# Has the time come for Remote PHY in the HFC network?

With the advent of DOCSIS® 3.1, it would be beneficial to implement a digital forward link in Hybrid Fiber-Coax (HFC) networks to improve the performance of cable modems and similar devices. Several options are available for implementing a digital forward link, which can be categorized as described in this paper. Proposes tenets and implementation objectives are presented and evaluated for each of the options.

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

**Jorge D. Salinger**
VP, Access Architecture
Comcast
1701 JFK Boulevard
Philadelphia, PA 19103
+1 (215) 439-1721
jorge_salinger@cable.comcast.com

# Overview

NOTE: The concepts and proposals presented in this paper are those of the author and presented for discussion purposes only, and do not reflect any actual plans from Comcast. Similarly, all examples presented are only provided for illustrative purposes.

The analog forward optical link in Hybrid Fiber-Coax (HFC) networks is becoming a limiting factor in achieving the advanced modulation profiles included through the data over service interface specification (DOCSIS®) 3.1 specification. Replacing it with a digital forward link would enable more cable modems to run at the high-order modulation profiles included in DOCSIS 3.1.

There are several ways to implement a digital forward link in the HFC network, including Remote Physical Layer (Remote PHY), a technique consisting of distributing the physical layer components (e.g., modulation and demodulation) of the access network equipment (e.g., the Converged Cable Access Platform, or CCAP™) into the HFC network (e.g., nodes). Many cable multiple system operators (MSOs) and access equipment vendors think that the time has come to implement Remote PHY. But this is a very big step for the cable industry, which has kept these functions in the headend.

This paper will discuss why a digital forward link is beneficial, the various options available and how each would work, why Remote PHY might be the best solution, and why the time may have finally arrived for implementing this technique.

First, the paper will cover the background. This will include the continued and projected demand for narrowcast services, especially video services through both the IP transport as well as over the traditional Motion Picture Experts Group Transport Stream (MPEG-TS) infrastructure. The background will also cover how MSOs continue to expand the access network capacity dedicated to those services through use of the traditional tools at their disposal.

The paper will then describe the various options available for deploying a digital forward link, including why and how these techniques will help to increase the overall potential capacity supported by HFC networks. The discussion will focus on the technologies themselves (e.g., supporting modulation orders higher than 4096 Quadrature Amplitude Modulation (4096-QAM) for the downstream and higher than 1024-QAM for the upstream with Remote PHY), and also on the operational improvements afforded by a digital forward versus keeping the analog forward link in place.

# 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 cloud-based digital video recorder (cDVR) 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 HD and SD versions of the program are not distributed unless a subscriber is requesting them, 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 cDVR, and a full-array of HD video-on-demand services. For the former, initial observations suggest that network requirements for cDVR 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 continues. Network operators have observed an increased use of HSD service capacity for well over a decade now, as shown in Figure 1, 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.

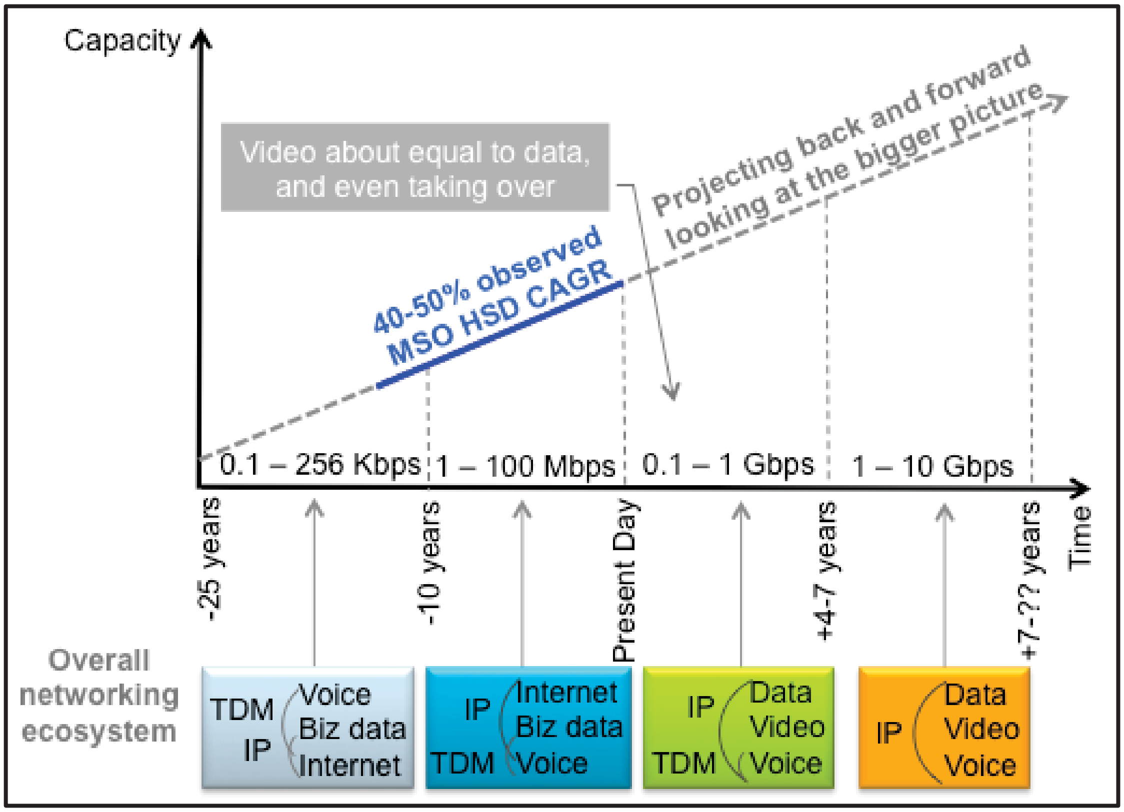

Figure 1: Example of HSD capacity utilization over time

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 same rate as in the past, 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 narrowcast capacity, as shown in Figure 2 below.

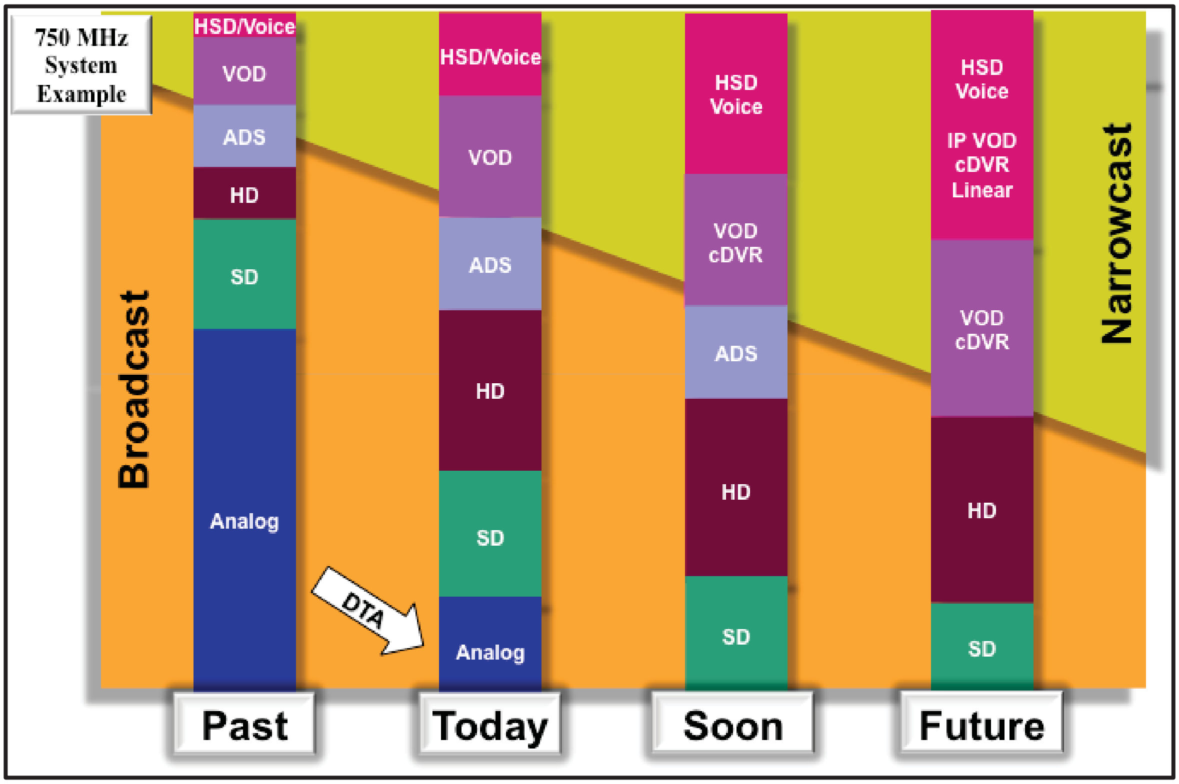

Figure 2: Example of narrowcast service growth over time

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, and their operational complexity and cost 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]. Furthermore, with the recent release of the H.265[3] standard in April of 2013, it is possible to achieve a 50% improvement over H.264. And the use of VBR promises to result in a capacity efficiency gain of as much as 70% versus CBR[4]. The combined gains from using the above 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.

---

[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] ITU-T Recommendation H.265: 2013, High efficiency video coding

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

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 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.

# The Advent of DOCSIS 3.1

As it has been pretty well advertised in the media, DOCSIS 3.1 is under development.

NOTE: For further details on the DOCSIS 3.1 technology and its implementation, see the DOCSIS 3.1 Symposium planned for the SCTE Expo 2013 show.

The key motivation for the new version of the DOCSIS specification is, in a nutshell, to scale DOCSIS more efficiently, both from the cost and operations perspectives.

While for the first 10 years or more it was possible to offer Internet services and support its growth with just 1 DOCSIS channel, services today require many more channels. This is because 1 DOCSIS channel provides almost 40 Mbps, which was well above the data rate of the services offered in the past. However, the year-over-year growth drove service speeds well above the initial levels, to 20, 50 and even higher Mbps tiers today, which can’t be supported by the single channel. MSOs then went to multiple DOCSIS channels, now reaching 4 and even 8 channels, and soon requiring well beyond that.

To that end, the 3 key goals and features of DOCSIS 3.1 are:

1. Much more efficient use of spectrum, with up to 50% improvement in bandwidth efficiency (or bps/Hz, resulting from:
    a. the use of more efficient forward error correction (i.e., replacing the older and less efficient Reed-Solomon approach for the more modern and far more efficient Low Density Parity Check, and
    b. the addition of the higher-order modulations 1024 and 4096 QAM downstream and 256 and 1024 QAM upstream.

    These new modulation schemes provide 2 and 4 bits/Hertz/second improvement in both upstream and downstream, while the use of the new forward error correction approach provides approximately 5 dB better RF performance. The end result is that MSOs will be able to transport 1 Gbps of DOCSIS capacity in about 120 MHz of spectrum while doing the same with the current DOCSIS approach using single-carrier QAM requires about 180 MHz of spectrum.

2. Cost reduction, mainly by leveraging technologies commonly used in other transmission media, such as the inclusion of Orthogonal Frequency Division Multiplexing, which is used extensively in wireless and wireline transmission media. Specifically, the addition of OFDM for the downstream and OFDMA for

the upstream should enable MSOs to reduce costs while "packing" more bits in the HFC network more efficiently since these technologies likely result in a larger supplier ecosystem, increasing innovation and fueling competition.

3. Enable a simple and orderly transition strategy, both with respect to compatibility with previous generation CMTS and CM equipment while supporting an expanded spectrum capacity in the HFC network. Specifically, DOCSIS 3.1 cable modems will operate with DOCSIS 2.0 and 3.0 CMTS/CCAP equipment, enabling deployment of DOCSIS 3.1 CPE as soon as available. Similarly, DOCSIS 3.1 CCAPs will support DOCSIS 2.0 and 3.0 CPE allowing MSOs to upgrade headend equipment without having to change any of the existing CPE. And, both DOCSIS 3.1 CM and CMTS equipment will support the currently required upstream and downstream spectrum, plus an expansion of the upstream to 85 MHz and beyond, and of the downstream up to 1.2 GHz.

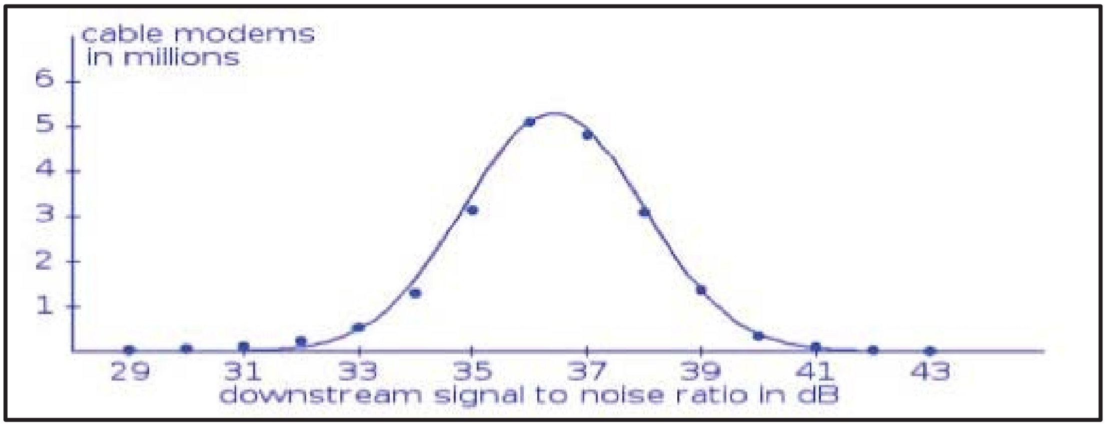

Figure 3: Example of downstream SNR for a large population of cable modems[5]

Figure 3 depicts the downstream SER as reported by a very large population of cable modems. This data verifies that many cable modems will be able to support the high-order modulation profiles included in DOCSIS 3.1. However, many others will not without an increase in SNR.

Assuming an 8/9 coding ratio, Table 1 shows the SNR requirements for the new modulation rates included in the DOCSIS 3.1 specification:

| 512 QAM | 1024 QAM | 2048 QAM | 4096 QAM | 8192 QAM[6] | 16384[6] QAM |
|---|---|---|---|---|---|
| **27 dB** | 30 dB | 33 dB | 36 dB | 39 dB | 42 dB |

Table 1: SNR required for DOCSIS 3.1 modulation rates

---

[5] Data collected by Comcast and reported to the DOCSIS 3.1 working group
[6] 8192 and 16384 QAM are considered for future inclusion in the DOCSIS 3.1 specification

Applying the SNR requirements from Table 1 to the population of modems shown in Figure 3, we can easily see that a large population of cable modems would not achieve sufficient SNR to operate at 4096 QAM. Furthermore, if sufficient headroom is allowed to account for environmental fluctuations, the population of cable modems that would not receive signals with sufficient SNR to operate at 4096 QAM would be significant.

## The Analog Modulated Forward Link in HFC

As their name indicates, Hybrid Fiber-Coax networks use of a fiber transport between the headend and the coaxial cascade. This fiber link, intended to reduce the size of cascades, mainly driven to improve performance, was originally developed with analog modulated lasers and receivers in both directions, upstream and downstream.

Over time, the performance of the upstream link was improved by replacing the analog modulation with a digital transport. This change improved performance significantly, and allowed for longer distances between the headend and the node. Different vendors implemented their own methods and technical capabilities to implement a digital transport, which resulted in incompatible systems and required the use of the same vendors' components for both the node and the headend.

However, the downstream link remained almost unchanged over time, with the only enhancements focused on improving distance and RF spectrum capacity. Performance has not really been an issue like it was in the upstream.

But more importantly, while the digital capacity of the upstream was limited to a few megabits per second, well under a gigabit of digital capacity which could easily be digitized and carried with Ethernet optics, the downstream digital capacity necessary to transport the downstream spectrum has been considerably higher, reaching and even exceeding 10 gigabits per second.

Because of the above, analog forward links continue to be used to date. And, while headend equipment is currently capable of launching signals with >47 dB MER performance, which would be sufficient to generate and transport 16,384 QAM signals, analog lasers are limited to about 35-38 dB of MER performance, which would limit end-of-line performance to barely enough for 2,048 QAM or 4,096 QAM in short cascades the best of the cases.

## Description of Options for Digital Forward Link

As time has gone by, technology evolution and certain developments as described below have enabled options for implementing a digital forward link. These include:

1. Evolution of QAM edge modulators which have gone from single and/or a few modulators to supporting 32, 64 or even more modulators,

2. Development of the CCAP, combining the functions of the video QAM modulator and DOCSIS into a single platform, and
3. Migration to digital video, either partially for now or already completely.

With this technology evolution, it is conceivable to remove the RF combiner network, and instead implement it digitally in the edge device, such as the CCAP. Figure 4 conceptually depicts the output of a CCAP device.

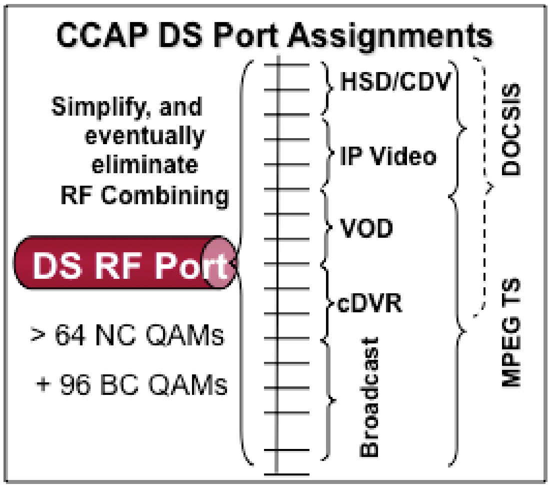

Figure 4: CCAP Downstream Port Functions

This evolution of the edge headend devices makes it possible to envision several options for digitizing the forward link.

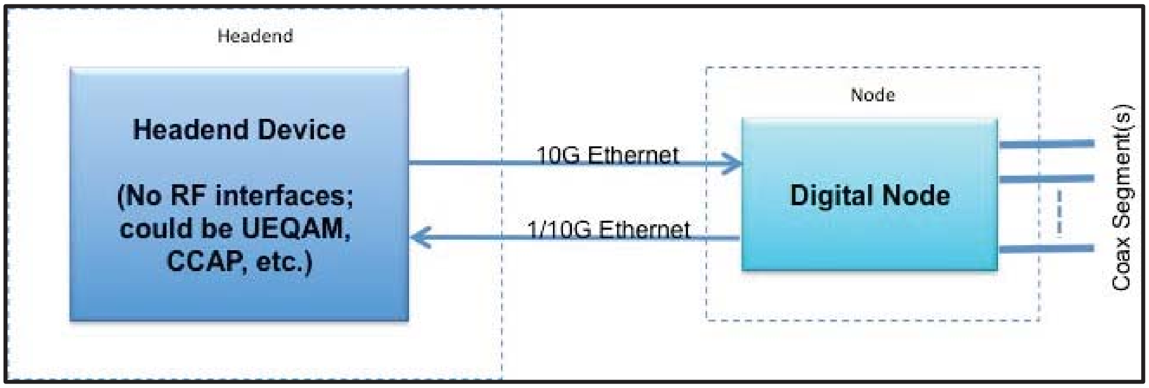

Figure 5: Digital Forward – High Level Architecture

Fundamentally, the migration to a digital forward includes the components included in Figure 5, as follows:

- The headend device, such as a CCAP, which would be a high-density edge QAM comprising QAM modulation for the entire spectrum,

- The node would contain components normally implemented in the edge QAM or CCAP which generate the RF signals,
- The link between the headend device and the node would be comprised of a digital interface, such as an Ethernet link.

There are then various approaches for how a digital forward link can be implemented to replace the currently used analog link. These various approaches for distributing the various components can be categorized into 4 groups, plus 1 option that would still leave an RF generation at the headend device, as outlined in Table 2:

| Option | Description and Approach |
|---|---|
| 1. Maintain RF output in the headend | 1.a Headend equipment remains unchanged<br>1.b Headend RF output is digitized, transported digitally, and RF is regenerated in the node |
| 2. Remote the DAC from the PHY | 2.a The DAC is removed from the headend<br>2.b Digital samples are transported digitally to the node where the DAC generates the RF signals |
| 3. Partition the PHY and remote the lower portion of the PHY | 3.a The PHY is split between the headend and the node<br>3.b The digital bit stream between upper and lower PHY is transported from headend to node |
| 4. Remote the entire PHY | 4.a The entire PHY modulation is moved to the node<br>4.b The MAC remains in the headend, and MAC frames are transmitted from the headend to modulator that resides in the node |
| 5. Remote the entire PHY and MAC | 5.a The entire PHY and MAC is removed from the headend device and placed in the node<br>5.b IP frames are transported from the headend to the node. |

Table 2: Categories of options for implementing a digital forward link.

# Comparison of Options for Digital Forward Link

There are pros and cons for each of the options. The following sections outline these trade-offs.

<u>Option 1:</u> RF remains in the headend

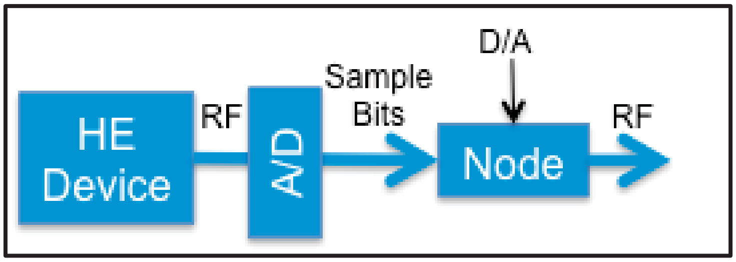

Figure 6: Block diagram for Option 1

- Equivalent to digital return, the RF output from the headend device is digitized, transported digitally, and converted back to RF in the node.
- Maintains HFC transparency
- This option results in the highest bitrate over fiber; the capacity for multiple nodes would not fit into the available capacity of one 10G fiber
- There is a loss of MER in the double conversion, so this option provides the least performance improvement
- Results in the least intelligence placed in the node, but an additional conversion stage is added in the headend

<u>Option 2:</u> Digital-to-analog conversion is moved to the node

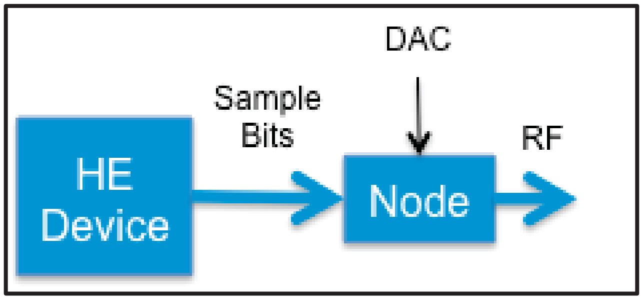

Figure 7: Block diagram for Option 2

- Requires separation of the digital-to-analog conversion from the modulator
- Together with Option 1, results in the least intelligence in node
- Similar high bitrate over fiber as Option 1; capacity for multiple nodes would not fit into the available capacity of one 10G fiber

Option 3: Lower PHY is moved to the node

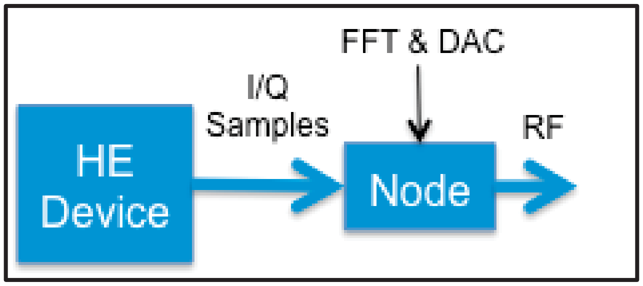

Figure 8: Block diagram for Option 3

- The PHY layer needs to be split into two components: upper and lower PHY
- More intelligence than in either of the previous options is placed in the node
- Although lower than the previous options, this option also results in a very high bitrate over fiber
- This option would require and industry proprietary point-to-point link between the headend port and the node to transport the I and Q samples
- Implementation of this option would require the definition of interfaces which have never been defined in previous versions of the DOCSIS specifications, which in turn would result in modification of the silicon used and/or planned to date, and therefore results in the highest implementation complexity

Option 4: Entire PHY is moved to the node

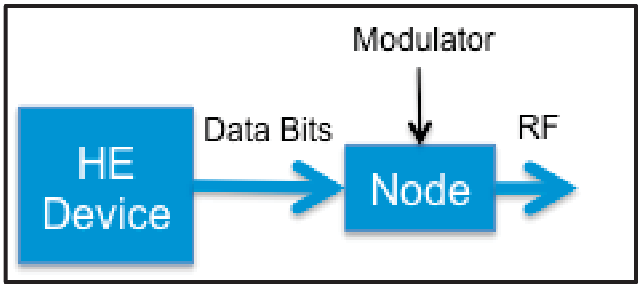

Figure 9: Block diagram for Option 4

- More intelligence is placed in the node than with all previous options
- This option results in the lowest bitrate over fiber; multiple nodes fit into the capacity of a 10G fiber
- Enables a packet-based link between the headend and node, which results in significant benefits outlined later in this paper
- Could use existing/planned silicon devices, and thus may be the easiest and quickest to implement
- Offers the best MER performance improvement over analog

Option 5: Move PHY and MAC to the node

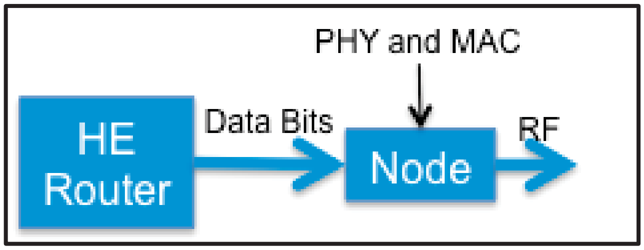

Figure 10: Block diagram of Option 5

- This option puts the most intelligence in the node
- The data rate between the headend and the node is equivalent to the actual data transmitted, except for the addition of ancillary network data
- Same packet-based network benefits as Option 4
- Same highest MER performance as Option 4

# Proposed Tenets for Digital Forward Link

In considering the various approaches for implementing the digital forward link, and the 5 options in which these approaches could be categorized, it might make sense to consider tenets for its implementation.

The following list outlines proposed basic, underlining tenets for digital forward link:

1. Headend and node devices for digital forward link should be interoperable
2. Limit the specifications to the areas that are absolutely needed for interoperability
3. Minimize the electronics that is housed in the node to the extent possible
4. Minimize the software that is placed to run in the node
5. Minimize the amount of capacity needed in the optical link
6. Keep as much of the higher layers as possible in the headend
7. Make the timing requirements for the node as simple as possible
8. Keep the independence between the DS and US as much as possible
9. Maintain the digital forward link independent from the DOCSIS version

Additional objectives could be established that would further limit the options to be considered for the digital forward link. What follows are additional proposed objectives:

A. Develop an architecture that enables scalability as capacity is needed over time

B. Minimize the need for replacing the node components as additional capacity is needed

C. Leave system components that scale with capacity in the headend

D. Use technologies used in other communications protocols when possible

E. Minimize space and power requirements in the headend

F. Minimize power requirements in the node, targeting the power consumption of a line extender as the maximum power requirement

G. Enable the use of the digital forward link for other networking functions

## Comparison of Digital Forward Link Options

Table 3 is an analysis of the pros and cons for each of the 5 options considering the tenets and additional objectives outlined above:

| | | Option 1 | Option 2 | Option 3 | Option 4 | Option 5 |
|---|---|---|---|---|---|---|
| Basic Tenets | 1 | 🟩 | Should be | Should be | 🟩 | Should be |
| | 2 | 🟩 | 🟩 | 🟩 | 🟩 | Should be |
| | 3 | 🟩 | 🟩 | 🟩 | 🟩 | 🟥 |
| | 4 | 🟩 | 🟩 | 🟩 | 🟩 | 🟥 |
| | 5 | 🟥 | 🟥 | Not likely | 🟩 | 🟩 |
| | 6 | 🟩 | 🟩 | 🟩 | 🟩 | 🟥 |
| | 7 | 🟩 | Should be | Should be | Should be | Should be |
| | 8 | 🟩 | 🟩 | 🟩 | 🟩 | Not likely |
| | 9 | 🟩 | Should be | Should be | Should be | Not likely |
| Additional Objectives | A | 🟩 | 🟩 | 🟩 | 🟩 | 🟥 |
| | B | 🟩 | 🟩 | 🟩 | 🟩 | Not likely |
| | C | 🟩 | 🟩 | 🟩 | 🟩 | 🟥 |
| | D | 🟥 | 🟥 | 🟥 | 🟩 | 🟩 |
| | E | 🟥 | Not likely | 🟩 | 🟩 | 🟩 |
| | F | 🟩 | 🟩 | 🟩 | Should be | Not likely |
| | G | 🟥 | 🟥 | 🟥 | 🟩 | 🟩 |

Table 3: Comparison of tenets and additional objectives for digital forward link options

Given the comparison of the applicability of each of the proposed tenets and proposed additional objectives outlined above, it seems as if Option 4 might be the best target for the proposed tenets and objectives.

## Summary and Conclusions

As the projected demand for narrowcast services continues to grow at approximately the same rate as currently observed, especially related to high-speed data and video services through the IP transport, it will become useful to support the higher modulation orders that will be supported in DOCSIS 3.1 to support higher capacity in the available HFC network spectrum.

However, the analog forward optical link in HFC networks will become a limiting factor in achieving these advanced modulation profiles in most cable modems. Replacing the analog forward link with a digital forward link would improve the link MER, which given the advances in headend equipment would enable most cable modems to support the high-order modulation profiles included in DOCSIS 3.1.

As outlined in this paper, there are several ways in which a digital forward link can be implemented. These various approaches can be categorized into 5 options, which are described in this paper. One of the options for implementing a digital forward link is called Remote PHY, which is a technique consisting of distributing the physical layer components (i.e., the downstream modulation) of the access network equipment (e.g., the CCAP) into the HFC network (e.g., nodes).

The paper outline proposed tenets for the definition of the digital forward, and proposes additional objectives for its implementation. A comparison of each of the 5 options for the digital forward link is presented for each of the various proposed tenets and implementation objectives, from which Remote PHY may appear to better target such proposed tenets and objectives.