NEXT GENERATION ACCESS ARCHITECTURE OPTIONS FOR ADDITIONAL UPSTREAM AND DOWNSTREAM TRANSMISSION ON HFC

Jorge D. Salinger
VP, Access Architecture
Comcast Cable

Abstract

HFC networks currently have an US capacity limited by the available spectrum. Typically, North American networks are limited to the 20-30 MHz of spectrum these systems have available between 5 and 30 or 42 MHz, and European systems have an additional 25 to 45 MHz of available spectrum. At 6.4 MHz of channelization and QAM64 modulation, assuming complete use of the available spectrum, this capacity allows transmissions of 100-300 Mbps.

The potential migration of North American systems from the current 42 MHz-split to 65 MHz, 85 MHz, or even higher split systems, and of European systems from 65 or 85 MHz to higher splits, are complex to implement, service impacting and financially and operationally onerous. The capacity gain from such migration is limited to approximately doubling, or perhaps tripling, the maximum transmission capacity. But, an indirect impact of performing such spectrum allocation change is a reduction of the downstream capacity of the HFC systems, which is larger as the split is moved further up into the downstream spectrum used today.

Taking into account other factors may result in concluding that other solutions may be more appropriate and no more operationally impacting than just changing the spectrum split to increase upstream transmission.

Such other factors may include:

- *Use of a different portion of the spectrum for US transmission to simplify the migration and avoid impacting currently available HFC capacity.*
- *Enable additional DS transmission capacity, not just additional US capacity, which will have the additional benefit of alleviating the expected peak capacity needed when new services (additional programs, additional technologies such as IP video simulcast, etc.) are added to an already fully allocated HFC network.*
- *Establish a new, more efficient HFC transmission mechanism for coupling with home gateway devices, which would be based on existing technologies, such as those being considered for advanced video service planned by most MSOs.*
- *Solve certain limitations of the current transmission mechanisms, such as enabling more native support for IP-based, enhanced business, wireless and other services and applications.*

When considering the above factors in addition to increasing US transmission capacity, alternate solutions to just changing the US/DS spectrum split may become more attractive and valuable to MSOs because the burden of such implementation could be attributed to the deployment of services not otherwise possible, with the resulting potential increase in both ARPU and RGUs.

This paper will then presents an analysis of various such options, and expands on one particular alternative that offers the following characteristics:

- *Leaves all current services untouched in the currently allocated HFC spectrum.*
- *Unleashes new spectrum in the HFC network.*
- *Could enable transmission of up to an additional 3 Gbps combined between upstream and downstream.*

To do so, the proposed approach makes use of traditional and already available technology and network strategies in a new way, such as:

- *Continue pushing fiber deeper into the HFC network, up to the current last active, without changing the current HFC architecture, such as not requiring cascade reductions.*
- *Use of existing baseband fiber-based technologies, such as PON, and EPON in particular, and expansion in the use of WDM, for transmission from the headend to the current location of the last active,*
- *Implementation of existing technologies for higher-order PHY modulation and encoding and light-weight MAC for transmission from the last active to the home gateway using silicon devices already under development.*
- *Superimpose the new transmission method in a portion of the spectrum above the band currently being used, such as above 1GHz, and still utilize spectrum allowing the use of passive devices such as below 1.8 GHz.*
- *Use of a new, simpler home gateway to terminate the HFC network in the home, bridging the new transmitted capacity described above, and using consumer network technologies inside the home.*

The proposed approach described above leaves existing services untouched, enables considerable new capacity in the HFC network (almost doubling the overall current HFC capacity), solves the current US capacity limitation, leverages existing technologies, builds upon the current MSO strategy of segmenting the HFC network to increase overall capacity, and should be neither more costly for initial deployment nor operationally more impacting than simply changing the HFC US/DS spectrum split to provide a marginal increase of US transmission capacity.

TYPICAL HFC NETWORKS

Most MSO's 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 high HSD service tiers, additional HD programs for both broadcast and especially narrowcast services such as VOD and SDV are deployed, or new services such as network-based 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 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 videiservices, especially nDVR, and a full-array of HD video-on-demand services. For the former, initial observations suggest that network requirements for nDVR 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. As shown in Figure 1, network operators have responded to competition by increasingspeed tier offerings, which over the long term amounts to year-over-year compounded growth of 40% to 60%, as observed for well over a decade now. The applications have changed, but the demand has continued to increase. One thing MSOs have learned, as have network operators of more traditional Internet network services, is that it is very difficult to predict what subscribers will use the additional capacity for, but that the bandwidth growth is somewhat predictable.

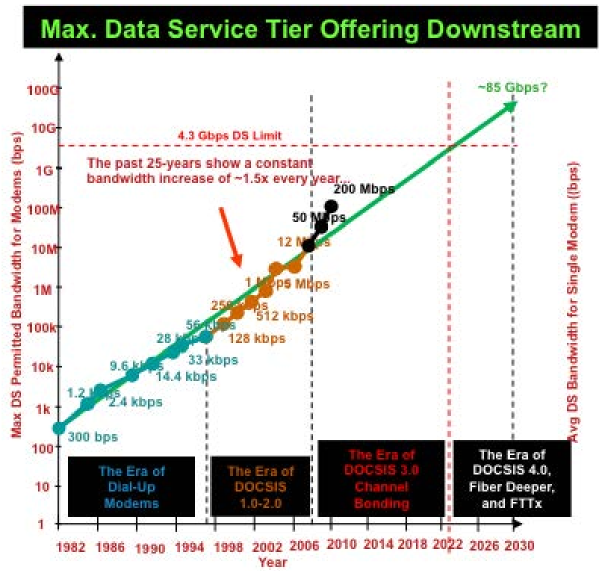

Figure 1: HSD Service Tier Trends[i]

In fact, such increase in demand for HSD capacity shows no evidence of decreasing. Quite to the contrary, all the evidence gathered over many years indicates that the growth will continue, and perhaps even increase when considering the observed increase in the consumption of video-based content. When projecting this experienced growth into the future, as also shown in Figure 1, it is easy to see that should the usage growth pattern continue networks will be required to provide services in the range approximating 1 Gbps within the next few years.

To support this growth, MSOs have deployed, or are considering deployment of, bandwidth reclamation tools such as SDV for digital broadcast, 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 AVCs and VBR. In the case of AVCs, coding efficiencies of approximately 50%, depending on implementation and content type, can be obtain with H.264[ii] and/or MPEG-4 Part 10[iii]. And the use of VBR could result in a capacity efficiency gain of as much as 70% versus CBR[iv].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 set-tops still support AVCs and VBR, especially the latter tool. However, because of the lack of widespread support across set-top populations, MSOs are only able to use these tools for unicast-based services, such as VOD, for the small minority of the set-tops that support the tools. These tools will likely enjoy significant support in newer, IP-video based services equipment moving forward.

NEW TREND: IP-BASED ADVANCED VIDEO SERVICES

Industry-wide, MSOs are now considering deployment of video services supporting IP devices. For this, 3 approaches are possible:

a. Reuse of existing legacy content by re-encapsulating at the edge,

b. Use of the existing legacy infrastructure to distribute IP-based content, and

c. Leverage the DOCSIS infrastructure to distribute IP-based content.

The first approach makes very efficient use of existing spectrum for a new application, but has significant draw-backs, such as: limits content to that currently distributed, need for CAS termination in a gateway device, complex and seemingly expensive gateway, proprietary implementation, etc.

The second requires the development of special-purpose software in the gateway, inability to benefit from certain network efficient DOCSIS features such as channel bonding, lack of standardization across potential suppliers, etc.

With the advent of lower cost DOCSIS CMTS gear in recent years and especially moving forward, the third has become the one garnering the most attention. It offers significant network efficiency factors that can't be simultaneously achieved with the other two methods, such as immediately benefiting from VBR gains, immediate use on all applicable clients of AVCs, multiplexing gains from channel bonding, standardization of the gateway implementation, benefit from multi-industry approaches for web-based video distribution, multicast and caching, and a very broad availability of development talent and supplier ecosystem.

When considering the options and their trade-offs, all of the above is believed to make the third approach, that of a DOCSIS-based gateway, significantly more efficient and long-term beneficial than that of the other two approaches.

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. This increase in simultaneous use of advanced video services while maintaining legacy services will be especially impacting over time as its penetration increases.

Initial target for most MSOs seems to be 2nd screen devices, such as PCs, tablets, game consoles and non-traditional customer-owned devices. This target subscriber use and device base likely only require low-resolution video services (e.g., 1.5 to 3 Mbps streams), possibly expanding to higher stream data rates for TV-attached devices and larger screens.

It is likely that if the deployment of these advanced video services is successful, a possible additional target would be to deploy IP-based set-tops as part of the mainstream MSO device profile, for which higher bandwidth full-resolution services would likely be used (e.g., 9 Mbps and higher resolution streams).

ALL-IP VIDEO CAPACITY EXAMPLE

What follows is an example of capacity required in a network to support an extensive deployment of advanced video services.

Assume the following scenario and circumstances:

- Typical service group of 400 HHP, which is an accepted average across the industry.
- 50% penetration of IP-based devices, which is typical in legacy video today.
- 2.5 devices per home, which is a high average for legacy video, but considered likely with IP-devices.
- 70% of homes active at peak, which is a typical industry average for video services
- 70% of secondary devices are active at peak, which is also a typical average
- Assume all-unicast video at 6 Mbps. This assumption is based on a number of factors, including: initial deployment of IP video will likely not benefit from multicast; over time it is expected that video viewing will trend towards more individualized viewing; rate could be

considered a high average, but video quality is constantly being enhanced requiring additional bandwidth; the use of ABR would cause the simultaneous transmission of content at multiple rates for different subscribers; etc.

Using the above assumptions, peak capacity could be established by the following formula:

$$400 * 0.5 * (1 + 1.5 * 0.7) * 0.7 * 6 \text{ Mbps}$$

The above would yield a total capacity of over 1.7 Gbps.

Assuming an additional capacity of 1 Gbps for HSD, the total capacity required would be over 2.5 Gbps, or an equivalent of over 70 DOCSIS channels.

When all services and all devices in the network are based in IP transport, a 750 MHz plant should suffice to support such services. In fact, a 550 MHz plant would suffice as well.

And, even a higher US capacity could be supported. A 200MHz split could yield >1 Gbps US, which could be accommodated in a 750 MHz system. In such case the remainder of the spectrum implementing 10 b/Hz yields >5 Gbps. Furthermore, implementing any change in split would likely be accompanied by an expansion of the spectrum to 1 GHz, yielding significantly more DS capacity than would seemingly be required.

SIMULTANEOUS IP AND LEGACY

At issuemay not be whether the existing network would support a set of all-IP services, but rather that, if implemented, a transition to an all-IP set of services will likely take considerable time, and during such transition it would be necessary that both advanced video services and current legacy services be provided simultaneously.

For financial and operational reasons, the migration to all-IP services will likely take a long time to be completed. The reasons for this are the same as always: technologies take long to be developed and longer to be adopted;operational readiness is a long road for any new approach, especially in replacement cases; amortization of equipment dictates a need for preserving the investment; and even when equipment is amortized it is by no means trivial to fund a replacement.

Moreover, as support for advanced video services and IP-based devices is deployed and grows over several years, additional legacy devices will continue to be deployed. For example, viewership of HD VOD will increase, requirements for HD services in SDV increase, and nDRV with its expected high usage may be accompanied by a high use of content in HD, all of which will expand the need for spectrum for legacy services,which must be preserved and expanded.

Therefore, both legacy and IP-based services would need to be supported simultaneously. In that case, which appears likely, the deployment of IP-based services will occur while supporting a full array of legacy services, which is still growing.
For that, more HFC capacity than is used today would be needed. Current spectrum utilization for legacy services will likely not decrease, but rather need to increase as described above. In parallel, spectrum would need to be allocated for IP-based services, which is likely to grow over time while the capacity allocated to legacy services has not yet decreased.

ADDITIONAL SERVICE TARGETS

In addition to the likely need to simultaneously support existing legacy services, for which capacity needs continue to expand, and the deployment and parallel growth of advanced video services to IP-based devices, it might make sense to consider additional service targets.

For example, a list of desirable targets could include the following:

- Native support for TDM and IP wireless services, which are currently supported with fiber-based technologies, mainly for now via MetroE using mostly dedicated fiber for the time being, and migrating over time, at least in some cases, to PON technologies.

- Higher-capacity IP-based commercial services than those available via DOCSIS cable modems, which are normally provided via dedicated fiber-based network services (again, MetroE).

- As described in the above sections, an expansion ofboththe upstream and downstream capacity, not just the upstream.

- Continue to simplify operations and increase reliability as all cable-based services have become primary and fundamental for subscribers. For this, fiber deeper into the HFC network has been a goal and driver for quite some time.

LIKELY DESIRABLE PREFERENCES

If industry operators and vendors could just freely do what they thought would be best, especially irrespective of financial considerations, what might they decide to do?

At very first blush one might argue that the most desirable path would be to expand network capacity without any cost or any change in the plant. While the first one could possibly done (i.e., no cost), the second is realistically unlikely (i.e., expand capacity without any change).

So, assuming that some change needs to be made, but still the financial aspects of doing so could be disregarded for the time being, then the most logical outcomes would be:

- Unleash more capacity in the HFC network to expand its use further

- Continue the current path towards smaller service groups since fiber deeper increases capacity and reliability and is already business-as-usual for MSOs

- Leave all current equipment and services unaffected, including leaving STBs in place for the services these provide until a natural/organic transition takes place, and not forcing the removal of analog and/or broadcast channels in favor of deploying DTAs or SDV where not already planned, etc.

- Develop technology that could be deployed on a success basis, incrementally, minimizing cutover changes, and deployed to the extent needed and justified.

- Assign any new DS and/or US spectrum in a flexible way, such as: starting with some spectrum and grow over time as needed, enable future change even as new infrastructure is deployed, change DS:US ratio as needed without service interruption, etc.

OPTIONS BEING CONSIDERED

Let us review the 3 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 opportunitythe details of theoptions and on how these could be deployed.

1. Increase US capacity by moving the US/DS split to a higher portion of the spectrum and simultaneously expand DS to 1 GHz

   From an equipment perspective, this option is generally readily available to MSOs. 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.

2. Additional US via top-split, up to 1 or ~1.2 GHz

   Unlike the 1st option, this approach involves equipment not currently available. Instead, the preferred 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 or 860 MHz and up to 1 or ~1.2 GHz, for US capacity.This new equipment could be built in the form of a new network gateway that would be installed in the vicinity of the node, which would provide the 'translation' from optical transmission from the headend up to the node location into electrical signals, and RF transmission from the node location through the coaxial portion of the HFC network. Additionally, this option would require the deployment of WDM equipment where not already installed to enable the use of the existing node backhaul fiber for communication back to the HE.

This approach would increase US capacity considerably, likely providing an additional 1 Gbps of net additional US bandwidth. In the process it leaves legacy services and existing CPE untouched, but it does not provide any additional DS capacity.

In order to implement this approach,new actives and return equipment would be required and some passive changes might be necessary, especially for older 750 MHz systems that may not have been built with 1 GHz capable actives and passives.

Because of the necessary equipment development and plant impacts, this option is not readily available to MSOs for execution as the first option was. In fact, if standardization were desirable, which all parties considerto almost be a requirement, this option requires considerable development before it would be ready to MSOs for deployment.

3. Overlay DS & US network above ~1.2 GHz up to ~1.7 GHz

   Like the 2nd option, 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. Also like the 2nd option, this approach would require the development of a network gateway that would convert signals from electrical to optical to bypass the analog optical link from the headend to the node.

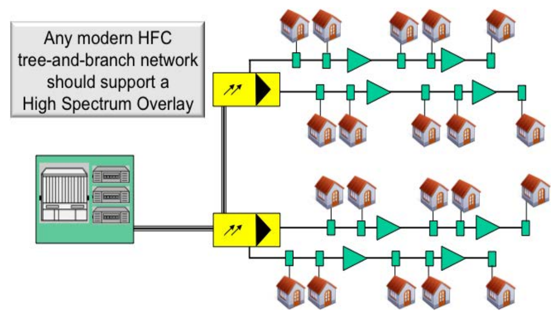

Figure 2: Typical HFC Network

This option could be implemented in two fundamental ways: one where the network gateway is located throughout the HFC network, and the other where the network gateway is deployedin the vicinity of the node.

In the first case, the network gateway would be installed in the vicinity of each active component where advanced services are to be provided. Therefore, this option would require the deploymentof 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.

As shown by Figure 2, any modern HFC network should support a High Spectrum Overlay.

Figure 3 depicts an initial deployment of High Spectrum gateways, for which PON equipment is deployed in the headend, a separate optical wavelength is used in the trunk fiber to carry PON 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.

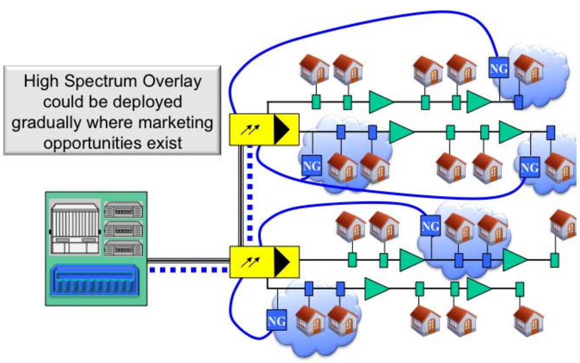

Figure 3: Initial High Spectrum Overlay

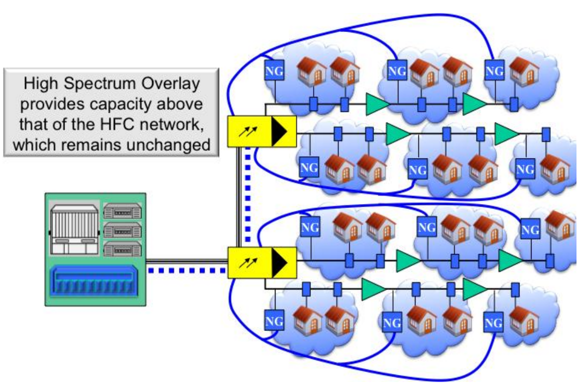

Figure 4: Complete High Spectrum Overlay

In the second case, signals in the new portion of the spectrum would be transmitted through the HFC amplifiers, for which these amplifiers would need to be modified. Figure shows a diagram for this approach.

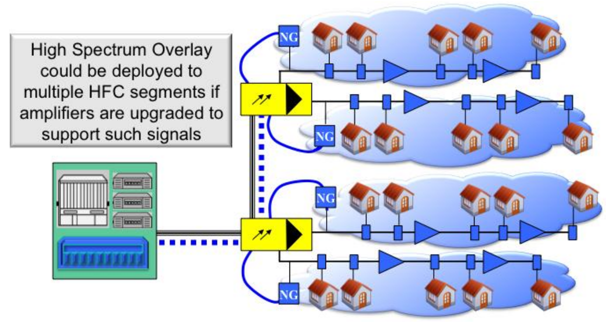

Figure 5: Multi-segment High Spectrum Overlay

Similar toOption 2, Option 3 leaves existing legacy services and current CPE untouched, and would take considerable amount of time to be developed.

However, unlike Option 2, this approach would enable more US andmore DS capacity.

COMPARISON AND APPLICABILITY OF EACH OPTION

For ease of evaluation of the options as versus the previously outlined ‘ideal criteria’, Table 1 below depicts a comparison of each of the 3 options considered above as a function of the ‘likely desirable preferences’ outlined before.

If only additional US is needed, #1 and #2 are good options. Option 1 will likely beless costly and quicker to implement, but could prove to become operationally more complex. Option 2 is probably 2x more expensive that Option 1, but it is operationally easier to deploy and offers greater US capacity.

If instead more DS is needed, then only #1 and #3 are viable. Option 1 is also less costly than option 3, but offers less peak capacity for both US and DS. Instead, Option 3 is probably 4x more expensive than Option 1, but offers significantly more capacity and long-term headroom and is the only option that can be deployed entirely on a success basis. Additionally, it could be argued that Option 3 is the only one that could be implemented with true flexibility in the allocation of US and DS, while Option 2 would be much more restrictive in achieving this goal.

IMPLEMENTATION ALTERNATIVES FOR OPTION 3

While multiple alternatives could exist for implementation of Option 3, only 2 of themare outlined in this paper. In addition, a combination of these 2 alternatives and its benefits is described.

A. Fiber to every active

Option	Added capacity	Path to smaller SG	Legacy unaffected	Success-base deployment	Flexible US/DS allocation
1. Increase US and expand DS to 1 GHz	~ 200 Mbps US	Yes	No	No	No
2. Increase US w/top split	~ 1 Gbps US	Yes	Yes	Partial	Yes
3. Overlay above ~1.2 GHz up to ~1.7 GHz	~ 1 Gbps US plus ~2 Gbps DS	Yes	Yes	Yes	Yes

Table 1: Evaluation of Options vs. Ideal Solution

This first implementation alternative consists of using fiber to the last active to transport US and DS signals optically between the headend and the coaxial part of the HFC network, converting these signals into RF via a network gateway at the location of the last active, combining the resulting RF signals onto the coaxial plant after the active HFC component, distributing thesignals via the cascade of passives to homes corresponding to that portion of the network (e.g., the signals from one network gateway are combined into the coaxial plant after an HFC active and do not traverse the following amplifier in the cascade), and finally terminating the RF signals in a home gateway where the HFC network is bridged onto the home using standard home networking technologies (e.g., Ethernet, MoCA, WiFi).

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 operating between ~1.2 and ~1.7 GHz of the gateways deployed in the HFC network 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 between ~1.2 and ~1.7 GHz on the distribution network, it is likely that a large 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 using this new portion of the spectrum 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 ~6-8 b/Hz might bepossible, therefore resulting in a combined US/DS payload transport capacity exceeding 3 Gbps.

B. Fiber to the node

This second alternative consists of only using fiber to the node, andinstead of over-lashing fiber onto the coaxial distribution hardline RF signals would be used on the coaxial plant from the node through the active components of the HFC network. The same tap faceplate replacement would be required as in the previous implementation alternative.

This approach needs little to no fiber construction, but as a trade-off requires the development and installation of modified actives which are likely complex to develop and expensive to deploy.

Assuming currently available modulation and encoding techniques, and a reduced operating spectrum, an effective transmission of~4-6 b/Hz is expected, which could result in a payload transport capacity of 2-3 Gbps.

C. Combination of alternatives A. and B.

An alternative approach could be that of combining the two alternatives previously described. This could be done for a number of reasons, but perhaps the most important one could be the relative value of each approach. The first approach has the benefit of its increased capacity and lower RF complexity, but has the drawback of requiring the fiber over-lash throughout the coaxial hardline in the HFC network. The second has the converse benefits and drawbacks. Perhaps it might make sense to utilize the first approach where aerial plant exists, and deploy the second approach in cases where underground plant exists. In this way the best performance/cost trade-off can be achieved.

Furthermore, it might make sense to progressively deploy fiber deeper into the network up to the last active in areas where construction costs are relatively high as compared to the cost of the modified actives and their installation. Fiber could then be added progressively as additional capacity is needed over time, likely coinciding with the continuedcascade reduction business-as-usual strategy, eventually reaching the last active, or deploying to last active when additional capacity is needed in the narrowcast portion of the spectrum.

COMBINING OPTIONS

As a further refinement of the approaches suggested, it may make sense to develop a technology strategy that implements the benefits of each option in various stages, and progressively leverages them as these become necessary.

For example, while US capacity beyond 200 – 300 Mbps in the upstream is sufficient and no significant additional DS spectrum is needed, deployments of Option A (e.g., move split to 85 MHz and extend forward spectrum to 1 GHz) might be sufficient, and could be followed by deploying network gateways from Option C using either of he proposed approaches as additional capacity is required.

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 CCAP architecture already supports the modularity necessary to upgrade line cards progressively as new technologies become available. Furthermore, the CCAP architecture provides support for EPON, such that even the Option 3 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.

Option 1 might likely not need silicon development, but the other two options would, for which technology design decisions would be important.

CONCLUSIONS

HFC networks currently have an US capacity limited by the available spectrum, typically limited to the 20-30 MHz of spectrum available between 5 and 30 or 42 MHz in North American systems. At 6.4 MHz of channelization and QAM64 modulation, this capacity allows transmissions of ~100 Mbps.

While US capacity is limited, analysis of DS capacity needs indicates that additional DS capacity will likely be necessary, and perhaps be needed even before additional US capacity is required.

Additional US capacity could be achieved via a change in split to 65 MHz, 85 MHz, or even higher split systems, but these are complex to implement, service impacting and financially and operationally onerous. The capacity gain from such migration is limited to doubling or perhaps tripling the maximum transmission capacity, and has an indirect impact in the reduction of the downstream capacity of the HFC systems.

Taking into account other factors may result in concluding that other solutions may be more appropriate and no more operationally impacting than just changing the spectrum split to increase upstream transmission.

Such other factors include: use a different portion of the spectrum for US,enable additional DS transmission capacity, establish a more efficient HFC transmission mechanism for coupling with home gateway devices, enabling more native support for IP-based, enhanced business, wireless and other services and applications.

When considering theseadditional factors, alternate solutions to just changing the US/DS spectrum split may become more attractive and valuable to MSOs.

This paper presented an analysis of various such options, and expanded on one particular alternative that offers the following characteristics: leaves all current services untouched in the currently allocated HFC spectrum, unleashes new spectrum in the HFC network, and could enable transmission of up to an additional 3 Gbps combined between upstream and downstream.

To do so, the proposed approach makes use of traditional and already available technology and network strategies in a new way, such as pushing fiber deeper up to the current last active, use of existing baseband fiber-based technologies, implementation of existing technologies for higher-order PHY and light-weight MAC, superimpose the new transmission in a portion of the spectrum above the 1 GHzband currently being used, utilizing spectrum allowing the use of passive devices such as below 1.8 GHz, leaving existing services untouched, enables new capacity in the HFC network, solving the current US capacity limitation, and increasingDS capacity.

ACKNOWLEDGEMENTS

Thanks to my colleagues Mitch Bowling, Doug Jones, John Leddy, Saif Rahmanand 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.

ACRONYMS

The following acronyms are used within the paper without being previously defined:

ABR: Adaptive Bit-rate

ARPU:Average Revenue per User

AVC: Advanced Video CODEC

BC: Broadcast

CAS: Conditional Access System

CBR: Constant Bit-rate

CCAP: Converged Cable Access Platform

DOCSIS: Data over cable service interface specification

DTA: Digital transport adapter

Gbps: Gigabit per second

GHz: Gigahertz

GigE: Gigabit Ethernet

HE: Headend
HFC: Hybrid fiber-coax

HSD: High-speed data

MAC: Media Access Protocol

MetroE: Metro-Ethernet

MHz: Megahertz

MPEG: Moving Picture Experts Group

MSO: Multiple system operator

NC: Narrowcast

nDVR: Network DVR, sometimes referred to as RS-DVR for remote storage DVR

OOB: Out of Band

PHY: Physical

PON/EPON: Passive optical network/Ethernet passive optical network

QAM: Quadrature amplitude modulation

RGU: Revenue Generating Unit

RF: Radio frequency

SDV: Switched digital video

STB: Set-top Box

VBR: Variable Bit-rate

VOD: Video on-demand

WDM: Wave Division Multiplexing

---

[1] Presented at Cable Congress 2011 by Arris

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

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

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