAs has been just as, or perhaps even more, striking. MSOs deploying DTA devices are able to eliminate the need to distribute the analog channels in the network. Even if DTAs are distributed to top analog tier customers, such as only to subscribers of the traditional expanded basic subscribers, such deployment would reduce a channel line up from perhaps 50 EIA channels dedicated to 50 analog programs to perhaps as little as 4 EIA channels dedicated to transport the 50 programs in their equivalent digital transport. Using the same comparison method as the above SDV case, this is a >90% efficiency. If extended to the entire analog tier the efficiency gains are very significant.

Despite the availability of these tools, they are not universally applicable. With respect to SDV, in general it is not likely that all broadcast programs will be switched since experience shows that many broadcast programs are constantly viewed by someone in the service group during peak hours, which will leave a large portion of the spectrum still used for broadcast. Similarly, not all analog channels can be removed in the short term due to operational and/or cost constraints.

Additionally, while many MSOs will use one or both tools, in general these tools won’t be used by every MSO for all applications. Finally, there are also significant potential gains to be achieved from the use of advanced video CODECs (AVCs) and variable bit-rate (VBR). In the case of AVCs, coding efficiencies of approximately 50%, depending on implementation and content type, can be obtain with H.264[1] and/or MPEG-4 Part 10[2]. And the use of VBR promises to result in a capacity efficiency gain of as much as 70% versus CBR[3]. The combined gains from using both approaches could be very significant.

However, these are difficult tools to take advantage on the network since proportionally relatively few legacy set-tops still support AVCs and VBR, especially the latter. These tools will likely enjoy significant support in newer, IP-video based services equipment moving forward.

But, this approach will require additional capacity on the network. This is especially true when considering that the deployment of these advanced video services will result in an additional simulcast of video programs, at least initially, which is expected since its deployment will not at least initially replace the currently deployed services. Furthermore, ubiquitous support for such devices would require considerable spectrum if the legacy services are maintained for an extended period, as it is expected since legacy devices are and will continue to be deployed. Moreover, this increase in simultaneous use of advanced, IP video services while maintaining legacy services will be especially impacting over time as its penetration increases.

All of the above, coupled with the success experienced by MSOs in recent with business services, will likely require the deployment of IP capacity beyond what can be supported

---

[1] ITU-T Recommendation H.264: 2005, Advanced Video Coding for generic audio-visual services

[2] ISO/IEC 14496-10: 2005, Information technology – Coding of audio-visual objects – Part 10: Advanced Video Coding

[3] Capacity, Admission Control, and Variability of VBR Flows, CableLabs Winter Conference, February, 2009

today, requiring the development of tools for increased efficiency in the use of spectrum and/or unlashing of additional spectrum in the HFC network. The following sections of this paper will enumerate ways in which this can be achieved.

OPTIONS BEING CONSIDERED

Let us review the categories of options being considered throughout the industry, and evaluate how each one fulfills the above desirable targets. In the process, let us review the key implementation aspects of each option, leaving for another opportunity the details of the options and on how these could be deployed.

The categories of options are:

1. Traditional service group segmentation

2. Move QAM generation downstream into the network

3. Implement PHY and MAC improvements

4. Increase HFC downstream capacity beyond currently deployed, and/or move split to higher spectrum for increased upstream capacity

5. Develop technology that would operate in unused portions of the spectrum, and even unleash spectrum above current top range (e.g., above 1 GHz)

1. Traditional service group segmentation

This option is readily available and has been in use for many years. It consists of decombining service groups (SGs) when possible, or dividing nodes into smaller groups when decombining SGs is no longer viable.

Traditionally SGs have consisted on a number of nodes combined together in the cable headend, and nodes include a number of homes passed and corresponding subscribers. Therefore, service group segmentation normally is achieved initially by separating nodes into smaller SGs, and when SGs consist of a single node these are segmented further by separating a number of the homes in a node into a new, separate node.

For example, assume that a SG consists of 2,000 homes passed (HHP), which results from combining 4 nodes, each with 500 homes passed. The SG decombining could be initially achieved by dividing the SG into 2 SGs, each consisting of 1,000 HHP. The segmentation could continue by separating each of the 4 nodes into a separate SG, consisting of 500 HHP/SG. Beyond that, SG segmentation would include “breaking up” each of the nodes into a smaller group by adding 1 additional node, creating nodes (and SGs) consisting of 250 HHP.

The following a key options for SG segmentation:

I. SG decombining is generally achieved by adding equipment in the cable headend. This re-uses the spectral HFC network capacity in smaller SGs.

II. Node segmentation requires the same additional equipment in the headend, but also requires that additional nodes, and/or fiber be installed in the plant.

And, the following are key factors to consider regarding SG segmentation:

A. SG segmentation usually involves the same decomposition in the upstream (US) and downstream (DS).

B. The relative cost of SG segmentation is higher for node segmentation than for SG decombining. This is because the work requires for the former requires the installation of additional nodes and/or fiber in the network (node splits), which in some cases is substantially more expensive. Conversely, in general SG decomposition is significantly less expensive than node segmentation.

C. However, when additional peak capacity is needed, such as in high-speed data (HSD) services, the SG segmentation is not a viable solution since it does not inherently add peak capacity.

2. Move QAM generation downstream into the network

This option would require including the PHY, part, or all, the MAC, or all of the CCAP functionality into a line-gear device which would be installed in the HFC network. Depending on the functionality being ‘remoted’ into the HFC network and the desired interoperability, this option would require the creation of specification. Connectivity back to the headend would be achieved via a baseband laser, such as point-to-point Ethernet, as opposed to an analog modulated laser as used now in HFC network.

The advantage of this approach is the migration to a baseband laser, and the operational simplifications that this entails. This approach would also result in additional capacity given the inherent segmentation that would be implemented. And, given the reduction in noise sources (e.g., removal of the analog laser, shortening of the links especially upstream, and reduction in the number of components), it should be possible to achieve higher order modulation rates than are possible to achieve with the PHY located in the headend.

From an operational perspective, however, the proliferation of intelligent devices that would need to be maintained, upgraded, and supported, might result in an increased complexity.

3. Implement PHY and MAC improvements

Clearly, cable systems today are capable of supporting higher order modulations, resulting in greater bit transmission capacity in the same spectrum. For example, it is considered possible to support 1,024 QAM downstream modulation in current cable systems. In fact, it should be possible to support even higher downstream modulations such as 2,048 and perhaps even 4,096 QAM. In addition, it should be possible to support 256 QAM in the upstream, and perhaps even higher order modulation rates. These improvements would come at a cost of higher signal-to-noise requirements, which are believed possible to achieve in today’s cable systems.

Additionally, given advances in CPU performance in DOCSIS components, both in the CPE and the CMTS, it should be possible to replace the currently used Reed-Solomon forward error correction (FEC) for Low-density Parity Check (LDPC) FEC. This change is expected to provide an improvement in bitrate equivalent to 2 bits/Hz.

Additionally, it appears that it would be beneficial to migrate to multicarrier modulation techniques, such as orthogonal frequency division multiplexing (OFDM) for the downstream, and orthogonal frequency division multiple access (ODFMA) for the upstream, as opposed to the currently used single-carrier approach.

OFDM and OFDMA offer superior performance and benefits over the older, more traditional single-carrier QAM modulation methods because it is a better fit with today's high-speed data requirements. The use of OFDM, and OFDMA, has become widespread and their implementation well understood in recent years, which was not the case when DOCSIS was initially conceived 15 years ago OFDM when it was extremely difficult to implement with the electronic hardware of the time. These techniques remained a research curiosity until semiconductor and computer technology made it a practical method in recent years, and extensively used for cellular and Wi-Fi transmission. OFDM, and OFDMA, is perhaps the most spectrally efficient method discovered and implemented so far.

4. Increase HFC downstream capacity beyond currently deployed, and/or move split to higher spectrum for increased upstream capacity

From a headend equipment perspective, this option is generally readily available to MSOs. However, CPE equipment would have to be implemented to support the new enhanced upstream and/or downstream spectrum.

From a network perspective, this option involves the change of the diplexers throughout the network such that the frequency division crossover is moved from the 42-50 MHz up to a higher portion of the spectrum, plus the simultaneous expansion of the network capacity to 1 GHz via a retrofit of the active components with minimal changes to the plant spacing and passive components. However, from an operational perspective, this option requires perhaps the most operational change to existing services, such as the removal of analog channels in that portion of the spectrum. That may not be possible for many MSOs that are either required to maintain support for analog TVs directly (e.g., without DTAs), or are unable to remove the analog channel for contractual reasons, or some combination of the above two reasons.

Even if removing the analog channels is possible, this option seems to require the installation of CPE filters in most or perhaps all home CPE devices (e.g., TVs, VCRs, etc.) to both protect that portion of the spectrum from emissions from such home devices and to protect the devices themselves from the levels of transmission of the new CPE that would use that portion of the spectrum for transmission.

And, even if removing the analog channels and deploying the necessary filters were possible, this solution alone provides limited additional US capacity in the network, as follows:

- A move of the split to 65 MHz provides an additional capacity of just 15 MHz, which less than doubles the current capacity. By all accounts, this is a change not worth embarking on.
- A move of the split to 85 MHz almost triples the US capacity, and the simultaneous expansion of the DS network capacity to 1 GHz would add a net 15-30 new DS QAMs (this calculation considers the combined effect of expanding the capacity of the network to 1 GHz from 860 MHz or 750 MHz respectively, and the loss of DS spectrum with the move of the split into the current DS region).
- The shift of the split up to the 200 MHz is also being considered, but while this change would provide much more US capacity, it would reduce the next number of DS capacity significantly and would require the change of large numbers of non-DSG STBs (most of the STBs deployed to date) because the existing and

extensively deployed OOB carriers would become inoperable since the region of the spectrum these utilize would be used for the US. Additionally, this change has other plant implications, such as the US equipment currently deployed would not support such extensive US, and thus a new HFC return strategy/equipment would be required.

5. Develop technology that would operate in unused portions of the spectrum, and even unleash spectrum above current top range (e.g., above 1 GHz)

Unlike option #4, this approach involves equipment not currently available. Instead, implementation of this option will require the development of network components and corresponding equipment that would make use of the existing forward spectrum but would use an unused portion of the spectrum, above 750, 860 MHz, or even 1 GHz. This new technique, which we will call High Spectrum Overlay, would require new equipment that could be built in the form of a new network gateway that could be installed in the headend, or in the vicinity of the node, or even deep within the HFC network. This new equipment would provide the 'translation' from the optical transmission generated at the headend into electrical signals, and RF transmission from the location of the converter to the coaxial portion of the HFC network.

This approach would increase US and DS capacity considerably, likely providing multiple Gbps of net additional US and DS capacity. In the process it leaves legacy services and existing CPE untouched.

However, this approach will require considerable equipment development before it would become available for deployment. Such equipment would use spectrum above that being used today for both additional US and DS capacity.

This option could be implemented in three fundamental ways: where the network gateway is located in the headend, or where the network gateway is deployed in the vicinity of the node, or where the network gateway is deployed throughout the HFC network.

In the first case, the RF signals would have to traverse the entire HFC network, including the forward and return analog modulated lasers and receivers, thereby being limited to the spectrum manageable by the analog modulated lasers and receivers.

In the second, the RF signals would traverse the various amplifiers within the coaxial part of the network, but would not require of transmission via the analog modulated lasers and receivers.

And, in the third, the network gateway would be installed in the vicinity of each active component where advanced services are to be provided. This option is known as a Passive High-Spectrum Overlay system. Therefore, this option would require the deployment of additional fiber beyond what's already installed in the network, namely between the existing node and each of the active components in the HFC network. In that way WDM would be used to carry baseband signals up to the node, from which traditional PON technology would be used to interconnect each of the new network gateways back to the HE.

Any modern HFC network should support a Passive High-Spectrum Overlay. Figure 3 depicts an initial deployment of Passive High-Spectrum gateways, for which EPON equipment is deployed in the headend, a separate optical wavelength is used in the trunk fiber to carry the EPON signals up to the node (shown in dashed blue lines), additional fiber is deployed in the distribution portion of the network (shown in solid blue lines), and new Network Gateways that provide optical-to-electrical signal conversion are installed to provide the overlay within an HFC segment between amplifiers.

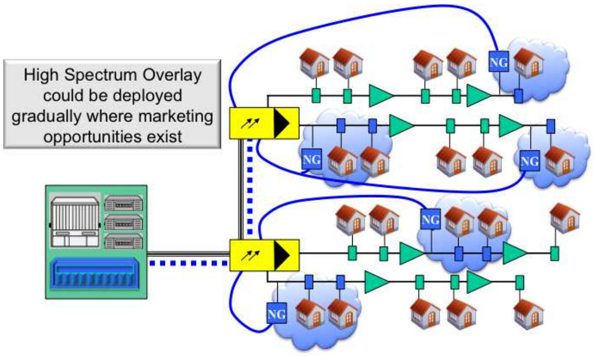

Figure 3: Initial High Spectrum Overlay

This approach should not be construed as resulting in a Node + 0 HFC cascade reduction. This is because the cascade of HFC actives is not modified. Instead the RF output of the gateways deployed in the HFC network and operating above 1 GHz are combined with the RF signals existing in the coaxial network which operate below 1 GHz, much in the same way as narrowcasting a set of signals on a per service group basis where the other signals are broadcasted to the set of service groups.

The following categories of work would need to be performed in the plant in order to achieve the above:

- WDM could be used from the headed to the location of the node to reuse the existing long-haul fiber.

- To provide the remaining optical link from the node to the location of each active, additional fiber would be over-lashed to the distribution coaxial hardline cable, which is generally a short to medium length span.

- Finally, in order to pass RF signals above 1 GHz on the distribution network, it is likely that a proportion of the tap faceplates would need to be replaced, although it is expected that the tap housing will likely support these new faceplates, and that only faceplates serving subscribers and upstream from it would need to be replaced.

Assuming a high-bandwidth optical network from the headend to the network gateway, such as 10 Gbps EPON, and a high-order modulation and encoding scheme, it is expected that a transmission achieving 8-10 b/Hz might be possible, therefore resulting in a combined US/DS payload transport capacity of approximately 3-5 Gbps.

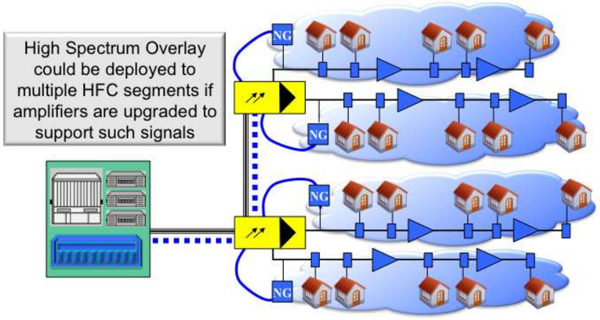

Figure 4: Multi-segment High Spectrum Overlay

Figure 4 depicts the case of a deploying Network Gateways at node locations. This option would require less fiber, but would necessitate a rebuilt with amplifiers that would pass the new RF signals.

POSSIBLE IMPLEMENTATIONS

In evaluating the possible approaches outlined above, and taking into account the technologies available to date, it makes sense to consider the following implementations:

enhancement to DOCSIS for residential applications, and development of a new transport alternative to EPON over HFC for commercial applications. Naturally, despite the primary target services, either of the two technologies could be used for either or both services.

DOCSIS Enhancements

Given the success and widespread use of DOCSIS-based services to date, and the advent of the technical advances outlined above, it seems plausible to consider the following enhancements to DOCSIS to enable additional capacity and a more efficient use of HFC spectrum:

- Use of higher-order (1/2/4K QAM) and modulation techniques (OFDM/OFDMA) to improve throughput and simultaneously reduce spectrum utilization by as much as 50%,

- Replace the current Reed-Solomon FEC technique with a more modern Low Density Parity Check (LDPC) FEC, which would improve overall efficiency by as much as 25%,

- Enable use of additional spectrum for the US, beyond the current 5-42 MHz, up to 100 MHz or even higher spectrum, to increase US transmissions by a 3x to 5x factor, and

- As capacity is enhanced, consider simplifications of the DOCSIS protocol that may reduce implementation complexity, accelerate the availability of newer implementations, and reduce costs.

Implementation of the above new functionality will have to be done taking into account backward and forward compatibility to maximize the benefit for current and new equipment.

New HFC transport for EPON

Similarly to the enhancements now available for DOCSIS, it seems possible to implement a new transport for EPON over coax. Envisioned in the past as a component of Comcast’s Next Generation Access Architecture[4], a new transport for EPON Protocol over Coax (EPoC) is now under development at IEEE. This new transport will make it possible to provide EPON services to end-devices attached via cable operator’s coax network rather than only via fiber cable.

The work currently underway, known as an IEEE 802.3 Study Group, is intended to demonstrate the feasibility of implementing a coax transport for EPON using technologies and approaches similar to those that would be applicable to DOCSIS. Once completed the work of the Study Group, a Task Force would be formed to define the new PHY for EPON over coax.

This work would lead to the availability of a coaxial-attached alterative to EPON, which would enable MSOs to deploy EPON services to customers already served by its HFC network. This should result in a more economical and operationally simpler way to provide Metro Ethernet (MetroE) services to business customers without having to deploy fiber to each potential customer premise.

OVERALL ACCESS ARCHITECTURE

The new edge platform devices currently under development by vendors, as specified by the CCAP architecture, will support either of the approaches described above. The

---

[4] What is CMAP? Jorge Salinger and John Leddy, CED Magazine, February 2010

CCAP architecture already supports the modularity necessary to upgrade line cards progressively as new technologies become available.

For the case of the enhancements to DOCSIS outlined above, it seems reasonable to expect that the current downstream line cards could be updated via field programmable gate array (FPGA) programming changes, such as Hardware Descriptor Language (HDL) or Register Transfer Level (RTL) programmable changes. In the case of the upstream, it is expected that new line cards could be developed that would take advantage of the new technologies.

Furthermore, the CCAP architecture provides support for EPON, such that even EPoC is supported in the overall access architecture.

SILICON DEVELOPMENT

One important consideration in evaluating the benefits of each approach is the need and availability of silicon components, or on the flip side the need for its development.

This is critical for the following fundamental reasons:

a. When silicon exists the availability of the system solution is quicker, whereas when it needs to be developed the timeline is significantly longer, and

b. If silicon devices, or at least some of their components, are used for multiple purposes, especially for multiple industries, then their production increase rapidly and costs decrease considerably.

Some of the new technology enhancements will likely require silicon development, but others would not, for which technology design decisions would be important.

CONCLUSIONS

Additional HFC network capacity will be required for narrowcast services for both residential and commercial service applications in years to come. The expected growth appears to be quite large.

New technologies are now becoming available that would make it possible to achieve higher throughout and more efficient use of spectrum. This includes higher-order and more modern modulation techniques, more sophisticated forward error correction, and the use of more spectrum than currently utilized.

This paper presented an analysis of these technology options and their corresponding pros and cons, and outlined how these technologies could be used to enhance the current transport options available to MSOs, such as DOCSIS, and to create new infrastructure options, such as EPoC.

ACKNOWLEDGEMENTS

Thanks to my colleagues Sam Chernak, Doug Jones, Saif Rahman, John Bevilacqua, Chris Bastian, and Dave Urban from Comcast for their contributions to the discussions leading to the proposals and conclusions presented in this paper.

Also, many thanks to my colleagues from CableLabs®, the many vendors and advisors of Comcast, and especially fellow MSOs, for the many discussions leading to the materials presented.