Trends in Cable Network Economics

Implications for Public Policy

David Reed

Interdisciplinary Telecom Program
University of Colorado Boulder
May 16, 2016

Authors Note: This is a DRAFT manuscript for discussion purposes only at the Silicon Flatirons INTX Academic workshop. This paper is not for citation without permission of the author.

David P. Reed

Interdisciplinary Telecom Program
University of Colorado Boulder
May 16, 2016

Trends in Cable Network Economics

Implications for Public Policy

Abstract

Introduction

Over the past 50 years the network economics of cable television networks engendered a profound impact upon the diversity of information available to modern society in the form of video and broadband content. Cost-effective one-way transport of analog and digital technology has increased the diversity of television channels to several hundred from a handful. Cost effective two-way transport of Internet access services over cable in substantial part established societal reliance upon always-on access to Internet content at broadband speeds. In response policy makers have enacted many laws and regulations over this time designed to promote and preserve open access to information either specifically or generally targeted to the flow of content over cable networks, with the Federal Communications Commission (FCC) Open Internet decision in 2015 the latest prominent example.[1]

Incremental technological innovation over time has been the key driver establishing favorable network economics to push this evolution of expanding services and information carried on cable networks. As the cable industry now sits upon the cusp of deploying yet another new generation of transport technologies and the new Open Internet rules start to take shape, it is an opportune setting to examine how well the network economic equations enabled by new IP transport technologies are likely to align with the principles and objectives of the FCC's Open Internet framework.

To answer this research question, this paper reviews how the economics associated with the allocation of cable capacity for video and data services has evolved over the past 20 years, and where it is likely to go in the next 5 years. This paper argues that significant past and

---

[1] In the Matter of Protecting and Promoting the Open Internet, Report and Order, FCC 15-24 (2015). [hereinafter "2015 Open Internet Order"]

ongoing changes in the economic framework for allocation of network resources over cable networks have been, and will continue to be, caused by significant new technological innovations in digital television and Internet Protocol based ("IP-based") transport. To accomplish this, we develop an economic framework to illustrate how the efficient allocation of network capacity for analog cable service can be described, and the changes introduced to the framework with the migration to digital television technology. Next, the paper provides a brief explanation of how the new IP transport technologies of Data Over Cable Service Interface Specification (DOCSIS) version 3.1 and network management technologies of software defined networking (SDN) and network virtual function (NFV) are likely to be used over cable networks. These new technologies will likely spark another significant change in the economic calculus or network economics associated with the efficient allocation of capacity over cable networks.

Based upon these anticipated changes, the paper next examines the potential policy implications with a particular focus upon the 2015 Open Internet order. The allocation framework helps to identify policy issues posed by the use of new IP transport capabilities. A preliminary finding is that as video becomes the dominant application carried by the Internet, the Open Internet framework will likely need to address the network access issues specific to IP video including the wholesale migration of video to IP, definition of services using QoS and reasonable approaches for managing public Internet congestion. such as the principles and rules addressing no paid prioritization and specialized services. The analysis also shows how the economics of new channel deployment over cable will become easier in the IP environment, but more complex with regard to the network operations associated with quality.

Literature Review

The amount of network capacity over cable television networks has been an ongoing point of public attention since their initial deployments in the 1950s due to its important role in the overall telecommunications infrastructure. Consumers look to the number of television channels carried for television service or the overall speed of broadband service as important inputs into their deliberations in selecting their TV and Internet services. Policy makers have considered the associated channel capacity of cable networks in considering public policy issues associated with broadcast retransmission, vertical integration, open access, and Open Internet proceedings to name but a few.

Starting in 1989 the cable industry deployed hybrid-fiber coax (HFC) networks that have benefitted from the large information carrying capacity of fiber optic and coaxial cable.[2] HFC networks allow cable operators to incrementally invest in new networking technologies that have significantly increased network capacity since this time.

To increase capacity cable operators generally have extended the fiber optic portion of the network closer to the home. This has occurred in incremental fashion over time, starting with a pair of fibers approaching to within a mile of several thousands households (e.g., a fiber node size of 5000) to reaching within 1000 feet of homes creating node sizes of 100 – 250 homes.

This progression of network "upgrades" extending fiber deeper into the network probably has been the largest driver of additional capacity over time. A significant body of knowledge on the incremental technological advances allowing this growth in cable network capacity can be found in the archives of the cable industry association over the past 30 years, though readers should note that these papers are often authored by cable industry vendors without peer review.[3] Likewise, due to the heavily regulated environment from which cable has emerged, many technical reports and studies have been produced on cable capacity that support of the public policy positions of interested parties.[4]

The topic of how the economic equation for deciding how the bandwidth has been or should be allocated over cable networks, both for new video channels within the video service as well as between different services carried over the network, has not been addressed in any great detail. While there is a presumption that the revenues associated with new channels and services need to outweigh existing revenue streams or network costs, there have been few, if any, allocation models published describing this economic calculus with any great detail.

---

[2] See J. A. Chiddix and W. S. Ciciora, "Introduction of optical fiber transmission technology into existing cable television networks and its impact on the consumer electronics interface," in *IEEE Transactions on Consumer Electronics*, vol. 35, no. 2, pp. 51-62, May 1989.

[3] The technical papers from the National Cable and Telecommunications Association (NCTA) can be found at http://www.nctatechnicalpapers.com/.

[4] See, for example, Steven J. Crowley, "Capacity Trends in Direct Broadcast Satellite and Cable Television Services" prepared by the National Association of Broadcasters, et al, October 8, 2013.

Economic Framework for Cable Network Allocation

As technology has evolved over time, what have been the incentives for cable operators in making decisions to carry new additional television channels or to add new services? Technological innovation and progress clearly have expanded cable network capacity. Yet as we will see, while this growth in capacity may have reduced the opportunity cost of adding additional content, it has not simplified the economic calculus for making this decision. If anything, the growing convergence in the digital transport of services serves to reduce the transparency and ability to precisely predict the impact, or opportunity cost of each additional video channel or new service carried by the network.

This section of the paper reviews how the economic framework for adding cable channels has evolved over the past 20 years, starting with analog television channels in the late 1990s.

Good Old Days: Analog Television Economics

A little over 20 years ago, most cable operators were still only offering their customers cable television service based on analog technology. The decision facing these cable operators to introduce a single, analog television channel could then be illustrated using a relatively straightforward economic framework that compared the cost of the network to the anticipated revenues of each channel.

In an effort to develop a framework reflecting this situation, we assume each channel occupies a fixed 6 MHz of spectrum. Thus, the number of channel slots can be estimated to be the overall downstream capacity divided by 6 MHz.[5] If we assume the network is used only for analog television services, then the total fixed and operating costs of the network is fully allocated to the television service. Dividing the total annualized network cost (by annualizing all fixed costs and then summing with all the operating costs) by the total number of channels supported by the network capacity gives the average total cost per channel. This figure represents a rough benchmark that the revenues of each channel need to exceed to justify the cost of network capacity to support each channel.

On the revenue side, each channel generates subscription revenues from the customers of the cable operators and advertising revenues from advertisers. The net total from these two revenue streams should offset the cost of network transport. Each channel, of course, generates different subscription and advertising revenues based upon the popularity and genre of its programming. Due to this variation it is possible to rank the different cable channels according to the generated amount of net revenue.

---

[5] For example, a low-split 550 MHz cable system typically used the frequencies between 54 – 550 MHz for downstream transmissions, for a total of 82 channels.

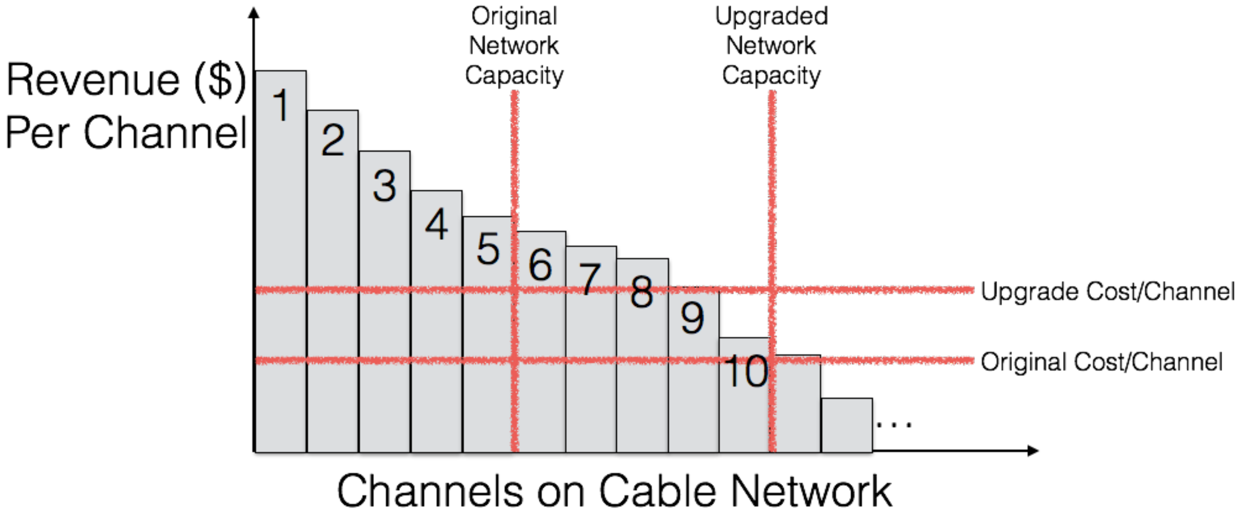

Figure 1: Allocation Framework for Analog Video

Figure 1 provides an illustration of this rank-ordering framework. In this hypothetical scenario, the net revenues of channels carried by the cable network are illustrated, with Channel 1 generating the most revenues, followed by Channel 2, *etc*. The total cost per channel of the network establishes a cost benchmark that the net revenues of all carried channels should exceed. In this case, the original network capacity is 5 channels and the network carries the set of 5 channels that exceed the average cost per channel benchmark and offer the most revenues per channel. Note that some channels may provide revenues that exceed the cost benchmark but are not part of the video service due to a lack of network capacity. How does the operator decide to expand capacity to carry more channels? Adding cable network capacity is usually "lumpy" in that cable operators cannot sequentially add a single channel as needed but instead must increase capacity through network "upgrades" that add 30 or more 6 MHz channels to the overall capacity (e.g., adding 200 MHz of network capacity that moves a 550 MHz network to 750 MHz). The network upgrade cost establishes a new cost benchmark that identifies the channels to be added to the video service.[6]

---

[6] Note the assumption in Figure 1 that the network cost per channel benchmark increases with the upgrade. This may not occur in all cases. Generally speaking with cable networks, as the network capacity increases closer to the home, the cost benchmark also will increase.

Complexity Introduced: Digital Television

The transition from analog television to digital television over cable networks, as well as the ability to compress digital video streams to reduce the overall bit rate, has been well covered in numerous scholarly works over the past 20 years and will not be covered here. What is important for this analysis is to list the technical characteristics of digital television and encoding which have caused major changes in the network economics of cable networks:

1. Can support improved picture quality through the introduction of standard definition television (SDTV) and high definition television (HDTV) resolution formats.
2. Allows for a much higher number of video channels to be carried for every 6 MHz of bandwidth. Depending upon the content of each video channel, a single 6 MHz channel over cable can carry 8 – 15 SDTV or 3 – 5 HDTV programs.[7]
3. Takes advantage of new standards for encoding through the Motion Picture Experts Group (MPEG) that have been realizing a reduction of about 50% in video bit rate with every new generation of the standard (while maintaining picture quality).

What is the allocation framework given these new characteristics of video transport? During the time frame of 2000 – 2015, a major impact of digital video technology was the need to support different video formats due to the wide variation in the type of television receivers found in consumer households.

1. Legacy analog format to support the base of analog television receivers used by their customers.
2. SDTV format for the most affordable digital TV sets during the early years of this time period.
3. HDTV format for the increasing number of HDTV sets that emerged over this time period as the cost of these sets plummeted.

In the late 1990s high-speed data service emerged as a core service element for the cable "triple play" of telephone, video and broadband services. Broadband service required a separate channel allocation that has been growing steadily over time since its inception.[8]

---

[7] The payload of a typical 6 MHz channel over cable using 256 QAM (Quadrature Amplitude Modulation) is roughly 38 Mbps. Thus, 10 SDTV channels could be carried at an average of 3.8 Mbps each, or 4 HDTV channels at 9.5 Mbps apiece.

[8] Most cable operators began offering telephone service during this time that required a small amount of high-speed data bandwidth. A few operators used "circuit-switched" telephone service that required separate, dedicated bandwidth to operate, but this approach was quickly scrapped in favor of the voice over IP (VoIP) approach employed today.

We can apply the same rank-ordering framework described above for the analog scenario in Figure 1, though we require some modifications due to the different transport characteristics of digital television. The network now is used for analog and digital television services, along with telephone and broadband service. The common fixed and operating costs of the network must therefore be shared across all these services. Dividing the total annualized network common cost allocated to each service, plus any service-specific fixed and operating costs, by the total number of channels supported by the network for the service gives the average total cost per channel.

For video services the same cost benchmark based upon the annualized cost per 6 MHz channel can be applied, though it has to be recognized that multiple video programs can be carried per 6 MHz channel. The overall impact of this is that the network cost per program, as opposed to the network cost per channel, now falls considerably depending upon the number of SDTV and HDTV programs carried per 6 MHz channel. On the revenue side each channel again generates subscription revenues from the customers of cable operators and advertising revenues from advertisers, but these revenues streams are separated based upon the technical format used by the customer (advertising is usually the same across each format). So now the net total from these revenue streams should offset the cost of network transport for all video formats.

For broadband services the network cost per channel benchmark used for video can be applied to reflect the opportunity cost for using some network capacity to carry broadband service. As long as the revenues generated by broadband service exceed the network cost per channel multiplied by the number of channels dedicated to carry broadband, then the use of the capacity is efficient as compared to using the capacity for additional video services.

Figure 2 provides an illustration of the new allocation framework as applied with the delivery of digital television over cable. In this hypothetical scenario, the network uses 70 channels to carry analog video, 15 channels to carry 150 SDTV programs, and 25 channels to carry 75 HDTV programs. This is a typical cable band plan for a 750 MHz HFC network circa 2005 – 2010. The problem here for the cable operator is that the revenue calculation has become much more complex. Each major cable channel might be carried in three separate formats since the operator risks losing customers by not carrying each format. While the revenues per channel can be rank-ordered within each format category, and then summed across the three categories, the operator may have contractual constraints in being able to keep one format while eliminating another. Also note that the lower network cost per program benchmark can afford the cable operator additional opportunity to add new digital programs assuming there is enough network capacity.

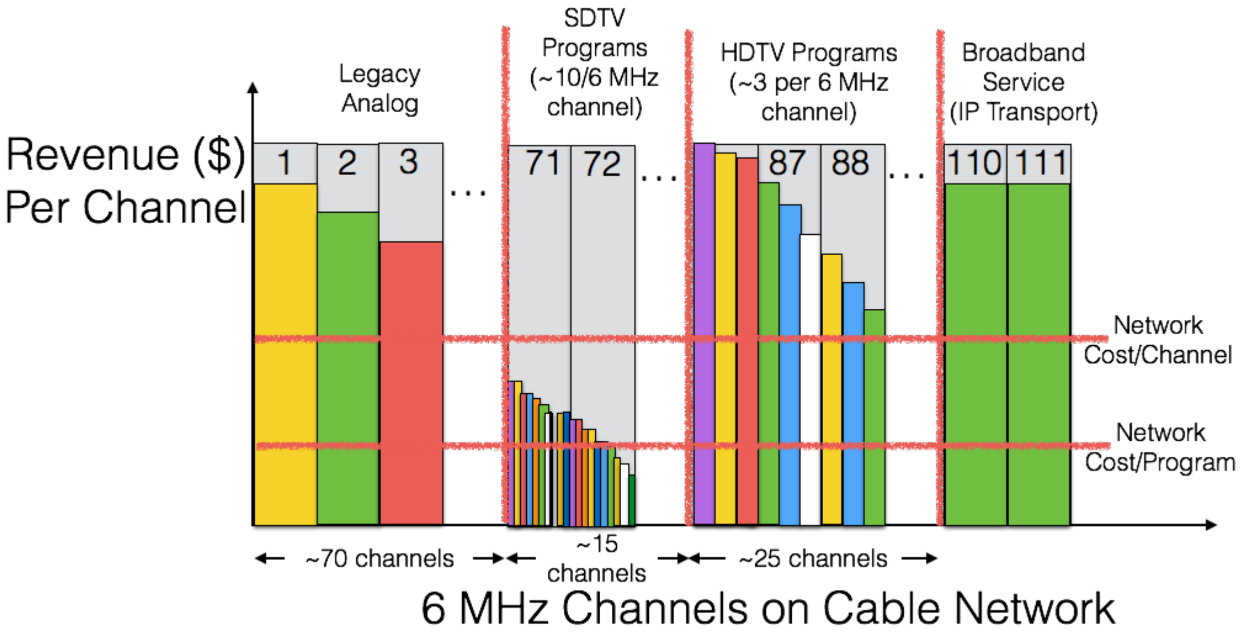

Figure 2: Allocation Framework for Analog and Digital Video

Two interesting trends emerged during this time frame that further changed the network economics and shifted the overall band plan away from analog to more HDTV:

1. Low-cost digital devices allowed operators to deploy cheap digital-to-analog (DTA) adapters eliminating the need to carry the analog television format.
2. Emergence of low-cost HDTV receivers reduced the size and demand for the SDTV platform as well.

As television viewing fully migrates to HDTV over time, it would appear initially that the economic equation for video channel selection should return to a simpler state after the transition from analog television and lower use of SDTV. The problem with this view, however, is that the next format for television known as "4K" is already emerging, and the transition to IP transport from the MPEG-based QAM is occurring as well. Moreover, the amount of network capacity now required for broadband services is increasing sharply. As discussed next, the transition to IP promises to simplify some aspects of video carriage, and introduce further complexities given the flexibility afforded by IP transport.

Finally, over the past 5 years observe that the allocation for broadband service as shown in Figure 2 has expanded from 2 channels (6 MHz apiece) to 4 – 8 channels. The use of DOCSIS 3.0 channel bonding technology permits cable operators to combine up to 24 separate 6 MHz channels together to offer higher speed broadband services above 100

Mbps.[9] To make economic sense, the revenues of existing and new broadband services need to exceed the network cost per channel benchmark adjusted to reflect the additional cost of channel bonding. Note that the revenue calculation may become more complex given the flexibility of the IP platform to offer multiple sources of revenue based upon existing and new speed tiers for broadband services, advertising revenues associated with broadband services and IP transport services for businesses.

Integrated Flexibility: IP Transport

Today the migration to IP technology is fully affecting the entire video ecosystem. Video service providers are now looking to use IP transport because the IP-based networking equipment is the lowest cost and the proliferation of IP-based consumer devices makes it easier for service providers to provide video content to any device at any location and time.[10] A full description of the transition to IP for the cable architecture is beyond the scope of this paper. Instead, we will focus upon the emergence and impact of two significant emerging IP transport innovations, namely DOCSIS 3.1 and SDN/NFV network management technology. Because how these innovations will be used over cable is less well understood at this point in time, we will first explain how these technologies are being developed, and proposals for how they will be used over cable in the next 5 years.

DOCSIS 3.1: The End of 6 MHz Channelization Legacy

The latest version of the DOCSIS specification is DOCSIS 3.1.[11] Wide-scale deployment of this new technology is likely to commence in 2017.[12] DOCSIS 3.1 contains several interface changes that hold significant implications for the amount of channel capacity available for IP transport over cable networks. The key changes relevant to this study are listed below:

---

[9] For example, combining 4 256-QAM channels (38 Mbps apiece) allows cable operators to offer broadband service at speeds above 100 Mbps. See Om Malik, "Finally, 100 Mbps Everywhere (If You Have Comcast)," April 14, 2011, https://gigaom.com/2011/04/14/finally-100-mbps-everywhere-if-you-have-comcast/.

[10] The migration is underway even though the transport efficiency of IP video (less than 85%) is actually much less than MPEG-2 transport (97%) over DOCSIS. See "Streaming over HFC: MPEG-2 or IP or Both?" accessed April 22, 2016, http://www.nctatechnicalpapers.com/Paper/2000/2000-streaming-over-hfc-mpeg-2-or-ip-or-both-/download.

[11] The suite of three DOCSIS 3.1 specifications is located at www/cableabs.com. Previous DOCSIS specification versions include 1.0, 1.1, 2.0, and 3.0. DOCSIS 3.1 equipment will be backward compatible with the previous DOCSIS 3.0 version equipment.

[12] See "In D3.1 First, Comcast Goes Gig in Atlanta," *Light Reading*, 3/15/2016, accessed April 7, 2016, http://www.lightreading.com/cable/docsis/in-d31-first-comcast-goes-gig-in-atlanta/d/d-id/721876.

1. *New Modulation*. DOCSIS 3.1 uses Orthogonal Frequency Division Multiplexing (OFDM) in the downstream and Orthogonal Frequency Division Multiplexing Access (OFDMA) in the upstream. This new multiplexing approach employs separate discrete subcarriers that are only 25 or 50 KHz wide but do not require guard bands between each other or the use of 6 MHz channels. The guard bands required between OFDM subcarriers are very small as compared to legacy technology.
2. *Larger Channel Size and More Spectrum*. In fact DOCSIS 3.1 affords operators considerable flexibility to define channel size between 24 – 192 MHz in the downstream and 6.4 – 96 MHz in the upstream. The larger channel size increases multiplexing efficiency, and sharply reduces the use of spectrum as guard bands between channels.[13] The specification also increases the amount of upstream spectrum to the range of 5 – 204 MHz and downstream range of 258 – 1218 MHz (optional up to 1794 MHz).
3. *Higher Order Modulation*. The highest order modulation formats supported by DOCSIS 3.0 were 64-QAM and 256-QAM for upstream and downstream, respectively, with each delivering a payload throughput of roughly 27 Mbps and 38 Mbps over a 6 MHz channel. DOCSIS 3.1 adds up to 4096-QAM (optional up to 16384-QAM) for the downstream, and up to 1024-QAM (optional up to 4096-QAM) in the upstream. To compare, the payload throughput of 4096 QAM over 6 MHz is about 54 Mbps. DOCSIS 3.1 also adopts more powerful error coding techniques that allow more use of higher order modulation formats.
4. *Modulation Flexibility*. A key benefit of OFDM/A is that all of the subcarriers can have a different modulation order.[14] In the past, the poorest part of the worst 6 MHz channel for the high-speed data service would dictate the modulation order employed for the full service. Modulation order decreases with an increasing noise environment encountered, so the ability to work around the noisy portions of spectrum with less efficient, but more robust, subcarriers allows for substantial boost in overall network speed of up to 40%.[15]

---

[13] See, for example, http://www.incognito.com/blog/a-technical-guide-to-docsis-3-1-and-beyond-part-2-technical-advantages/ estimating a savings of 6% in spectrum (1.44 MHz) per 24 MHz of OFDM space through less us of spectrum for guard bands.

[14] Wikipedia defines the modulation order as the number of different symbols that can be transmitted during a fixed period using a digital communication scheme. Using this definition, a higher order modulation scheme can transmit more symbols per fixed period of time, meaning more information is transported over the communication channel.

[15] See, for example, http://www.incognito.com/blog/a-technical-guide-to-docsis-3-1-and-beyond-part-2-technical-advantages/ estimating a 35.8% improvement in network efficiency due to variable modulation profiles.

5. *Hierachical Quality of Service (HQoS)*. Historically quality of service (QoS) over DOCSIS has been defined and managed as individual service flows over the HFC network.[16] HQoS enables cable operators to aggregate together groups of service flows to better manage and control QoS policies at the aggregate level. This permits, for example, the network operator to assure fairness for each of the different aggregated service flows (i.e., each aggregated service flow receives a percentage of overall traffic). A key use case for HQoS is the ability to manage the per subscriber aggregate QoS traffic.[17] With the migration to all IP, subscribers will have several different flows to support their different applications (e.g., different video, voice and video conferencing applications), each of which might have QoS associated with the service flow. HQoS permits the operator to better manage the allocation of network resources in the event of contention.[18]

The upshot of migration to DOCSIS 3.1 is that cable operators will have much more IP capacity on their networks to support the migration of video services to IP transport. IP-based video services over cable will still employ the same MPEG encoding, but will use different transport protocols developed to stream video programs over the Internet. Most major cable operators are choosing to deploy DOCSIS 3.1 in deep fiber configurations. This very likely could be the last incremental, major upgrade to HFC plant before fiber is extended all the way to the home.[19]

---

[16] Wikipedia defines quality of service as "the overall performance of a telephony or computer network, particularly the performance seen by the users of the network." For a useful discussion of the definition of QoS and its use on the Internet, see K. C. Claffy and David D. Clark, "Adding Enhanced Services to the Internet: Lessons from History," SSRN Scholarly Paper (Rochester, NY: Social Science Research Network, September 7, 2015), http://papers.ssrn.com/abstract=2587262.

[17] See "Hierarchical QoS, The Next Step In Enhancing DOCSIS QoS Architecture - NCTA Technical Papers," accessed April 22, 2016, http://www.nctatechnicalpapers.com/Paper/2013/2013-hierarchical-qos-the-next-step-in-enhancing-docsis-qos-architecture.

[18] *Ibid*, p. 4. For example, a household might subscribe to a 20 Mbps broadband service and a separate video service with a maximum traffic rate of 9 Mbps that has higher priority. The aggregate QoS settings for the home might limit the combined traffic rate of both flows to 20 Mbps, which means the whole 20 Mbps would be available for the broadband service when the video service is not in use, but this could fall to 11 Mbps when the video service is fully active. Without an aggregate QoS limit, both services would run separately and could consume bandwidth up to 29 Mbps, which could be an inefficient use of network resources.

[19] While still early in terms of planning, some cable operators are stating that they prefer upgrading to fiber-to-the-home networks instead of deploying DOCSIS 3.1. See "Why D3.1 Isn't on Every Cableco Agenda," *Light Reading*, 3/3/2016, accessed April 7, 2016,

Coming back to how DOCSIS 3.1 might change the economic equation for allocation of cable bandwidth to different services, the following observations are important:

- The concept of a 6 MHz channel no longer applies. Very large swaths of cable spectrum (up to 192 MHz apiece) can be provisioned to deliver IP bandwidth.
- The cost benchmark for cable capacity therefore transitions from a "cost per channel or program" metric to a "cost per Mbps" measure for the *application* being carried over the network.
- The cost of IP transport is highly sensitive to the latency requirements of the application. Thus, a QoS regime will likely be necessary to prioritize the traffic needs of real-time applications. This means the cost per Mbps metric needs to reflect the anticipated level of priority required by the application.

Before applying these insights into a new economic framework for cable allocation, we also look at another important development in IP transport – the trend to software-centric networking – that will heavily influence the new model.

Network Management Using SDN and NFV

An in depth explanation of software defined networking (SDN) and network function virtualization (NFV) architecture and technology is beyond the scope of this paper.[20] The focus here is to describe how these new technologies might be applied over cable, and the implications for the economics of capacity allocation.[21]

Briefly, SDN is network management technology that moves the control of network resources (e.g., routers and switches) through a standardized interface to a new element called a "controller" providing centralized control. The abstraction of the SDN control layer between the application and network resource layers is achieved through standardized

---

http://www.lightreading.com/cable/docsis/why-d31-isnt-on-every-cableco-agenda/d/d-id/721514.

[20] See, generally, ITU, "Framework of software-defined networking", ITU Recommendation Y.3300, June 2014, <http://www.itu.int/rec/T-REC-Y.3300-201406-I/en>, and ETSI, "Network Functions Virtualisation (NFV): Architectural Framework", ETSI GS NFV 002, October 2013, <http://www.etsi.org/deliver/etsi_gs/nfv/001_099/002/01.01.01_60/gs_nfv002v010101p.pdf>.

[21] Some of the information in this section was derived from research conducted during the Fall 2015 semester by Gauri Kanitkar, a master's student in the Interdisciplinary Telecom Program at CU Boulder.

interfaces or application programming interfaces (APIs).[22] Some key benefits anticipated by using SDN include:

- *Lower cost equipment*. Networking equipment will be lower cost due to simpler devices without the need for control logic. SDN utilizes so-called "white box switches" that are commodity devices customized for SDN applications and not related to any costly proprietary brand or functionality. For cable, this benefit is available for use of SDN technology in the backbone and middle mile portions of their networks, but less so in the local access network where DOCSIS equipment is customized for the HFC network configuration that is a smaller, less commoditized market.

- *Ease of provisioning new services and maintenance*. New service provisioning and maintenance of network devices becomes fast and inexpensive using only downloadable and deployable software. Maintenance costs of networks are comparably less as software can be updated with minimal downtime of a network.

NFV technology decouples software from hardware, meaning that network elements no longer have to be an integrated hardware and software entity. Detachment of the software from hardware allows for hardware resources to be shared by multiple software entities, which permits better scaling and utilization of network resources overall. Using NFV, network functions are migrated from proprietary hardware to a virtual platform without performing all network functions on the hardware device. These network functions can be networking devices such as firewalls, routers and switches or networking applications like content delivery networks (CDNs) and voice over IP (VoIP). Some key benefits anticipated by using NFV include:

- *Lower cost*. NFV fundamentally changes the way networks are built by incorporating virtualization to improve capital efficiencies. Network operators can use low-cost, general-purpose servers and storage devices to provide multiple network functions using software virtualization techniques. For cable, this benefit is available for use of NFV technology mainly in the network cloud through the efficient provision of new and existing cloud services on a virtual platform. NFV helps to lower capital expenses by reducing the need for dedicated, proprietary hardware implementations;

---

[22] These standardized APIs removes the need for network administrators to have to learn the syntaxes for all the propriety devices. Instead, s/he can configure these devices using commonly defined API calls. Because these APIs are implemented in software, the growing usage of SDN applications is associated with the trend to "software-centric" networking.

it also helps in lower operational expenses on space, power, cooling and hardware maintenance through more efficient use of hardware.

- *Rapid deployment*. NFVs can be deployed quickly on a virtual platform while deployment of hardware equipment is time consuming. Deployment of hardware devices involves multiple steps like ordering, purchasing, deploying, configuring, and purchasing support. NFV often eliminates most of these steps.

SDN and NFV technologies provide different but complementary functionalities. SDN is a technology by which the control plane and forwarding plane of devices can be separated. This gives benefits like centralization of control and ability to program the behavior of devices using well-known APIs. NFV is a technology that is used to speed up the deployment of network services and help in reducing capital and operating expenses. NFV technology aims at consolidating networking devices into industry standard high volume servers and deploying them in data centers and customer premises.

In summary, the ways that SDN and NFV are likely to change the economic equation for allocation of cable bandwidth are the following:

- Rapid deployment of new cloud services can be achieved at very low cost.
- Agile support and deployment of new IP services and applications such as video, telephony and video conferencing featuring different levels of quality of service are possible.
- New network management techniques that optimize the speed of broadband service to each customer based upon the quality of the received signal.[23] Depending upon implementation, the amount of spectrum consumed by each customer to receive a given level of speed could vary depending upon the geographic location and quality of the nearby network. Alternatively, the overall speed available to cable broadband customers theoretically could vary depending upon these same factors, similar to the speeds provided by ADSL today.

---

[23] See Saifur Rahman, Joe Solomon, Jason Schnitzer, and David Early, "DOCSIS 3.1 Overdrive: Dynamic Optimization Using a Programmable Physical Layer," forthcoming INTX paper, May 2016.

- Use of congestion management techniques that implement dynamic, rule-based policy controls to manage access to network resources during times of network congestion.[24]

Figure 3 shows how the cable allocation may evolve with the use of DOCSIS 3.1, SDN and NFV technologies. Now, the 750 MHz HFC network in this hypothetical scenario is fully digital, with 32 channels (each 6 MHz) carrying legacy HDTV video service and the remainder of the network being used for IP transport to send streaming video, telephony and video conferencing applications, as well as high-speed broadband services. Other notable differences in the model:

1. The revenue metric is changed from $/channel to $/MHz to reflect the decreasing use of 6-MHz channelization in the band plan.
2. The cost benchmark for HDTV channels remains as the average cost per 6 MHz channel for the portion of bandwidth used for HDTV service, as in prior models. More advanced compression and more broadband service capacity will likely further lower the cost benchmark for video programs to facilitate a large number of broadcast and interactive video service options.
3. Using DOCSIS 3.1 the broadband service portion of the spectrum is no longer divided into 6 MHz channels, and instead is split into large blocks of spectrum up to 192 MHz in size. Such a block size can carry between 1.2 – 1.7 Gbps depending upon the modulation order.
4. Lacking 6 MHz channelization, the cost benchmark for the broadband service block of spectrum is the annualized average cost per MHz for the portion of bandwidth used for broadband service. Note, however, that this metric presumes all costs associated with providing broadband service at all times, including the busy hour of usage. To the extent all traffic is best effort, this would distribute IP transport costs uniformly across all applications creating traffic. Because QoS is applied to some traffic, the true cost of QoS could be underestimated using this metric. In this case a metric that allocates cost across the aggregated service flows would be a more accurate reflection of the distribution of IP transport costs across applications.

---

[24] See Samuel Patel, Mohammad Chowdhury, Jason Schnitzer, and David Early, “SDN Ground Truth: Implementing a Massive Scale Programmable DOCSIS Network,” forthcoming INTX paper, May 2016.

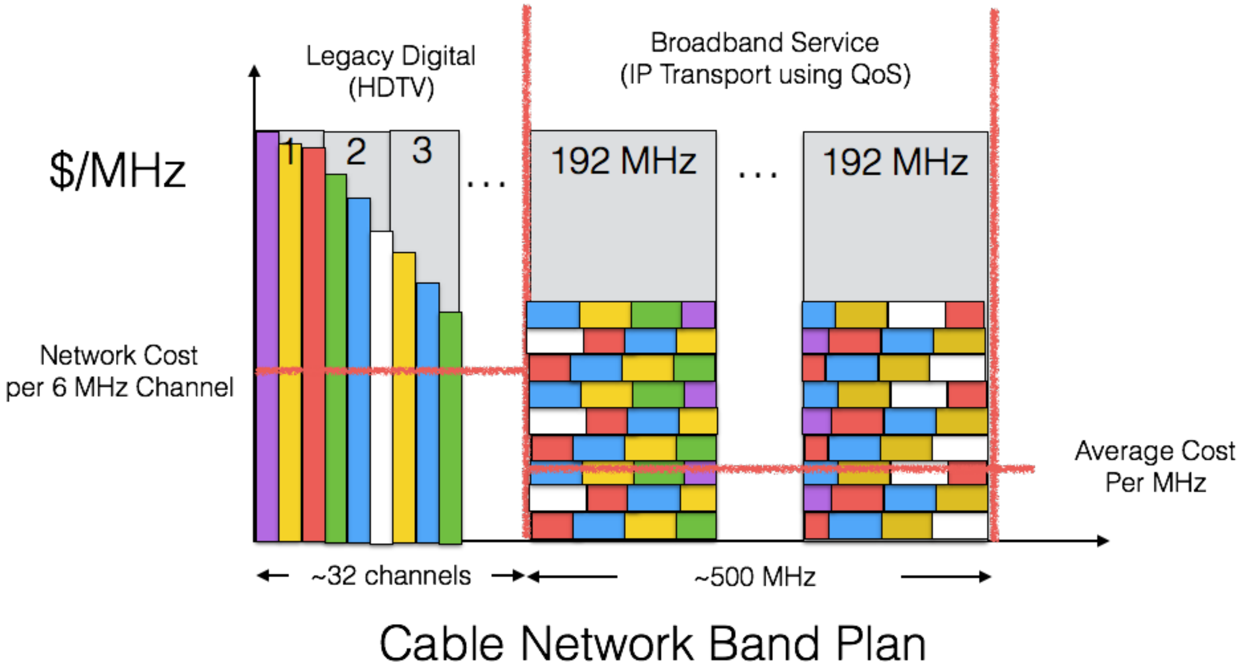

Figure 3: Allocation Framework for IP Transport

In the long run most cable operators expect to migrate to an "all IP" network platform where all network capacity uses IP transport using QoS to differentiate service flows amongst the different applications carried over the network. In this scenario the cable allocation model becomes simpler in that it uniformly consists of only IP transport, but simultaneously more complex in the management of network resources according to the level of priority afforded each service flow being carried over the IP network. This will lead to the creation of a "public" Internet as well as "private" Internet transport over the same network, sometimes with traffic from both categories being carried over the same networking equipment. How much QoS is utilized and serves as a differentiator in consumer experiences will depend upon the amount of congestion experienced on the network.

This analysis illustrates how more efficient IP transport using DOCSIS 3.1 and the new network management technologies of SDN and NFV will facilitate changes in the services and applications being carried over local access networks providing residential broadband services. From a policy perspective, the primary guideposts in place today to direct this transition can be found in the 2015 Open Internet Order, which we turn to discuss now to better understand the policy implications of the new network economics likely to emerge over cable networks over the next 5 years.

Public Policy Implications

This examination of the evolution of how capacity is allocated over cable networks over the past 50 years has been instructive by establishing how the network economics associated with cable networks is largely driven by underlying video and transport technologies. Using a simple model, this analysis showed how the economic calculus for adding new television channels and services has evolved with the transition from analog to digital video formats, followed by the current transition to IP transport. Going forward, the cable networks are migrating to become a single, shared IP "pipeline" with the use of QoS to differentiate, or prioritize, the many different service flows being carried over the network. Modern network management technologies such as SDN and NFV are moving to implement this vision with highly agile and dynamic software-centric solutions.

What are the public policy implications of this transition? Claffy and Clark in large part have provided many insights to this question, beginning with the notion that network neutrality emerged in part out of concerns that QoS can benefit services with stringent transport requirements, though potentially at the expense of other service flows on the network.[25] Claffy and Clark further provide an overview of the general policy issues generated by the use of QoS over the Internet in the current environment. Within the context of this study, some of these issues include:

- Identification of the potential incentive by network operators to defer investment in overall broadband capacity in order to grow revenues for managed services that will work better using QoS in the face of congestion. [26]
- The need for transparency in the application of reasonable network management techniques through disclosure and evaluation of the risks and benefits associated with QoS. [27]
- Identification of the FCC's prohibition on paid priority in the 2015 Open Internet order, noting the potential problem for QoS tools to be used mostly on private networks since it cannot be used on the public internet due to this ban. [28]

Building upon Claffy and Clark's discussion, this paper looks at three policy topics upon which this analysis provides some unique insights including: 1) the implications of the increasing migration of video to IP, 2) the need for clearer definition of services using QoS and 3) overall management of public Internet congestion.

---

[25] See Claffy and Clark, *op. cit.*, p. 17-18.
[26] See Claffy and Clark, *op. cit.*, p. 18.
[27] See Claffy and Clark, *op. cit.*, p. 20.
[28] See Claffy and Clark, *op. cit.*, p. 21-22.

Video over IP

Two main drivers of the transition of video services to IP transport will be: 1) through use of over-the-top (OTT) services and applications delivered over the public Internet, and 2) increasing use of IP video by Multichannel Video Programming Distributors (MVPDs) to transport their current video services currently using the IP protocol in private networks.

OTT services use IP-based streaming video protocols over broadband Internet access service (BIAS) to reach customers. Most streaming video applications use adaptive streaming protocols, meaning that the speed of the download is adjusted to best utilize the bandwidth available over the path of connection. If the application estimates additional bandwidth is available in the network path, it will switch to a higher-resolution copy of the video that will use more bandwidth but also improve the user experience; alternatively, if the application senses network congestion over its path, it will shift to a lower resolution copy to reduce the speed of the download but also improve the user experience by reducing buffering delays.[29] One upshot of the heavy use of adaptive streaming protocols is that it makes it very difficult for the network operator to manage network congestion given the constant adjustments occurring in the streaming services. This may create network management issues over time as networks become more congested due to increasing video usage.

IP video services use IP transport to carry the video to cable customers. The reason that cable operators are moving this direction is that the cost advantages of IP transport are now making this technology the lowest cost option to build and operate. No major cable operators are providing pure IP video today for all their cable television services, so it is not clear how this service will be offered. The operator will need to consider how it may offer the service without violating the Open Internet rules. One model that it might follow could be the “specialized services” category defined in FCC’s 2010 Open Internet order. Specialized services are defined to be services that are not part of BIAS, and the migration of the existing video platform to IP transport might fit this definition if appropriately disclosed to the FCC.[30] The FCC has indicated that it would oppose a specialized service if

---

[29] In adaptive streaming, if the video player determines that the download speed over the network is greater than the bit rate of the segment downloaded, it will request the next higher bit rate segments. If the client finds the download speed for a segment is lower than the bit rate for the segment, it will request a lower bit rate segment.

[30] The FCC’s example of a specialized service is the VoIP service provided by cable operators. Cable VoIP typically applies QoS for telephone calls using the same DOCSIS infrastructure used for broadband service. The telephone application does require some different network facilities to manage the phone call, though transport facilities are generally shared to some degree.

it finds that the service is "providing the functional equivalent of broadband Internet access service or is being used to evade the open Internet rules."[31]

With all the trends in technology pointing to a future migration of most video services to IP transport, some additional policy implications become clear:

1. The cost of building and maintaining IP networks will likely be the lowest cost option for providing video services. This will cause an increasing number of MVPDs to transition to IP to transport their video services, while at the same time many OTT video services will be using BIAS to deliver their services, often from BIAS offered by the same MVPD. The FCC states that the 2015 Open Internet rules only apply to IP services that travel over BIAS, which will increasingly include OTT video services. The potential for confusion among policy makers, industry and consumers seems high given these circumstances.
2. Refinement to the reasonable network management standard may need to occur due to the increasing use of OTT video services employing adaptive video streaming protocols. The benefit of these protocols is that they permit OTT providers to maximize the viewing quality of their video programs in automated fashion. The downside is that the protocols can have a significant demand for bandwidth that can be inefficient.[32] We can anticipate the OTT and MVPD space to become more competitive over time as the growing popularity of OTT services continues. As this competition intensifies, it is not unreasonable to anticipate this the quality of experience of streaming video consumers will be an area of concern. The key issue to resolve will be to establish expectations for the expected service level agreement (SLA) associated with best-effort BIAS.

Defining QoS-Based Services

If the forecasted migration to all IP occurs over time as described above, then IP transport technology will be used as the dominant transport platform for cable networks over the next decade. Whether or not the use of QoS will be a prominent attribute of this platform will depend upon two conditions: 1) the amount of congestion experienced as part of BIAS, and 2) the ability of the operators to deploy specialized services as described above.

---

[31] See 2015 Open Internet Order, *op. cit.*, para. 35.

[32] For example, a lab experiment I give students is to measure the download speed of a popular OTT service to different devices including a smartphone, tablet, and 50" HDTV. The students find that the download speed is the same for all devices, even though the large variation in screen size means that the smartphone could be delivered a much lower speed copy of the video program than the HDTV without any impact on the viewing experience.

With the continuing growth in popularity of OTT video services driving more viewing of videos over BIAS, it will be a safe assumption to forecast continuing strong growth in broadband usage for the foreseeable future.

Whether or not specialized services offer a good model going forward is an open question. At the heart of the issue is the definition of specialized services – namely, that "they are not BIAS." If one interprets this to mean to limit a specialized service to a specific IP service or application, then it conceivably would permit the cable operator to deploy a number of different specialized services, each carrying a specific service flow associated with a disclosed service or application such as broadcast video, on-demand video and network DVR services. Each specialized service would have a prioritized QoS attached to its service flow of packets over the network.

There may be a further issue to address, however, if this approach were to occur. Given the paid prioritization ban that would prevent the operator from receiving any revenues for specialized services from third parties (and this type of IP transport would not be given away), then the only likely beneficiary from specialized services would be the operator itself. The regulations themselves would create this outcome by preventing any market to form for IP transport using QoS. To avoid this outcome, further definition of how QoS can fit within the definition of BIAS needs to occur.

Avoiding Public Internet Congestion

Claffy and Clark raise a concern that the emergence of the use of IP transport for private networks or specialized services might create a "shadow" activity outside the regulated sector that may hinder policy makers' ability to respond to regulatory issues in ways that may cast the public Internet as a "place for activities of lower importance and lower profitability".[33]

These concerns are to be taken seriously if we are to insure the continued benefits of the Open Internet virtuous cycle. QoS functionality is an important technical differentiator in user experience in congested network conditions. The benefits of QoS have to be balanced against the implied priority of usage for a service flow over other service flows. To date, the FCC has yet to wrestle with this specific question, and as mentioned earlier there is some concern that the paid prioritization ban could have some unintended consequences on the development of specialized services or private networks.

---

[33] See Claffy and Clark, *op. cit.*, p. 22.

To date, the FCC has largely left the provision and interconnection of IP transport services to market forces. The result of this approach has been the establishment of a successful IP transport infrastructure for the United States, though one missing element of the system has been the establishment of a functioning end-to-end QoS regime. As a result, to optimize control of network resources, some broadband application providers have built their own private content distribution networks (CDNs) in part to give their users the best possible experience for their applications.[34] In effect, this way is a topological approach to managing network quality by avoiding a long series of network interconnection points that might not support QoS in a unified or efficient manner. For these providers, CDNs provide complete control of the resource policies of the network.

What this analysis shows is that the transition to IP video may serve as the case study that forces the FCC to consider the treatment of QoS under the Open Internet regulatory framework.[35] IP transport and network management technologies have advanced far enough to start serious planning for the migration of cable to a fully IP transport network. The transition to video will also cause broadband usage to dramatically increase such that the broadband capacity will continue to experience some congestion pressure.

To meet this migration, an Open Internet framework that recognizes the benefits of openness and QoS will be needed. Edge providers and broadband customers should be reassured through disclosure and transparency rules that the Open Internet will continue; likewise network operators should be able to deploy efficient IP technology to carry their services without fear of regulation or concern that they are undermining the Open Internet with these actions.

---

[34] See David P. Reed, Donny Warbritton, and Douglas Sicker, "Current Trends and Controversies in Internet Peering and Transit: Implications for the Future Evolution of the Internet," SSRN Scholarly Paper (Rochester, NY: Social Science Research Network, August 20, 2014), http://papers.ssrn.com/abstract=2418770.

[35] Indeed, this process might already have started as the FCC sent letters to Comcast and T-Mobile to learn more about their "zero-rating" video services in January of 2016. See Jon Brodkin, "FCC Had 'productive' Net Neutrality Talks with Comcast and T-Mobile," *Ars Technica*, January 15, 2016, http://arstechnica.com/business/2016/01/fcc-had-productive-net-neutrality-talks-with-comcast-att-t-mobile/.