

Lehr, William; Sicker, Douglas

Conference Paper

Would you like your Internet with or without video?

27th European Regional Conference of the International Telecommunications Society (ITS): "The Evolution of the North-South Telecommunications Divide: The Role for Europe", Cambridge, United Kingdom, 7th-9th September, 2016

Provided in Cooperation with:

International Telecommunications Society (ITS)

Suggested Citation: Lehr, William; Sicker, Douglas (2016) : Would you like your Internet with or without video?, 27th European Regional Conference of the International Telecommunications Society (ITS): "The Evolution of the North-South Telecommunications Divide: The Role for Europe", Cambridge, United Kingdom, 7th-9th September, 2016, International Telecommunications Society (ITS), Calgary

This Version is available at:

<https://hdl.handle.net/10419/148683>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Would you like your Internet with or without video?

William Lehr¹
MIT

Douglas Sicker²
CMU

**** Draft -- Comments Welcome****

Abstract

According to Cisco's VNI forecast, "consumer Internet video traffic will be 85 percent of all consumer Internet traffic in 2020, up from 76 percent in 2015," and the majority of this traffic will be entertainment-oriented video. Many might view this as the (near) realization of the promised convergence of digital broadband delivery platforms that has been coming since first generation broadband services started becoming available in the mid-1990s. A question we should ask is whether this is the Internet we want? Even if one concludes that the marriage between entertainment media and the Internet is a foregone conclusion, it is worthwhile to consider what this may mean for the design, regulation, and economics of the Internet.

In this paper, we critically examine the proposition that the conventional wisdom that convergence toward "everything over IP," or even stronger, "everything over the Internet," is efficient, inevitable, or desirable may be wrong. Convergence means different things in technical, economic, and policy terms. Building a single network that is optimized for 80% entertainment video traffic might disadvantage other services. Moreover, the economics of media entertainment are distinct from, and potentially in conflict with, the economics motivating many of the usage cases most often cited as justification for viewing the Internet as an essential infrastructure. Finally, separately managing the traffic for Internet and video services may be advantageous in addressing regulatory agenda items such as performance measurement, set-top boxes, universal service, OVD reclassification, and Internet interconnection. While most of the traffic may share the same physical (principally, wired) conduit into homes, it may be more efficient and flexible to segregate traffic into multiple logically distinct networks; and doing so may facilitate technical, market, and regulatory management of the shared resources.

¹ wlehr@mit.edu. Corresponding Author.

² sicker@cmu.edu.

Table of Contents

1. Introduction	3
2. Background Notes on Convergence and Working Assumptions	4
3. Technical	11
3.1. Background.....	11
3.2. Traffic Trends	14
3.2.1. Video Trends: traffic, growth and resolution	14
3.2.2. Traffic types/categories	16
3.3. Evolving the network for video	16
3.4. Traffic Percentages and a Future Internet.....	18
3.4.1. Current Traffic Composition of Broadband Internet Access	18
3.4.2. Usage Scenarios of the 20%	19
3.4.3. Videos Impact on Other Services	20
3.4.4. Differences of Voice and Video as Specialized Services.....	22
3.5. Architecture	23
3.5.1. Getting video to the customer.....	25
3.5.2. VIAS.....	27
3.6. Summary.....	29
4. Economics	30
4.1. Benefits of Bundling Services	31
4.2. Entertainment Industry Economics and Broadband	36
5. Policy	42
5.1. Separate policy domains and the challenge of Convergence	42
5.1.1. PSTN Regulation.....	42
5.1.2. Broadcast, Cable and Media Regulation	46
5.1.3. Data Communications and Internet Regulation	48
5.2. Relevance to Current Policy Issues	49
5.2.1. Performance measurement and managing consumer expectations	50
5.2.2. Universal service	50
5.2.3. Internet Interconnections	51
5.2.4. MVPD Regulatory Reform.....	52
5.2.5. Set-top boxes	53
5.2.6. Summing up the policy challenges.....	54
5.3. A framework for separating Broadband and Video	54
6. Summary Conclusions	56
7. Exhibits (landscape)	57
7.1. Figure 1: Qualitative Characterization of Video Traffic Types	57
7.2. Figure 2: Entertainment Video versus “Other 20%” traffic characterized	58

1. Introduction

According to Cisco's VNI forecast, "consumer Internet video traffic will be 85 percent of all consumer Internet traffic in 2020, up from 76 percent in 2015," and the majority of this traffic will be entertainment-oriented video.³ Many might view this as the (near) realization of the promised convergence of digital broadband delivery platforms that has been coming since first generation broadband services started becoming available in the mid-1990s. A question we should ask is whether this is the Internet we want? Should we assume that our video, voice or other services will inevitably come to us over a single Internet connection, and is such convergence desirable? How might the design, economics, and regulatory policy for the Internet diverge if this video traffic were not part of the Internet's future?

This paper engages these questions and addresses the deeper underlying question of what should be the optimal convergence path for our digital communications infrastructure, at least with respect to our last-mile access services. Convergence means different things in technical, economic, and policy terms. At the technical level, convergence requires us to provision our IP access networks to handle a mix of traffic that is heavily weighted toward video that is mostly one-way, cacheable, and associated with large-sized files. At the business/industry level, most of this traffic is supported by the economics of media entertainment. Media entertainment, competing for consumer attention and leisure expenditures, is an important contributor to economic activity but the supply/demand concerns are distinct from and, in many cases, only distantly related to other sectors of the economy (e.g., healthcare, government, energy, transportation) that are critical to the growth of our smart, Information and Communications Technology (ICT)-connected economy. The balance of non-media-entertainment traffic that is associated with everything else -- while smaller in aggregate volume -- is more heterogeneous with respect to its requirements for connectivity, message size, and delivery specifications (i.e., routing, data rates, and QoS requirements). Finally, with respect to communications policy, convergence requires us to address the challenge of transitioning from a world of separate regulatory rule-sets for broadcast (TV/mass media), telephone (PSTN), and Internet (which is in the midst of transitioning from mostly deregulated to partially regulated). Each of these rule-sets has sector-specific concerns with little obvious overlap.

In this paper, we critically examine the proposition that the conventional wisdom that convergence toward "everything over IP," or even stronger, "everything over the Internet," is efficient, inevitable, or desirable, may be wrong. We explore ways in which building a

³ See Cisco (2016), "Cisco Visual Networking Index: Forecast and Methodology, 2015-2020," white paper, June 6, 2016 (hereafter, Cisco VNI). The corresponding global shares of traffic that will be Internet video in 2015 and 2020 are 70 and 82 percent respectively. This includes both fixed and mobile video; and ambient video (Nannycams, petcams, and home security cams), but excludes video conferencing/chat and multimedia gaming, which are included in other categories.

single network that is optimized for 80% video traffic might disadvantage other services, as well as limiting options for efficient video distribution. Relying on entertainment industry economics to drive broadband investment may be convenient, or even necessary, but may not best serve to promote the Internet uses that are most often cited as justifying regarding the Internet as an essential basic infrastructure for the economy.⁴ A world in which most of the Internet traffic is *not* entertainment video may deliver greater choice of content, more industry/market structure flexibility, and more competition than a world of converged Internet traffic flows. Finally, separately managing the traffic for Internet and entertainment video services may be advantageous in addressing regulatory agenda items such as performance measurement, set-top boxes, universal service, OVD reclassification, and Internet interconnection. While most of the traffic may share the same physical (principally, wired) conduit into homes, it may be more efficient and flexible to segregate traffic into multiple logically distinct networks; and doing so may facilitate technical, market, and regulatory management of the shared resources.

The balance of this paper is organized into five sections. In Section 2, we further elaborate the motivation and focus of this paper, explaining the various senses in which we plan to examine the convergence question, which we undertake in the following three sections. In Section 3, we discuss prospects for technical convergence, discussing whether separating entertainment video traffic from Internet traffic, which will include everything else that is not the entertainment video stream, might change how we design, provision, and manage broadband access networks. In Section 4, we address the implications of relying on entertainment industry economics to drive broadband investment and the potential tensions that may pose for other anticipated – and generally regarded as more economically productive – uses of the Internet. In Section 5, we focus on the regulatory challenges posed by convergence and discuss ways in which these challenges might be ameliorated by delaying convergence at the regulatory level (even if convergence at the technical and economic levels is largely regarded as a *fait accompli*). We also sketch out one possible regulatory model for managing the traffic separately. Section 6 concludes with summary thoughts.

2. Background Notes on Convergence and Working Assumptions

Anyone following developments in broadband or media entertainment might well ask whether the convergence of broadband and entertainment video is a foregone conclusion. Both mobile and fixed broadband service providers have been expanding their capacity to handle streaming video from services like Netflix, YouTube, and Hulu – and a host of

⁴ We recognize that trying to predict the future of the Internet may be a fool's errand. However, in light of the fact that entertainment video with distinct traffic characteristics accounts for the dominant share of Internet broadband access traffic today, and it is certainly possible, that this will continue to be true into the foreseeable future, we believe it is worthwhile considering how this might be impacting the architecture, market economics, and regulatory policy for the Internet.

others.⁵ According to Sandvine, as of December 2015, 70% of peak Internet traffic consisted of real-time streaming entertainment, up significantly from the 13% share reported for 2008.⁶ The proliferation of multiple screens per user and per household capable of accessing Internet-delivered programming anywhere,⁷ the expanding selection of programming available in multiple quality formats (SD, HD, and ultraHD),⁸ and the vertical integration of programming and broadband service providers⁹ are all trends that have been in evidence for years. Collectively, these and related trends have contributed to the rapid growth of the entertainment video traffic that has induced broadband providers to expand capacity and increase per-subscriber data rates.¹⁰ Whereas much of the

⁵ Most of the traditional programming providers are launching services to access video content on-line over mobile and fixed platforms. Web-based and/or application (usually mobile) streaming is available from most broadcasters of over-the-air and PPV programming (NBC, CBS, ABC, Fox, Comedy Central, ESPN, HBO, Showtime), most providers of broadband access services (Xfinity for Comcast, FiOS TV for Verizon, etc.), as well as host of alternative and/or user-generated streaming sources (YouTube, Twitch, Meerkat, Periscope, BitTorrent). Although these sources also provide streaming news and other programming that some might not identify as “entertainment,” most of this includes entertainment programming.

⁶ Sandvine reported that real-time entertainment traffic accounted for the following shares of peak traffic: 12.6% (2008), 26.6% (2009), and 70.4% (2015) (see Sandvine reports for 2009 and December 2015). Real-time entertainment traffic includes streaming video and audio entertainment, but the bulk of the traffic is video.

⁷ The iPhone inaugurated the consumer smartphone revolution in 2007 and big screen LCD TVs that cost north of \$4,000 were costing less than \$2,000 by 2005 (see http://money.cnn.com/2010/09/23/technology/lcd_tv_prices/). Increasingly, households have and are using multiple screen devices to consumer video content. A 2015 Accenture survey found 40% of consumers owned a tablet, laptop/desktop, and a smartphone; 16% also owned a connected TV; and 89% accessed long-form video content over the Internet (see Accenture (2015), “Digital Video and the Connected Consumer,” available at https://www.accenture.com/t20150523T021027__w_/us-en/_acnmedia/Accenture/Conversion-Assets/Microsites/Documents17/Accenture-Digital-Video-Connected-Consumer.pdf).

⁸ The resolution at the high-end continues to increase, with a growing volume of programming available in HD (high-definition) or ultra HD (4K) formats. At the same time, expanded options for SD (standard-definition) and other lower quality (lower resolution) format programming (e.g., animation, user-generated content) is being made available. Many popular programs are available in multiple formats to accommodate different subscriber equipment needs and viewing contexts.

⁹ The announcement that Comcast is acquiring Dreamworks and the FCC is expected to approve the AT&T-DirectTV merger are just two recent examples. The Comcast-NBC merger was consummated in 2011. Moreover, as of January 2016, SNL Kagan reports that 20% of MVPDs are now providing access to over-the-top Internet video content (e.g., Netflix, YouTube) as part of their regular video offers (see <https://www.snl.com/InteractiveX/Article.aspx?cid=A-35020212-15157>).

¹⁰ Peak and average data rates for fixed and mobile services have increased year-to-year. For example, FCC reported that the average peak rate was 31Mbps in September 2014, up from 10Mbps in 2011 (see <http://www.enterprisenetworkingplanet.com/netsp/u.s-broadband-speeds->

discussion touting the Internet's importance as essential infrastructure has focused on the role of the Internet in enabling "Smart X" (where X can be replaced with energy grids, homes, cities, healthcare, supply chains, education, etc.),¹¹ most of the traffic growth has been from video entertainment.

Although entertainment video already comprises the dominant source of traffic on our broadband networks, and market trends and strategic plans by key value chain participants are propelling us further in that direction, it is worthwhile to consider what a counter-factual world might look like: one in which the Internet carried the *other 20%* of non-entertainment-video traffic; and the 80% of the traffic that comprises entertainment video were delivered to homes via a separate access network service.¹²

Our focus here will be on the "last-mile" access networks that provide fixed communication services (Internet, video, and telephone) to consumer households – the so-called "eye-ball" networks operated by AT&T, Charter, Comcast, and Verizon,¹³ because these have been the principal focus of communications policy concern and remain the principal Internet "on-ramps" for most consumer households.¹⁴ In focusing on fixed access networks we are already accepting a significant degree of convergence, at least at the physical layer or with respect to much of the network and business facilities that will be shared by the Internet and such other services as may be delivered over these last-mile networks.¹⁵ (Mobile broadband, although growing rapidly still accounts for only

accelerate.html); which is significantly higher than the Akamai estimated average peak data rate in the U.S. of 11.9Mbps (see Akamai, SOTI December 2015).

¹¹ The National Broadband Plan touting the importance of broadband for the Nation's future has chapters devoted to discussing the role of broadband in healthcare, education, energy and the environment; in promoting economic opportunity and civic engagement, in enhancing government performance and public safety. Although it mentions "entertainment" as a common consumer use, it does not emphasize this as a key reason for why the Nation needs a national broadband plan. (See FCC, "Connecting America: the National Broadband Plan," Federal Communications Commission, Washington, DC, March 16, 2010).

¹² These shares are based on the VNI forecast cited earlier, see Note 3 *supra*.

¹³ There are also many smaller broadband access providers that collectively provide services using a variety of network technologies, including FTTH, HFC, xDSL, and wireless (terrestrial and satellite).

¹⁴ In so doing, we are mostly ignoring the network services, including Internet services, provided to businesses, government and anchor institutions (e.g., libraries, schools) that typically share much of the local network infrastructure and collectively account for significant (mostly non-entertainment) Internet traffic. We are also mostly ignoring those portions of the Internet that are upstream of the access networks (in the Internet "cloud") or downstream (part of the consumer's home network).

¹⁵ We note that for the foreseeable future, there will be multiple facilities-based network infrastructures capable of serving entertainment video to broadband households (i.e., offering competitive alternatives to whatever is available over a households broadband Internet access service). For most households, this will include at a minimum the legacy ILECs and cable

a small share of total Internet traffic,¹⁶ and for most subscribers remains a complement to fixed broadband service.¹⁷)

Furthermore, we recognize that the *other 20%* is an approximate estimate of the traffic that is not comprised of entertainment video. First, some of the entertainment video may be downloaded as files from on-line stores (e.g., iTunes or Amazon) or via file-sharing applications like BitTorrent (that include a large share of copyright infringing media) and may not be appropriately identified as video entertainment or separately identifiable in the relevant data sources.¹⁸ Second, video traffic may be associated with many non-entertainment applications such as video conferencing, surveillance/monitoring, or integrated into interactive multimedia applications.¹⁹ Third, the category "entertainment

company multiservice networks, as well as multiple mobile networks. There are also single-service network alternatives such as direct broadcast satellites and over-the-air television. Obviously, these networks are not fully independent (e.g., when the mobile and fixed networks are owned by the same provider) and offer different user-quality experiences. Although the implications of multimodal competition in delivery platforms is of great interest to policymakers, we will not focus on those issues here beyond noting that one argument in favor of delivering entertainment video via the Internet is precisely because of the potential for the Internet to serve as a "spanning layer" to support interoperability across these different entertainment video infrastructures.

¹⁶ According to Cisco VNI, in 2014, global IP traffic was 59.8 EB/month. Of that, 4% was mobile; 80% was consumer; and 36% was consumer IP video.

¹⁷ Fixed broadband access services with Wi-Fi handle much of the mobile data traffic already; and with the trend toward smaller wireless cells (driven in part by spectrum scarcity), wired infrastructure is increasingly important for backhauling traffic in local distribution networks (see Lehr, W. and M. Oliver (2014), "Small cells and the mobile broadband ecosystem," Euro ITS2014, Brussels, June 2014, available at <http://econpapers.repec.org/paper/zbwitse14/101406.htm>; or Chapin, J. and W. Lehr (2011) "Mobile Broadband Growth, Spectrum Scarcity, and Sustainable Competition," TPRC 2011, available at SSRN: <http://ssrn.com/abstract=1992423>.) Regarding point that fixed and mobile broadband are more likely to be complements for most subscribers, see Horrigan, J. (2016), "Smartphones and Home Broadband Subscriptions: Substitutes, Complements, or Something Else?," draft INTX workshop, May 2016; or Lehr, W. (2009), "Mobile Broadband and Implications for Broadband Competition and Adoption," available at SSRN: <http://ssrn.com/abstract=2446011>.

¹⁸ All available traffic and share estimates are imprecise due to data availability issues. The Cisco VNI, although widely cited, relies on a plethora of assumptions to develop its aggregate estimates. It is unclear how the Cisco VNI treats file sharing or purchases of digital media files from online vendors like iTunes. In any case, some of those files are likely to be encrypted and not identifiable as video. The estimates from Sandvine, Akamai and some of the other sources cited herein focus on real-time streaming of entertainment media, which includes audio. Although music files comprise a significant share of the file counts, the files are typically much smaller and video remains the dominant form of traffic.

¹⁹ Although as noted before, the Cisco VNI estimates exclude video-conferencing and multimedia gaming from the estimates of Internet video (see Note 3 *supra*).

video" includes content targeted at many different audiences and purposes such as news, home-movies (user-generated content not intended for mass audience distribution), educational videos -- as well as what might more generally be regarded as traditional leisure-time entertainment video. As we shall explain further in Section 4, distinguishing between entertainment video or Internet uses and non-entertainment uses (education, job-search, eCommerce, or Smart-X uses) is inherently ambiguous.²⁰

Our focus is on this latter class of video traffic which we believe comprises the bulk of the video traffic today and is likely to continue to do so in at least the near to midterm future. This is the traffic that is presently delivered to consumers as linear "television-like" or video-on-demand services as part of the video program offerings that comprise one leg of the triple-play bundle of services (Internet, telephony, and video) sold by the eyeball network operators.²¹ We recognize that with the rise of "New Media," definitions of what constitutes video entertainment are changing. This is perhaps most evident with respect to the user-generated content that is often produced with smartphone cameras and webcams, and is distributed via services like Periscope, YouTube, Vine, and Snapchat.²² A lot of this is shorter-form or involves real-time broadcasting from diverse locations. The ways in which it is produced, distributed, and consumed challenge traditional notions of television or entertainment video. Those developments notwithstanding, we believe that more traditionally produced types of video entertainment will continue to comprise a significant share of traffic; and in some cases, the new media channels are themselves becoming portals for access to legacy-format content.²³

Our goal is not to precisely estimate the share of total traffic that is comprised of this type of entertainment video, but rather to characterize in general terms what a future might look like where a very significant share of the traffic (upwards of two thirds in volume) is not carried as part of the broadband Internet service, but is delivered to homes over a logically (and capacity-isolated) separate access network. In most cases, we expect that the video services may be provided over the same physical infrastructure, as is the case

²⁰ For example, whether having three channels showing basketball, football, and fishing provides more diversity than three different football games depends on how one assesses diversity.

²¹ By "television-like" services we mean programming that may be delivered as linear programming (in which the viewer receives a stream of content that is organized into a sequenced "channel" that may be recorded to facilitate time-switching but is otherwise not decomposable by the user in real-time) or video-on-demand (where the user can select the particular video content to view). One justification for moving from traditional television to Internet-based delivery of television programming is to better enable video-on-demand.

²² For example, see Klym, N. (2015), "The Ambiguity of Disruption: Discovering the Future of Video Content," MIT Communications Futures Program white paper, September 2015, available at <http://cfp.mit.edu/Ambiguity%20of%20Disruption%20Klym%20Sept%202015.pdf>.

²³ For example, YouTube and other new media outlets are now offering channel programming that is not materially different from watching traditional television channels from a user experience perspective.

today in cable and telephone-based access network architectures, which begs the question of how separate the networks may actually be.

With Internet Quality-of-Service (QoS) technologies (e.g., DiffServ or MPLS) and with modern cable modem technologies (e.g., DOCSIS3.x),²⁴ it is feasible to prioritize packets and allow multiple types and flows of traffic to share the same physical network infrastructure, while allowing the traffic to be segregated into separately managed flows, or equivalently, logically separate networks. These technologies may be configured to allocate underlying network resources to share capacity between the logical flows in a variety of ways. For example, capacity may be shared among logical flows dynamically on a real-time basis; or alternatively, the flows may be allocated hard capacity assignments, mimicking what would happen if the traffic were segregated onto separate physical networks. While we will discuss the implications of capacity sharing, for much of the discussion it will be useful to imagine that the traffic is logically and capacity separated on distinct IP networks, one of which will carry the "other 20%" of broadband Internet traffic and one that will carry the entertainment video in the access network.²⁵

²⁴ There are multiple QoS mechanisms available on the Internet, and all of the modern broadband technologies have expanded capabilities to allow more fine-grained and flexible resource assignment to support QoS requirements for diverse traffic types, including video. For an explanation of DiffServ (short for Differentiated Services), see <https://technet.microsoft.com/en-us/library/cc787218%28v=ws.10%29.aspx>; for MPLS (short for Multi-Protocol Label Switching) see <http://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/4649-mpls-faq-4649.html>. <https://www.nanog.org/meetings/nanog49/presentations/Sunday/mpls-nanog49.pdf>; and for DOCSIS (short for Data Over Cable Service Interface Specification) which is a group of standardized technologies developed by CableLabs to support digital services over cable modem networks. The most recent standard, DOCSIS 3.1, provides a host of features to enable fine-grained, dynamic and flexible control of IP-based services over the hybrid coaxial-fiber-based networks operated by cable providers. For more information about DOCSIS, see <http://www.cablelabs.com/?s=DOCSIS>. For a discussion of how these will significantly increase the capacity and capabilities of cable-based broadband access networks to deliver increasing volumes of broadband traffic that includes lots of entertainment video, see Reed, D. (2016), "Trends in Cable Network Economics: Implications for Public Policy," draft paper presented at NCTA INTX Workshop, May 2016. Similar capabilities are available with other broadband access network technologies such as FiOS and VDSL.

For further discussion of these QoS mechanisms, see BITAG (2015), "Differentiated Treatment of Internet Traffic," Broadband Internet Technical Advisory Group (BITAG), October 2015, available at http://www.bitag.org/documents/BITAG_-_Differentiated_Treatment_of_Internet_Traffic.pdf; and BITAG (2013), "Real-time Network Management of Internet Congestion," Broadband Internet Technical Advisory Group (BITAG), available at http://www.bitag.org/documents/BITAG_-_Congestion_Management_Report.pdf.

²⁵ As we will explain further below, if any resources are shared (wires, IP routing, back-office support, conduit, etc.), it is generally not feasible to fully separate decision-making about the networks at the technical, economic, or policy level (by nature of the sharing). See, for example, Knieps, G. and V. Stocker (2016), "Price and QoS Differentiation in all-IP Networks." *International Journal of Management and Network Economics*, forthcoming.

From an economic demand perspective, asking the question of what might be different if entertainment video traffic is separate from Internet traffic, might strike some as either uninteresting or foolish. One might argue that attempts to forecast future Internet traffic requirements is a fool's errand because of the inability to predict what tomorrow's "Killer Application" may be. Not very long ago, we thought Peer-to-Peer (P2P) services like Napster and BitTorrent would dominate traffic flows, resulting in much more symmetric two-way flows. Although today (asymmetric) streaming entertainment video is the dominant traffic type, in the future, the traffic of some new application may become dominant.²⁶ Much of the economic benefit of the Internet as a "General Purpose Technology" (GPT)²⁷ stems from its ability to support traffic with diverse QoS requirements that are well-represented by video traffic, which may span a broad range of requirements (e.g., in terms of data rates, latency, MBs, or direction of flows).²⁸ While the future remains uncertain and a key component of the Internet's value proposition is to be a GPT, it seems hard to dispute that in the foreseeable future entertainment video is likely to remain the dominant traffic driver *and* the convergence of video entertainment and the Internet may bias the direction of the Internet's design away from being a GPT.

Another potential critique of the focus of this paper is that the marriage between entertainment and the Internet is a practical necessity since consumers desire entertainment. Infrastructure investment is being driven by market demand. Indeed, most of the investment that underpins the Internet was undertaken to satisfy the demand for ubiquitously available telephone and cable television services – not Internet access. Providing an Internet that mostly delivers entertainment video today allows us to build a highly capable broadband network that could not be sustained from an investment perspective without this traffic. One might fully accept this argument and still find the question we address in this paper interesting – if only as a thought experiment. On the other hand, the implicit explanation that convergence is a necessary byproduct of market evolution because that has been the case in prior generations of telecommunications infrastructure investment is hardly conclusive. We might determine that the social benefits of deploying the Internet we want could warrant supporting its investment with public funds if the only alternative is to rely on entertainment industry revenues. Alternatively, we might conclude that having entertainment revenues fund a significant share of broadband investment does not require convergence and may be consistent with managing entertainment and Internet traffic separately.

²⁶ Of course, if it is just another form of entertainment like gaming or entertainment-based Virtual Reality – the network requirements may change substantially, while the economic implications may change much less.

²⁷ See Helpman, E. (ed) (1998), General Purpose Technologies and Economic Growth, MIT Press: Cambridge, MA.

²⁸ Video comes in many formats and may differ widely in terms of its cache-ability (e.g., real-time sports or breaking news versus movies). The rise of interactive media and user-generated content, as well as non-entertainment-based, video-requiring applications such as video-conferencing and video-monitoring services may require support for symmetric and dynamic routing control.

Finally, it is worth noting that we view this paper as an exploratory effort. We are not personally persuaded yet that convergence does not offer the most suitable result, but are uncomfortable with what we view as the under-examined presumption that it is desirable or inevitable. Even if we ultimately conclude that a different future might be more desirable, we recognize the potential infeasibility of realizing a better future in light of real world practicalities.

In subsequent sections, we paint a picture of how the world might be different from a technical, economic, and policy perspective if Internet traffic were separate from the entertainment video traffic. That discussion is followed by consideration of possible options for how one might actually accomplish the separation if desired, and some of the issues that would need to be addressed. We conclude with a summary and discussion of future research directions.

3. Technical

In this section, we consider a number of technical questions relating to the role of video over broadband Internet access service. We start by providing background on video services and then look at trends in video services over broadband access. Our goal here is to characterize the (roughly) 80% of access traffic that is entertainment video, and the “other 20%” that includes everything else. We then look at the traffic running over a sampling of broadband access links, and consider what these links might look like if you were to remove the video traffic. This includes considering the impact of these video services on the access link, and the role that network architecture plays. We then close with a brief summary of the technical analysis.

3.1. Background

We start this section by discussing what we mean by video. Video on the Internet constitutes a broad range of services, where the basic uses include video for communications, sensing/monitoring, and entertainment. Figure 1 summarizes our qualitative characterization of these three main categories of video traffic, which we describe below.

Communications: videoconferencing is a key application that calls for bandwidth that supports symmetric, any-to-any real-time connectivity, and might in the future have patterns analogous to telephone calling. This category may also include groupware applications like WebEx, which provide collaborative, on-line, multimedia work platforms. It may include gaming, which can be a mix of characteristics, depending on the nature of the game. While communications video does demand considerable

bandwidth (on the order of 0.5 to 10 Mbps), as we describe in more detail below in “Trends,” at this point it represents much less than 1% of the traffic on access networks.²⁹

Sensing/Monitoring: Video monitoring cameras, which includes webcams,³⁰ are likely to present a wide range of traffic types that may mix cacheable and real-time content (e.g., traffic and security cameras). Unlike the communications uses, a lot of this traffic may be asymmetric and flowing upstream (in the opposite direction from most of the legacy entertainment media).³¹ Additionally, for many sensing applications, it may be possible to significantly vary resolution requirements and real-time access to peak resources.³² Once again, while sensing applications could represent considerable bandwidth in the future, today web-cams and monitoring camera traffic represents less than 1% of traffic on access networks. With the rise of the Internet of Things (IoT), this category could grow significantly, but is unlikely to exceed 5% of access traffic over the next five years.³³

²⁹ The implications for provisioning requirements are significant in light of the requirement to support symmetric, real-time, any-to-any traffic. The “any-to-any” implies that the capacity has to exist everywhere (and in a converged mobile/fixed world, that really may mean everywhere and not just for every possible fixed location). As one moves beyond the portion of the network that is not shared among end-users (a portion of the last-mile), calling behavior will determine opportunities for sharing resources, and the limits to multiplexing. For example, the Erlang distribution found wide application for provisioning legacy telephone networks (e.g., predicting the number of switching ports needed to provide a targeted level of non-blocking telephone calls, and for allocating network costs to services. However, the Erlang model was disrupted when the rise of dial-up Internet access resulted in much longer-holding-times for calls. It is unclear whether a new “Erlang” model for traffic is feasible or desirable in a world with video-conferencing and where the “telephone call” communications abstraction has expanded to encompass such a wide-range of applications and services (from real-time to voice-mail, from SMS to high definition, multi-party video-conferencing).

Additionally, while real-time communication applications limit the potential to cache traffic to economize on peak capacity resources, the desire to record for later playback or review implies that video-conferencing, as well as other applications like sports television that require real-time resources, may also have demand for caching support. (Whether such storage should be provided depends on the content-delivery strategy, and many possible options are feasible and/or desirable, depending on the nature of the content. For a further discussion of the diversity of CDN strategies, see Stocker, Smarakadakis et al., 2016, forthcoming.)

³⁰ Many webcams are not actively monitored, which begs the question as to their purpose and the value of the traffic they generate. They may actually be monitoring something that is potentially of interest (home security, fish tank) or they may be on in lieu of setting up a videoconference call.

³¹ Of course, in a future world of machine-to-machine automation, it is likely that automated communications may inextricably blend communications and sensing functionality and traffic.

³² Many sensing applications are interruptible, depending on what the goal of the sensing is. Time-lapse video recording is often adequate, and real-time access to the recorded video may not be necessary.

³³ Cisco forecasts that machine-to-machine (m2m) connections will rise to represent 46% of total connections, but the traffic will account for only 3% of traffic by 2020 (and only a portion of this

Entertainment: Television and long-form programming (movies) comprise the bulk of entertainment media traffic and the vast majority of this is currently still delivered and consumed via non-IP networks.³⁴ The entertainment-based IP video we are concerned about here includes the above types of content, but is broadly construed to include: entertainment content that is (1) commercial, as well as user-generated (e.g., YouTube); (2) offered by program aggregators (e.g., Netflix, Comcast), as well as specialty networks (e.g., HBO, ESPN); and (3) offered as linear TV where the channelization of programming and timing is set by the programmer and Video-on-Demand (VoD), so that the timing/selection of when a show is viewed is determined by the viewer. Entertainment video includes programming with an array of variable characteristics such as: delivery-time-sensitive (e.g., sports, live performances, and breaking news) and long-term cacheable (e.g., movies); long or short form programming; content available in several resolutions; content subject to varying digital rights management (DRM) treatments; and content that may be delivered via IP or the Internet in multiple ways. Delivery methods include streaming Video-on-Demand (VoD) or subscription services (e.g., Netflix, Hulu, Xfinity), or downloading of bulk files³⁵ at faster-than-real-time (e.g., iTunes downloads).³⁶ From a traffic characterization perspective, this traffic tends to be asymmetric (mostly downstream, requiring only limited upstream capabilities),³⁷ and mostly cacheable.

will be video). See Figure 2 and 3, Cisco (2016), “The Zettabyte Era: Trends and Analysis,” Cisco Virtual Networking Index (VNI), June 2016, available at http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/VNI_Hyperconnectivity_WP.html.

³⁴ Most cable and direct broadcast satellite and terrestrial over-the-air television services are not IP-based today, although this is rapidly changing. Additionally, non-networked entertainment media consumption remains significant (e.g., DVDs, movie theaters).

³⁵ This class of traffic presents a difficult challenge for our analysis since it may be difficult to distinguish from other types of file downloads, which we expect to remain an important component of the “other 20%” traffic. If there were a separate IP service for delivering entertainment video, it is plausible that this might also prove attractive for these sorts of bulk downloads of entertainment video, as we discuss further below.

³⁶ For the most part, we will ignore New Media applications like Periscope or live-video sharing applications being supported by a growing number of social networking services like Facebook, Twitter, and others. While traffic from these services is growing in importance, full consideration of the implications of such traffic is beyond the scope of this paper, and it is unclear how or to what extent such traffic represents a new form of traffic or legacy traffic shifting (or being replicated) in new streams (e.g., when a Periscope subscriber re-broadcasts entertainment video available via other services); and if it is a new form of traffic, how does it impact demand for legacy entertainment?

³⁷ VOD requires an uplink channel for interactivity to select programming, but that could be provided via a distinct and separate network from the network used to deliver the video (e.g., when one orders goods by telephone or the Internet that are delivered via UPS). User-generated content needs one upstream channel to upload the content (for each item), but most of the end-user traffic remains downstream to multiple users downloading content for viewing. For

<<INSERT Figure 1 here >>

In this paper's thought experiment, we are not assuming that all entertainment video traffic will be eliminated from the Internet (or, excluded from the “other 20%”). First, as we have already noted, drawing the boundary between what constitutes entertainment and communications or other types of traffic is not easy, and may become less so with the emergence of new forms of entertainment media. Highly interactive entertainment media may have very different traffic patterns than the sort of legacy television/movie-like entertainment video that is dominant today.³⁸ Second, because the Internet will still support large file downloads (e.g., for updating software or databases, or backing up files), and because many of those files may be video (e.g., iTunes/Amazon rental or purchase downloads), a significant volume of traffic may still be entertainment video. Third, the Internet may be used to support limited interactivity for video control purposes (e.g., basic VoD and VCR functionality for fast-forward, rewind, program selection, and billing). (However, those control functions could be handled via a gateway that would communicate with the entertainment video distribution network, if the entertainment and other Internet traffic were indeed to be separated). Furthermore, were entertainment video to be carried via a logically separate IP distribution network or service, that service might also prove attractive for other types of traffic that shares the general features of video traffic (e.g., software downloads).

3.2. Traffic Trends

In this section we discuss how the character of video traffic has been changing as a consequence of changes in traffic composition, program resolution, and the types of video available.

3.2.1. Video Trends: traffic, growth and resolution

Data from Sandvine's Global Internet Phenomena Report finds that (1) Entertainment traffic (streaming video and audio) accounted for over 70% of downstream peak traffic on fixed access networks; (2) This is twice as much compared to just 5 years ago; (3) Traffic sources include major video services such as Netflix (33%), YouTube (17%), and Amazon Video (4%); and (4) BitTorrent traffic, which is declining and now is at 5% of total peak traffic.³⁹

individual pieces of content, the distribution is very fat-tailed (a small number of YouTube videos account for a disproportionate share of the traffic). Real-time streaming of user-generated content (e.g., via Periscope) is an exception.

³⁸ This may also include user-generated media, which also presents a more difficult challenge with respect to differentiating entertainment traffic from other sources of traffic, and with respect to its technical characterization.

³⁹ See Sandvine (2016), “2016 Global Internet Phenomena: Latin America & North America,” June 21, 2016. The share of entertainment video on mobile networks is lower but growing rapidly, and during the first half of 2016 exceeded 35%.

A variety of factors have contributed to producing this increase in video traffic, including increased consumption of higher resolution video and increased per subscriber time viewing Internet video. Today, the average consumer is spending more time watching Internet video compared to just four years ago.⁴⁰ Part of this increase is attributable to the growing use of mobile devices (smartphones, tablets) for accessing video both in the home and outside of the home. The proliferation of higher-resolution devices, lower (quality-adjusted) pricing for viewing devices, and the expansion of Internet programming options have all contributed to the growth in on-line viewing. While entertainment video viewing via the Internet has grown rapidly, it appears to have had only a relatively modest impact on legacy TV viewing, with most Internet video viewers also subscribing to legacy TV services, and with only limited evidence of cord-cutting thus far.⁴¹

When considering the impact of video traffic on the access network, it is important to recognize the data demands of higher resolution video streams. Video streams are moving from Standard Definition TV (SDTV) at 1Mbps to 3.7Mbps depending on compression,⁴² to High Definition TV (HDTV) at 3 to 20 Mbps, and even to Ultra HDTV (UHDTV) at 20 to 320Mbps. It is expected that HDTV and UHDTV content will increase over the next 5 years and SDTV will decrease. The result is a growth in the overall traffic, but largely in the downstream direction.⁴³

⁴⁰ eMarketer reports that the average time watching video has grown from 0:47 (2012) to 1:08 (2016) hours per day, while legacy TV watching has fallen from 4:38 (2012) to 4:05 (2016) hours per day (see “Growth in time spent with media is slowing,” eMarketer, June 6, 2016, available at <http://www.emarketer.com/Article/Growth-Time-Spent-with-Media-Slowing/1014042>).

⁴¹ The data on Internet video viewing substituting for legacy viewing is mixed. Nielsen seems to show a significant shift from legacy to Internet viewing, especially among younger audiences; however other studies indicate that Internet viewing is complementing traditional viewing (see “Traditional TV Viewing: What a Difference 5 Years Makes,” Marketing Charts, July 5, 2016, available at <http://www.marketingcharts.com/television/are-young-people-watching-less-tv-24817/>; or, Kiefl, B. (2011), “Trends in TV and Internet Use: the Impact of Internet TV on Canadian Programming,” Canadian Media Research Inc. (CMRI), July 2011, available at http://www.omdc.on.ca/Assets/Research/Research+Reports/Trends+in+TV+and+Internet+Use/Trends+in+TV+and+Internet+Use_en.pdf). As legacy TV providers shift programming to the Internet, however, the substitution effect is likely to become much larger. This will make it feasible to shift resources on fixed networks from delivering legacy TV to delivering broadband Internet.

⁴² MPEG2, H.264 or HEVC compression technologies are commonly utilized, with HEVC being the most aggressive in terms of compression.

⁴³ Sandvine reports that on fixed networks, real-time entertainment accounted for 20% of upstream and 70% of downstream traffic in 2016 (see Note **Error! Bookmark not defined.** *supra*).

⁴⁴ For example, legacy HTTP used store-and-forward encoding. More recent versions of HTTP

3.2.2. Traffic types/categories

Having characterized what we mean by entertainment video, we turn to characterizing the traffic. Convergence requires us to provision our IP access networks to handle a mix of traffic that is heavily weighted toward video, which is mostly one-way, cacheable, and associated with large-sized files. As we just described, the balance of non-media entertainment traffic that is associated with everything else, while smaller in aggregate volume, is more heterogeneous with respect to its requirements for connectivity, message size, and delivery specifications (routing, data rates, and QoS requirements). Figure 2 summarizes our qualitative characterization of the two major streams of traffic.

<<INSERT Figure 2 here >>

3.3. Evolving the network for video

The advent of video over the Internet is a relatively recent phenomenon that required a number of earlier developments. Widespread adoption of broadband access was one of the necessary precursors for Internet video consumption to grow. Previously, the limited speeds and quality of dial-up access and the limited capabilities of consumer devices (personal computers with low resolution monitors) resulted in a poor video experience via the Internet. Content providers who were unable to offer a compelling user experience via the Internet, naturally focused on other distribution channels (principally legacy TV and VoD). While the initial driver for broadband adoption was not video, but instead, the desire for faster web access and always on connectivity, the higher data rates offered by broadband made it feasible to embed video clips in what had previously been mostly static web content.

Faster broadband data rates necessitated that network servers, routers, switches, and consumer device capabilities all be enhanced to support and take advantage of the higher data rates supported over broadband Internet access. Many of these improvements, while benefiting video, were not necessary solely for video, but contributed to enhancing the performance of all Internet applications, and made it feasible for video and other types of traffic to dynamically share network resources. Faster hardware processors, more capable software, and enhanced network services have all played a part in enhancing the overall Internet experience, and at the same time, made it feasible to offer an increasingly compelling Internet video experience. Because video traffic has been growing so rapidly, it is hard to separate improvements that were motivated by the need to better support video that also benefit other types of traffic, from those that are more video-specific (e.g., the proliferation of video coding/compression technologies). For example, improvements to HTTP to support better encoding for dynamic content enhance the performance of both video and other non-video Web 2.0 applications.⁴⁴

⁴⁴ For example, legacy HTTP used store-and-forward encoding. More recent versions of HTTP support “chunked encoding,” which enhances the user experience by allowing the browser to serve streaming content to the user without first knowing how large the file is to be served. This allows the end-user to start the viewing experience sooner. The optimal chunk size depends on

The rise of content distribution network (CDN) services that were developed as overlay network services on the Internet to identify how to optimally distribute content (both dynamic and static) provided new techniques and capabilities for distributing content of all types, including video. The introduction of CDN services has re-written the interconnection model for the vast majority of Internet traffic, making it feasible to reduce the costs of content distribution across the Internet cloud.⁴⁵

Finally, end-user devices have improved significantly, with higher resolution screens, faster processors and image processing hardware, and enhanced operating systems with application support for interactive and mobile network services. These enhancements have benefited all classes of applications, including video. Each of these developments, together with multicast (push and pull), enhanced congestion control (load balancing), and more granular support for QoS have brought us to the current state, wherein the broadband Internet is supporting increased volumes of both video and “other 20%” traffic, and generally offering enhanced capabilities for all types of applications (i.e., higher resolution, lower latency, better reliability, mobility, etc.). Users are both streaming video, and using interactive (delay-sensitive) communication services (including video-conferencing) and other “cloud” services (e.g., eCommerce, IoT, file backups, social networking, etc.). However, it is unclear whether this will continue if the entertainment media that is currently mostly consumed over legacy channels migrates to the Internet, and if the nascent applications that compete for resources (e.g., video-conferencing, IoT) take-off. It is worthwhile considering whether sustaining the trajectory of improvements and capacity expansion that has proved successful thus far will prove equally successful in the future (when total Internet traffic loads are potentially orders of magnitude larger); or whether an alternate architecture that segregates the traffic might prove superior.

When we look across the network, we can highlight the bottlenecks that video creates, and identify how the providers have responded. As mentioned, network and end devices have employed a diverse set of incremental fixes to address many of the problems that video presents. While these fixes do a fairly remarkable job, there are still gaps in terms of the user's quality of experience (QoE). There are features that will continue to challenge the current model. These challenges are also exacerbated by a host of potential failures that can and do occur across the video content distribution path, e.g., publishing failures, CDN failures, server failures, peering and transit disputes and failures, ISP provisioning inadequacies, and end host issues.⁴⁶

the application and may be different for entertainment video streaming and other Web 2.0 interactive content.

⁴⁵ See Clark, D., W. Lehr, and S. Bauer (2011) "Interconnection in the Internet: the policy challenge," 39th Research Conference on Communications, Information and Internet Policy (www.tprcweb.com), Alexandria, VA, September 2011 (available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1992641); or Stocker *et. al* (Note 28 *supra*).

⁴⁶ Krishnamurthy, Balachander, Craig Wills, and Yin Zhang. "On the use and performance of content distribution networks." *Proceedings of the 1st ACM SIGCOMM Workshop on Internet*

So we have asked - has video traffic fundamentally changed the way that we are designing the Internet, and at what cost to other services? The first question, "Has video changed the design of the Internet?," is an easy one to answer: yes, as discussed above. To answer the second question, "How might separating video content impact last mile design?," is a challenge, that partly requires us to speculate what remains in the "pipe" and which applications and traffic might benefit from this new non-video Internet connection. In answering this second question non-video trends can inform a new model for characterizing the potential "other 20%" traffic.

3.4. Traffic Percentages and a Future Internet

In this section, we review traffic measurements over current broadband access networks, then use this analysis to describe what a future video-less link might look like.

3.4.1. Current Traffic Composition of Broadband Internet Access

To understand the impact of video on access networks, we measured the types of traffic running over a number of broadband access network connections.⁴⁷ The types of traffic were then classified into video, web, file sharing, gaming, voice and "other data" (data that we couldn't easily identify).⁴⁸ We found (on average) the following traffic on these links:

- video (65.1%),
- web (17.5%),
- file sharing (8.6%),
- gaming (1.1%),
- voice (<0.1%) and
- other data (7.6%)

Measurement. ACM, 2001. Jiang, Wenjie, Rui Zhang-Shen, Jennifer Rexford, and Mung Chiang. "Cooperative content distribution and traffic engineering in an ISP network." In *ACM SIGMETRICS Performance Evaluation Review*, vol. 37, no. 1, pp. 239-250. ACM, 2009

⁴⁷ We looked specifically at cable, DSL and fiber connections. While we believe that wireless likely presents a unique set of challenges, we will not deeply explore wireless access in this paper beyond the following comments: There has been some recent work looking at the role of broadcast technologies such as ATSC 3.0 and the LTE broadcast effort, and the Multimedia Broadcast Multicast Services (eMBMS) standard. Each of these seeks to enable video services in a more efficient manner for wireless (or over the air) broadcast. Capacity limits made more sensitive by spectrum costs may present an interesting perspective here. Of course, there is unlicensed spectrum, but this too presents its own set of challenges. Also, in some ways video has had a history different from most wireline networks as a result of usage caps on cellular services.⁴⁷ Lastly, how designers exploit crosslayer optimization may differ in a wireless scenario, where closely coupling of the layers might be justified.

⁴⁸ Note that among the 19 connections that we measured, the only service that had highly variable percentages was gaming. We omit the details of the measure method, as this was more of a informal tool to help consider the impact of video on the network.

These measurements are in reasonable agreement with other recent studies.⁴⁹ It is notable that even at peak periods, the network connection was at most 60% of peak capacity, meaning that these networks were not under heavy constraints. As stated, after excluding video, we have left web, file sharing, gaming, voice and “other data”. When we remove the highly asymmetric video data, we are left with a traffic mix that is more symmetric, and of a much lower data rate. We also have left a higher percentage of voice and gaming applications, which are less delay tolerant and not cacheable. Therefore, in an access Internet without entertainment video, traffic flows are:

- more heterogeneous;
- lower average data rate; burstier (peak/average higher);
- more symmetric (2-way);
- less latency tolerant;
- and less cacheable.

We next used historic data to extrapolate a model of what a future (video included) traffic mix might look like based on the trends described above (higher resolution video and more of it). We found that video could account for more than 85% of traffic, which is higher than many of the available capacities analysts are projecting. We conducted a similar extrapolation for non-video traffic and combined this with some usage scenarios described below. In this mix, the aforementioned differences of heterogeneity, data rate, delay sensitivity and cacheability were all pronounced.

3.4.2. Usage Scenarios of the 20%

To aid in our predictions about future traffic demands of the other 20%, we consider several likely usage scenarios. The goal is simply to think about a potential traffic mix for future non-video traffic. We consider several proto-typical scenarios – IoT in the home, healthcare, and public safety – where the value proposition driving the need for advanced communications taps into a very different set of social goals and personal needs. Focusing on these scenarios facilitates our ability to hypothesize about non-entertainment applications that may (or may not) suffer in an Internet of mixed traffic. The three scenarios are as follows:

- Case 1 –Home IoT: characterized by low bandwidth; variable latency; constant (control signals) to arbitrary connectivity.⁵⁰
- Case 2 –Healthcare: characterized by mobile, IoT and video-conferencing; low to high bandwidth, variable latency, and variable connectivity.

⁴⁹ See, the Sandvine and Cisco studies cited previously.

⁵⁰ Chapin, J. and W. Lehr (2010) "SCADA for the Rest of Us: Unlicensed Bands Supporting Long-Range Communications," 38th Research Conference on Communication, Information and Internet Policy (www.tprcweb.com), Alexandria, VA, October 2010, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1988184.

- Case 3 –Public Safety: characterized by ad hoc networking needs; low to high bandwidth, variable latency, and variable connectivity.

It is clear that the needs for low to high bandwidth, variable latency, and variable connectivity differ from video. Consequently, it is reasonable to conclude that a network optimized for video characteristics may not serve the needs of these non-video services. Each of these cases stresses the need for a heterogeneous/reliable/robust infrastructure. These services would put an emphasis on connectivity to a wider array of devices. It may even justify the need for multi-homing to provide robustness benefits against failure modes. These cases raise significantly different connectivity models compared to traditional video.

We recognize that while the above is speculative, the scenarios considered are important and oft-discussed usage cases cited to justify the importance of the Internet as essential infrastructure, and that each present traffic requirements vastly different from traditional video. (We also recognize that a portion of the traffic, even for those applications, will include traditional video, so any model of the “other 20%” would need to take into account provisioning for that video.)

3.4.3. Videos Impact on Other Services

Another major question we consider is whether the Internet is flexible enough to carry a majority of video traffic and still retain characteristics necessary to ensure the growth of other services? Of course, we cannot know what future Internet traffic will look like, but we can consider several scenarios with a variety of traffic types. Another factor to consider is the burstiness of (non-video) data.

At the core, the network topology has flattened and become more densely connected. Less traffic has to be routed from lower tier networks via Tier 1 providers, and more of the content is served from networks that are directly connected to access networks. Much of this has been driven by the rise of video content and the need to minimize the traditional flow of this type of content across the Internet. The role of the traditional Tier 1 and 2 backbone providers has diminished with the advent of more densely connected networks, the introduction of CDNs and the growing importance of edge networks. The most significant point regarding the core is that traffic resides in CDNs and other servers closer to the consumer, and therefore less of the traffic needs to transverse the traditional backbone. These changes have resulted in pressure on peering arrangements and, as a result, interconnection among the networks continues to evolve. Of course, at the edge of the network we see a better representation of what the consumer requests, in that video is traversing this link (unlike the CDN traffic cached at the edge). It could be possible to cache some of this video content at the home and shift this download to off-peak periods.

While we did not measure the traffic by type, sources have identified numbers in the range of 40% video and 60% other across backbone and edge networks.⁵¹

Other questions of interest include: does the source of the video matter? Will most of the entertainment content be coming from known providers in the entertainment industry or from random end points (e.g., be user generated)? YouTube has allowed everyone to become a production studio, where individuals can upload cacheable content to a server asynchronously. Most of this does not become popular and may only matter to a small group of friends, whereas a small number (out of the total) can become very popular and result in flash-crowds. We do not know precisely how Google (which owns Youtube) pushes popular content toward the edge but presume it is via a CDN mechanism for its most popular content. There are a range of CDN solutions that vary along multiple dimensions, including the volume of content the content providers need to distribute. For example, small video providers may be happy with just using the standard Internet, posting files to web-servers and accepting that the QoE of users may be variable; middle to large size commercial providers of video services may purchase CDN services from Akamai or other CDN to lower distribution costs and ensure better QoE for their subscribers; while the largest video distributors (e.g., YouTube or Netflix) are so large as to be undesirable customers for any commercial CDN, and need to self-provision their own CDN. They do this by negotiating bilateral agreements and using deep-network caching with eyeball access networks.

Another interesting question is, when does capacity isolation become an issue? This arises when broadband platform resources (whether in terms of the raw RF capacity on the wires or other network resources) need to be allocated between Internet and other services (which may or may not be other IP-based services). These other services have collectively been referred to as "specialized services" to signify that the traffic and how they are offered is distinct from the Internet. These specialized services may be subject to different sets of regulation.

While the scope of this paper precludes rigorous examination, we can look at a number of characteristics of specialized services and consider what this might mean for the overall characteristics of the broadband pipe and the different ways in which resources might be allocated (e.g., volume of isolated capacity, traffic type, and trends).⁵² A whole set of

⁵¹ See, for example, Alcatel-Lucent (2013), "Bell Labs Metro Network Traffic Growth: An Architecture Impact Study," an Alcatel-Lucent Strategic White Paper, available at <http://www.tmcnet.com/tmc/whitepapers/documents/whitepapers/2013/9378-bell-labs-metro-network-traffic-growth-an-architecture.pdf>; Norton, W. (2006), "Video Internet: the Next Wave of Massive Disruption in the U.S. Peering Ecosystem," Draft 0.91, available at <http://www-tc.pbs.org/cringely/pulpit/media/InternetVideo0.91.pdf>; or Adhikari, V., S. Jain, and Z. Zhang (2010), "YouTube Traffic Dynamics and Its Interplay with a Tier-1 ISP: an ISP Perspective," SIGCOMM 2010, available at <http://conferences.sigcomm.org/imc/2010/papers/p431.pdf>.

⁵² While beyond this thought experiment, it may be possible to characterize networks that are "tree-like" –where a limited number of sources may distribute to a large number of destinations,

related questions to consider include: How might we determine or measure the impact of capacity isolation on the rest of the broadband pipe? Is there a time when such isolation is good for the rest of the pipe (peak periods)? Does capacity isolation appear to be hurting the performance of current Internet traffic? Is there a role for the end user in deciding how capacity is isolated? Is there a role for new technology (SDN, NFV)⁵³ for a future network design - one where the edge of the network can dynamically morph to adjust to demands? These are open research questions to be addressed by the network research and policy communities.

3.4.4. Differences of Voice and Video as Specialized Services

From one perspective, the case for segregating video traffic from the Internet shares similarities with the case for segregating voice traffic. Today, many of the largest ISPs offer voice services that are implemented as Voice-over-IP (VoIP) over managed IP networks that are logically separate from the broadband Internet access traffic. This decision is often justified in order to enable providers to ensure adequate Quality of Service (QoS) for telephone services, and in order to facilitate the integration with emergency services (e911) and compliance with lawful intercept requirements (CALEA). Because these VoIP services allow calls to and from legacy telephone customers connected to the Public Switched Telecommunications Network (PSTN), these are sometimes referred to as "interconnected VoIP" services. In addition to such services, there are also VoIP applications and services that use their customers' broadband access service to carry voice packets. This includes services like Skype, Vonage, or Ooma. These services compete with the carrier-provided VoIP services, and are thus analogous to the situation being considered here, in which video entertainment traffic may be delivered to consumers via separate networks or via the Internet.

From a capacity perspective, the difference is that voice traffic requires far fewer bits and so consumes a relatively small share of capacity resources relative to video. However, voice traffic is communications-oriented and symmetric, and as such relies on the two-way communications functionality that characterizes much of the "other 20%" traffic and distinguishes it from entertainment video traffic. Moreover, from an economic/policy perspective, policymakers have long concluded that voice telephony services should be regarded as essential communications services; whereas it is unclear whether access to entertainment television is regarded as having as much social importance. In summary, therefore, while it is plausible to argue that voice and video traffic should both be

inherently broadcast/multicast versus fully-connected (any-to-any telephone)– and relate that to peak/multiplexing constraints and costs.

⁵³ SDN refers to Software Defined Networking and NFV refers to Network Function Virtualization. Both developments are examples of the softwarization of modern communication networks, which allows more granular, flexible, and dynamic resource assignment using software instead of hardware-based control. Moving the control of network functionality into software can also help reduce costs by enabling the de-localization of network functionality (facilitating the realization of scale and scope economies) and enabling commodity hardware to be substituted for specialty hardware.

separated from broadband into separate network flows, we do not make such an argument here. The two cases are sufficiently different from a technical, economic, and policy perspective that each should be considered independently. However, consideration of the case for entertainment video traffic separation may prove useful in informing an investigation of how voice services might best be managed in the future (and visa versa).

3.5. Architecture

In this section, we consider what convergence might look like, and the implications of different network architecture choices. We ask how might we provision entertainment video over a separate network service, one that might share the same physical conduit, or even the same virtualized pipe. A simple way of considering this is to think of “specialized services,” akin to the way IP Video or carrier VoIP services are currently delivered, using logically distinct portions of the RF on the last-mile. There are many ways that this separation could occur, and depending on the attention placed on maintaining the bandwidth and other performance characteristics of the non-video Internet connection; it might be little more than a traffic shaping mechanism, or it might involve complex cross-layer optimization.

This happens today over the access networks operated by cable companies and telephone companies when they provision triple play bundles of services (Internet access + TV/video + voice). For those that are still in the process of converting to all-IP broadband access platforms, the services share the same physical network (and RF on the wires) but are delivered using a mix of IP and non-IP based network technologies. When those networks are converted to all-IP, the traffic may continue to be delivered over separately managed IP networks that will continue to share the physical network, but with much more dynamic network control over how the available resources may be shared across the services. There is a continuum of potential provisioning models. At one end, there could be an IP network for best-effort Internet traffic plus one or more specialized IP networks for video, VoIP, or some other (future) segregated traffic type. In that case, the capacity allocations across the different IP flows may be managed at variable time scales and levels of granularity. For example, the allocations may be set according to relatively static constraints that are managed based on aggregate per-IP flow characteristics. At the other extreme, one might envision a single IP network with per-packet or per-flow QoS management.

Today, we accommodate the variety of traffic flows in the best-effort Internet with a complex mix of workarounds that includes techniques, overlays and value-added services such as VPNs, CDNs, MPLS, etcetera.⁵⁴ It is certainly possible that this adaptation by accretion may continue. However, it is unlikely that simply provisioning a single IP

⁵⁴ VPNs are Virtual Private Networks, and MPLS stands for Multi-protocol Label Switching. Both of these are techniques for adding QoS control and traffic management capabilities to the legacy Internet. For a discussion of overlays in the Internet, see Lehr, W., D. Clark, P. Faratin, R. Sami, and J. Wroclawski (2006) "Overlay Networks and Future of the Internet," *Communications and Strategies*, no. 63 (3rd Quarter 2006) 1-21.

network offering only best-effort traffic management would be adequate to address the needs of traffic in the future.⁵⁵ If the future is separate IP networks that mix a best-effort Internet and specialized IP networks for other traffic while sharing the same physical network, then the issue to resolve will be how best to allocate resources. This is more than just a technical question since the technologies make it possible for us to address this question at a variety of time, geographic and context-related scales. At one extreme, we could have static resource assignments (full capacity isolation), whereas with full-blown NFV/5G, we could have dynamic resource assignment on a per-packet basis and technically, the traffic would be fully converged (i.e., single IP network with QoS).⁵⁶ Or, we could have more limited integration, more closely approximating what we have today where video and VoIP are separate logical networks and integrated at gateways (at PSTN-SIP interconnection gateways for IP-TDM conversion, or set-top boxes where entertainment video and broadband traffic are split/integrated for distribution to appropriate devices in the home, etc.).

If we considered what a separate IP network optimized for entertainment video traffic might require in the way of specifications and functionality, a range of (potentially) important features may be identified, including: (1) Support for DRM management; (2) Edge-caching and CDN support (to optimize caching); (3) Multicast support; (4) Support for diverse video encoding standards (which might be price-tiered based on bandwidth utilization, etc.); and (5) other requirements.

There are various reasons that it might make sense for broadband access service providers, particularly cable companies, to move their current linear video onto IP. Some of these are technology enablers and others are motivators. Advances in the design of their networks has allowed use of wide band channels, as opposed to the legacy 6MHz channel, and these wider bands allow for much higher throughput. Also, methods such as pruned multicast supports more efficient and targeted use of bandwidth. Being in IP with the opportunity to dynamically adjust channel capacities could facilitate better signaling between edge devices (e.g., regarding screen size and desired resolution, DRM permissions, or other video features) and may allow operators to dynamically choose between video encoding on the fly (e.g., to address different resolution preferences) or

⁵⁵ Reed (2016) explains how modern cable networks may provide the QoS support needed to handle the projected growth in video and other types of Internet traffic (see, Reed, D. (2016), "Trends in Cable Network Economics: Implications for Public Policy," draft paper presented at NCTA INTX Workshop, May 2016).

⁵⁶ Even if all traffic shares a single IP network, we may still desire regulatory separation of the traffic. How to achieve that would present a challenge that would likely have implications for resource allocation decisions. For example, label all public safety traffic with a special tag and then allow that to pre-empt any other type of traffic; or label all best effort traffic with another tag and let that be buffered or dropped first whenever congestion occurs. The range of potential options enabled by technology is virtually unlimited.

selecting pre-encoded content, depending on what the content was and how popular it was.⁵⁷

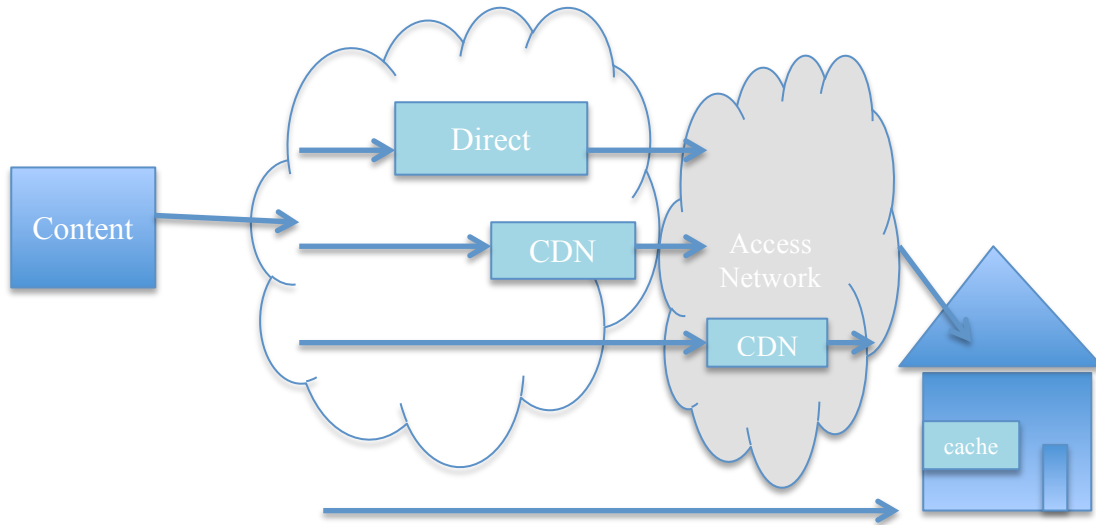
Being in an all-IP environment also offers benefits in terms of economies of scale, reduced complexity, and ease of implementation and deployment of features and service customization. Today's video service providers often rely on hybrid technologies, distributing some content via broadcast RF to all subscribers (and not necessarily even in digital format), distributing other programming via switched digital services (again, not necessarily IP), while offering "TV anywhere" access to selected programming over the Internet.⁵⁸ A cable company may choose to move its content into a single distribution network, served over one efficient network (CDN driving IP over the access network), serving all its customer's points of consumption (in home, out of home; set-top box, smart TV, tablet, cell phone, etc.). The operational simplicity of this paradigm and ability to switch to an agile/lean technology deployment model on this platform offers a number of advantages.

3.5.1. Getting video to the customer

In this section of the paper, we first consider how video content gets to the edge of the network. We then briefly consider how a Video Internet Access Service (VIAS) might work technically (addressing how this might work from a policy perspective in a subsequent section).

⁵⁷ Video encoding incurs processing costs. For popular content that may be desired in multiple resolution formats, it makes sense to encode the content in advance and store multiple copies in different resolutions; while for content that may only occasionally be desired in multiple resolution formats, it may make sense to encode the content on the fly. Being in an all-IP environment enables these sorts of decisions to be made on a per-content or per-customer basis.

⁵⁸ According to the FCC, 69% of MVPD households were served by all digital networks by the end of 2014, up from 57% in 2013. At the end of 2013, 45% of the digital cable subscribers were served using Switched Digital Video (SDV). SDV allows providers to conserve bandwidth by using shared bandwidth to transmit video-on-demand for less popular programming. Traditional video service providers like Comcast and others are offering "TV Anywhere" services to allow their subscribers to access certain content over any Internet-connection. This allows them to enhance the value of the service offered to their subscribers, and provides a competitive response to the threat posed by over-the-top programmers. See FCC (2016), *Seventeenth Report*, In the Matter of Annual Assessment of the Status of Competition in the Market for the Delivery of Video Programming," Before the Federal Communications Commission, MB Docket No. 15-158, released May 6, 2016, available at http://transition.fcc.gov/Daily_Releases/Daily_Business/2016/db0506/DA-16-510A1.pdf.



As the figure above depicts, at a high-level of abstraction, one might think about a content provider seeking to deliver video to edge subscribers by: direct streaming to the home without the use of a CDN (however, the content could still be cached at the home); placing content in a CDN close to the access network; or placing CDN within the BIAS provider's network. In each of these situations, the content provider could make use of transit, peering or other network deployment strategies to reach the customer. Furthermore, the content provider could even be the same company that owns the access network. A fourth way would be to have content delivered without the use of the Internet (think Netflix DVD delivery).

Even with the significant architectural enhancements just described, the performance shortcomings of over the top video are well documented.⁵⁹ These include:

- Set up delays of 5-10 seconds for a majority of sessions;
- Downlink bandwidth variation, which leads to buffering and rate adaption;
- Significant buffering and re-buffering (as a percentage of traffic and time); and,
- Variable quality CDN performance.

⁵⁹ Liu, et al., "A Case for a Coordinated Internet Video Control Plane," *SIGCOMM'12*, August 13–17, 2012, Helsinki, Finland.

3.5.2. VIAS

Here we briefly consider how a VIAS service might be offered to the customer. This is simply to highlight how such a service could be isolated from Internet service and to further demonstrate the varying degrees of convergence on a network. The options, going from completely separate to completely converged, might include:

- Independent connections (e.g., BIAS for Internet, with video delivered over a completely separate physical network such as over-the-air broadcasting or satellite)
- Same connection, separated at the PHY layer
- Same connection, separated at the Data Link layer
- Same connection, separated at the Network layer
- Same connection, separated at the Transport layer
- Over the top – OTT – (e.g., Netflix)

When considering how video is delivered to a customer, we might start with the historic norm of it coming through a separate conduit. For example, it could come in via broadcast TV or satellite transmission. This approach is still common, and it is even the basis of AT&T's motivation in acquiring the DirectTV satellite service. Alternatively, video can enter the home on the same connection as the Internet access service by combining streams. This combination of services can occur in a variety of ways at different layers of the network protocol stack as laid out above. For example, this could occur over MoCA⁶⁰, DOCSIS, a VLAN⁶¹, VRF⁶², or some traffic shaping method. The most familiar combined video and Internet access started with cable providers, who made use of DOCSIS technology to implement Internet access on top of the existing cable video distribution platform. In this model, all downstream data arrived in 6 MHz channels (traditional channel specification for broadcast video). The video content is

⁶⁰ MoCA stands for the Multimedia over Coax Alliance (<http://www.mocalliance.org/index.htm>). This industry alliance is focused on developing a suite of standard technologies for delivering a host of applications into homes over coaxial cables and within homes over coaxial, WiFi and other in-home networking infrastructures.

⁶¹ VLAN stands for Virtual Local Area Network which is a class of technologies that operate by modifying the headers of Ethernet packets to enable finer-grained traffic management capabilities (see <http://www.dummies.com/how-to/content/how-virtual-local-area-networks-vlans-work.html>).

⁶² VRF stands for Virtual Routing and Forwarding which is a technique for provisioning for and managing data traffic on a per flow basis over cable modem networks that operates at Layer 2 (see, http://www.cisco.com/c/en/us/td/docs/cable/cmts/config_guide/b_cisco_cmts_layer_2_vpn_features/b_cisco_cmts_layer_2_vpn_features_chapter_01000.html).

provided by Digital Video Broadcasting-Cable (DVB-C), which carries the video as MPEG-2 or MPEG-4 streams over QAM (the video is not carried over IP).⁶³

The Internet access traffic is also carried in these 6 MHz channels, but makes use of a differing data link technology to separate it out. Interestingly, cable operators are evolving toward a more IP-centric approach with plans to move linear channels and VoD over IP.⁶⁴ Even though all of this will be carried over IP, the video traffic will receive some differentiated treatment.⁶⁵ A similar all-IP approach is expected in the fiber and DSL space. Continuing up the stack it is possible to use layer 3 or layer 4 methods to separate and/or differentiate a video service from an Internet access service. Lastly, there is the possibility of simply placing everything as just another app running on the broadband access service – i.e., over the top video.

Before leaving the subject of video delivery, it is worth mentioning some recent research. Caching content at the edge in CDNs is increasingly becoming the primary source for serving video traffic to consumers. A recent paper by Brzozowski and van Doorn,⁶⁶ argues that the cost of the cable access network is so high that it makes sense to cache selective content at the subscriber's residence. This isn't a new idea, but the paper does present a compelling argument that video is dominating the access network, particularly during peak usage, and that caching could reduce this load. Earlier Liu et al. (2012)⁶⁷ argued in favor of implementing a separate Internet video control plane. That case could be extended to make the case for a separate data plane for video. While HTTP and related protocols have been tweaked to better carry video, it is still not the most obvious and efficient method of transport. Of course, the same argument was made against VoIP, but for a number of reasons, it worked well for transmitting voice.

⁶³ See cable specs available at <http://www.cablelabs.com/specification/physical-layer-specification>, which describes the way in which the cable PHY layer is evolving to handle higher IP data rates.

⁶⁴ It is expected that VoD service will migrate first, followed by less popular linear channels, and eventually the top linear channels. This will allow multicast services to carry these linear streams and sports VoD to customers, which cannot be done for OTT content. This will further allow a switched video architecture that supports even higher IP data rates, including much higher Internet access data rates.

⁶⁵ It is worth mentioning that even in early Data Over Cable Specification (DOCSIS) versions, a method to offer high quality voice was implemented as a QOS extension at the data link layer.

⁶⁶ See Brzozowski, J and J. van Doorn (2016), "Improving Customer Experience Through Cooperative In-home Caching And Pre-positioning (CIHCP)", NCTA Technical Conference, available at <http://www.nctatechnicalpapers.com/Paper/2016/2016-improving-customer-experience-cihcp>.

⁶⁷ See Liu, Xi, Florin Dobrian, Henry Milner, Junchen Jiang, Vyas Sekar, Ion Stoica, and Hui Zhang (2012), "A case for a coordinated internet video control plane," In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, pp. 359-370. ACM, 2012.

3.6. Summary

Summing up our preliminary technical discussion of how video traffic has changed the Internet and what this may mean for the future architecture of the Internet, we offer the following observations.

First, video traffic is continuing to grow as a percentage of total traffic on broadband access links. This growth is only expected to continue due to growing customer demand to access video via the Internet and the shift to higher resolution (higher bitrate) video content, more capable devices, and the opportunity to consume video in more locations (e.g., mobile, wireless). Traditional commercially produced entertainment video is a growing share of this traffic, but the volume of user-generated and non-traditionally produced content is also growing. The expectation is that if current trends continue, entertainment-dominated video traffic will comprise 80% or more of the traffic on last-mile broadband access networks (whether delivered via a single BIAS or multiple IP-based services).

Second, while we can model the current state of traffic, projections of future traffic are fraught with uncertainty. Nonetheless, it is possible to consider broad traffic types and the implication of the growth of each. Our simple model suggests that if entertainment video is extracted from the general mix, the resulting "other 20%" will grow at a slower rate (in terms of aggregate traffic), will be more symmetric, and include more delay-sensitive traffic than would be the case if all of the traffic were aggregated.

Third, the Internet (including backbone, interconnect, distribution and access networks) has evolved to accommodate the shift to more video traffic. These changes include CDNs, dense interconnections, higher backbone data rates, higher access data rates, and more. These enhancements, not completely attributable to the rise of video traffic, have delivered service quality improvements to all applications in terms of enabling higher data rates, reduced latency (delay), and reduced losses. This implies that a model where video continues to move into the Internet access pipe might imply a "rising tide" of service quality for all services.

Fourth, we expect video to further converge onto the IP pipe, as linear content is moved over IP in light of the benefits of moving to all-IP for wired last-mile networks.⁶⁸ The real question is how broadband network resources may be managed over all-IP networks and how edge and core functionality may be allocated. The manner in which video is delivered on a converged access service matters. For example, shifting from a unicast over-the-top model to one that made coordinated use of VoD or switched video and network and customer-premises caching could actually free up bandwidth resources for

⁶⁸ A case can also be made for all-IP in wireless networks, but that is rendered more complex by the requirements of operating over a wireless physical layer at different frequencies. For further discussion, see Lehr, W. and J. Chapin (2010), "On the convergence of wired and wireless access network architectures," *Internet Economics and Policy*, 22 (2010) 33-41.

non-video Internet access and improve the quality of experience for all services. In contrast, simply moving current linear video traffic wholesale into the IP pipe without traffic management could congest last-mile networks, resulting in reduced performance for Internet access services. The implication is that the architecture of this converged pipe will matter.⁶⁹

Fifth, the entertainment industries continue to evolve both new services and business models for capturing consumer attention and dollars. As entertainment content continues to shift from a broadcast model (which was the TV broadcast model of the 21st century) toward a model relying more on storable media and customized viewing experiences, the requirements for what is demanded from the electronic distribution networks is changing. The key question that emerges in a world that relies more on VoD is whether to store locally or in the cloud, which is a choice that often hinges on issues of costs, functionality, and market structure. Planning for a mix of both edge-based and cloud storage and localization of functionality, with consumers having the ability to substitute between the options, is likely to force cost equalization. That is, the ability of consumers to switch between cable TV-provider services, buying movies from iTunes, or subscribing to Netflix will impose reciprocal constraints on what features video service providers need to offer and at what prices if they are to remain competitive in the evolving markets for entertainment video.

4. Economics

In the preceding section we discussed some of the ways in which the design of the Internet might be different if video and data traffic were provided via separate logical networks. In this section, we consider how bundling services into the same broadband access service affects the business model for access ISPs, as well as for other providers of entertainment video services that might rely on the access ISPs in order to deliver their content to end-users. These are the Over-The-Top (OTT) providers that may choose to deliver their video either via Web-based streaming (enabled by a browser) or via an application (e.g., a Netflix application running on a tablet, smartphone, PC, or connected TV) that presently uses a broadband access connection to deliver the video content from the content-providers server (or cache) to the individual subscriber households.⁷⁰

⁶⁹ Our simple model also shows that a more converged network could benefit from potential traffic differentiation. As video traffic consumed more of the access service, all traffic suffered in terms of delay and loss (note that this assumes that the bandwidth of the connection does not increase, or does not increase proportionally). This raises questions about how this would be implemented and who decides the allocation and issues of fairness that surrounds this. How does a separate network for specialized services impact value of Differentiated Services for Internet (versus Best Effort)?

⁷⁰ In many cases, Wi-Fi may be used to connect the end-user device to the fixed broadband access service in the subscribers home or in a café, hotel, or other outside-the-home location. It is also possible that the content may be delivered via a mobile broadband service when Wi-Fi is not an option.

In the following two sub-sections, we first consider how the economics of bundled services impacts the business models of video providers; and then how entertainment economics might bias the evolution of the Internet.

4.1. Benefits of Bundling Services

Bundling is an important strategy for both the delivery of network services and entertainment media. Bundling can occur in multiple ways. For example, the content and conduit are bundled when the network provider also offers the services (e.g., Comcast provides cable television programming services over its broadband network platform); content choices are bundled when programmers offer a menu of selections (e.g., Netflix offers a library of movies, television networks offer suites of program choices); and bundling may occur at the retail level when a provider sells complementary retail services to consumers (e.g., Verizon sells telephone, broadband data, and video services as a triple play package). Finally, bundling occurs at the technical level through the choice of IP as the fundamental network protocol, which offers advantages in terms of interoperability and network layering.

Bundling facilitates resource sharing and the realization of scale and scope economies, which can lower overall costs and facilitate more dynamic and flexible capacity scaling.⁷¹ Bundling of content choices in entertainment media and with particular distribution channels facilitates addressing heterogeneous consumer demand and "windowing" by content providers.⁷² Bundling also provides strategic benefits for both network service providers and entertainment media companies that can impact competition.

For example, access ISPs prefer to sell a bundle of services to consumers, which includes voice, video, and data – or broadband Internet access – services for multiple reasons.⁷³ Once a provider has installed the last-mile facilities needed to provide any one of the

⁷¹ When traffic with imperfectly correlated peak capacity demands share resources, total provisioning costs are reduced; and sharing capacity enables flexible provisioning in face of uncertainty.

⁷² Consumer tastes for programming vary (across consumers, time, context) and to address those heterogeneous tastes, entertainment companies sell bundles of content (e.g., amusement parks have many different types of rides, Netflix has a library of movies, and channels offer streams of programming). "Windowing" is the term used to describe the practice of media companies using distribution channels to segment markets to facilitate price discrimination. Traditionally, theatrical releases and hard back books were released in the first window at the highest price per viewer, with subsequent windows designed to capture the demand of lower price consumers (e.g., paperbacks and broadcast TV). Today, the timing and ordering of media distribution windows has been disrupted, but the term is still used to refer to how bundling with distribution channels is used to segment markets.

⁷³ Access ISPs are expanding their bundled offerings with new services like home security and smart home management services, additional web services (e.g., access to premium content or on-line storage), or, in some cases, mobile services.

services, the incremental costs of providing additional services over that network are relatively small. The most significant components of access network costs are associated with deploying the wired facilities that increase with the number of households passed. Those costs are mostly fixed, and a large proportion of those costs do not vary with the number of actual subscribers, nor with the traffic utilization of those subscribers.⁷⁴ Selling multiple services to subscribers strengthens the relationship with those subscribers⁷⁵ and provides a larger revenue base over which to spread the shared and fixed (and sunk) costs of operating a facilities-based access network.⁷⁶ Finally, providing the multiple services over a common all-IP broadband platform (i.e., bundling at the technical level) offers both demand and supply-side benefits by making it easier to offer new, more capable, and better integrated services.⁷⁷ IP serves as a spanning layer that can be supported across diverse physical layer infrastructures (wired and wireless; copper, coaxial cables, and fiber) and can allow multiple applications to be supported (video, voice, and data). The use of IP also enables interoperability and interconnection across networks of different types.

The ISP's consumers also benefit from the opportunity to purchase bundled services since that enables one-stop-shopping that may simplify customer billing and interactions with customer support. This is of special importance when trying to diagnose service problems. Additionally, consumers who purchase bundled services typically receive a discount

⁷⁴ Note, we are not saying that access network providers do not also confront traffic-sensitive network costs that vary with the aggregate traffic loads that the operators must handle. The interconnection and middle-mile facilities of access providers vary with the aggregate traffic loads that providers must handle. The desire to manage these costs, while ensuring high-quality access to content induces operators to seek to locate cacheable content as close to end-users as possible and deploy other strategies (e.g., compressing content files, employing multicast or anycast routing, etc.) to economize on network resources wherever possible. Content Delivery Networks (CDNs) like Akamai (and proprietary ones deployed by large content providers like Netflix and Google) overlay ISP networks, assisting in the management of capacity to efficiently deliver content while preserving a good experience for the content provider's customers. For further discussion, see Stocker, Smarakadakis, Lehr, and Bauer (2016), *If Content is the King, (Peering) Location is the Emperor: an examination of CDN trends and evolution* (2016), ITS2016, Cambridge, forthcoming.

⁷⁵ Bundled subscribers churn less frequently (i.e., have longer subscriber lives with a provider). Reducing churn (the % of subscribers who leave each month) is a key strategic goal that contributes directly to profitability.

⁷⁶ The opportunity to realize scale and scope economies is not limited to network costs, but also applies to a significant share of other business cost elements. For example, brand advertising, billing, back-office, and other components of non-network costs are fixed to a significant extent, and do not vary with the number of subscribers or their traffic.

⁷⁷ For example, although it was possible to augment basic telephony service with enhanced features such as voice mail and call-forwarding even with legacy circuit-switched telephone networks, the ability to manage, define, and customize features in an all-IP, software-manageable network is much greater.

relative to the sum of the a la carte prices, which allows consumers to share in the cost savings afforded by bundling.

Content providers also benefit from bundling content choices and from bundling those choices with distribution channels. First, it is worth noting that providers of entertainment programming have multiple delivery options for getting their content into the hands of consumers. This includes movie theaters, over-the-air broadcasting, Digital Broadcast Satellites (DBS), and cable – as well as using broadband Internet service. Indeed, Netflix originally distributed its content by bundling it with a high-latency, high-bandwidth broadband network. That is, Netflix used the US Postal Service to deliver its content. Each of these distribution channels offers different benefits, and most content rights holders seek to maximize the revenue that can be captured by their content by using multiple channels through the process of "windowing," discussed earlier.

Although multiple distribution channels for media content exist, these vary in cost and quality and so are, at best, imperfect substitutes for IP-based delivery (although, depending on the circumstances, not necessarily inferior).⁷⁸ The benefits of digital distribution via wired broadband networks, and increasingly, over all-IP networks means that in many cases this will offer the lowest cost and most flexible network option for distributing entertainment video. This helps explain why so many entertainment content providers are interested in expanding their over-the-Internet service offerings.

In addition to the benefits of IP as a distribution channel, the Internet has expanded opportunities for end-users to access a greatly expanded selection of content from a much larger universe of content providers. The rise of user-generated content associated with services like YouTube, and the globalization of the Internet, the rise of search tools and social networking (with recommendation applications) have all contributed to making it feasible to present consumers with a larger library of usable entertainment video options than was ever possible before.⁷⁹ A number of researchers have identified the benefits from bundling digital media content (not just for access to video, but also other digital

⁷⁸ A general theme is that content providers of video entertainment are in the business of delivering a wide-range of consumer experiences from the content they offer, and whether high-resolution, mobile access, programming selection, with or without advertising, viewing screen, time, interoperability with other applications, price, or some other attribute is most important in a particular viewing context will vary, and may be impacted by the delivery option chosen. Thus far there has not been, nor do we think there will be, a single way to deliver entertainment video that is optimal for all situations.

⁷⁹ When the only way to see Hollywood movies was in theaters, the number of theaters limited the number of movies that could be viewed. With the rise of television, viewing options expanded but were still limited to linear programming available on a relatively small number of channels. With the rise of video-on-demand and now the Internet, the limits on consumer choices that may be presented, searched, and sorted has become effectively unlimited.

print and audio content), and the opportunity it provides for essentially unlimited shelf-space for consumer choices.⁸⁰

While there are many benefits from bundling video content in the multiple ways in which it may be bundled, there are also potential risks. Bundling may have anticompetitive effects to the extent that it may raise entry barriers and impose switching costs on consumers.⁸¹ To the extent this is the case, consumer choice may be limited if the market is not sufficiently large to support adequate competition among providers offering bundled services. As will be discussed further below, concern over the economic viability of sufficient facilities-based competition in access networks has provided a long-term justification for regulation of last-mile networks, although historically this has tended to be service *and* network specific because originally each service was supported by its own network (i.e., cable supported video media, the telephone network supported voice calls, etc.). Concerns that bundling might lead to market power over distribution channels (conduit), choice (content), or both have motivated a range of regulatory interventions, including common carrier regulation, media concentration and program access rules.

Bundling can also have ambiguous effects on service pricing from a consumer welfare perspective. For example, for continued investment in access network infrastructure to remain economically viable and incentive compatible, investors need to be able to recover their economic costs. From a pricing perspective, there is no single best way to allocate the shared and fixed costs of the local access network across the different services. The optimal allocation of costs depends on demand and competitive considerations. However, it is reasonable to expect that if an access ISP sold only broadband service, then the total costs of the access network would need to be recovered from that single service.⁸² Whether that would mean that total household payments to the last-mile facilities providers would increase or not is uncertain and depends on market dynamics.⁸³

Another important question is how other video providers (OVDs) that are not affiliated with the access ISP deliver their programming to end-users. Indeed, over the top

⁸⁰ See, for example, Bakos, Y. and E. Brynjolfsson (1999), "Bundling Information Goods: Pricing, Profits, and Efficiency," *Management Science*, 45(12), 1613-30; Bakos, Y. and E. Brynjolfsson (2000), "Bundling and Competition on the Internet," *Marketing Science*, 19(1), 63-82; or Crawford, Gregory S., and Ali Yurukoglu. "The welfare effects of bundling in multichannel television markets." *The American Economic Review* 102, no. 2 (2012): 643-685.

⁸¹ For example, consumers who switch a single service may lose their bundled discount; while customers who switch the entire bundle may confront service-specific switching costs (e.g., difficulty in migrating a customers email address).

⁸² Note that revenue for cost recovery need not come solely from end-users, but may also come from other value chain participants, including advertisers, content or application providers.

⁸³ Too much or too little competition could result in either higher total costs and/or prices and reductions in high-quality programming choices (although deciding what constitutes appropriate quality is likely to be highly contentious).

entertainment video providers like Google (YouTube) and Netflix were among the strongest advocates in favor of the FCC adopting its Open Internet Order⁸⁴ rules that are designed to protect consumers and edge providers (which include entertainment video as well as other content and application providers) from discriminatory treatment by access ISPs. They argued that in the absence of so-called Network Neutrality protection, access ISPs with market power might abuse their power to extract excess rents from edge providers and consumers, and might discriminate in favor of affiliated content. It is interesting to contemplate how the policy debate would have unfolded if the entertainment video content providers traffic had not been part of the picture.

The Internet's ability to present consumers with unlimited choice has ambiguous implications for the production of diverse quality content. For creators of commercial content, the fragmentation of markets may reduce the potential revenue that may be captured by a program, which may reduce ex ante incentives to invest in its creation, which could result in fewer and/or lower quality choices being available. Moreover, the marginal benefit of additional programming choices is likely to decline and may actually become negative (too much choice, especially if much of it is of very low quality, may actually result in diminished consumer surplus and may crowd out more socially desirable content).⁸⁵

Today, access ISPs provide Internet, video (and voice) services via separately managed networks.⁸⁶ This does not preclude them from providing bundled service offerings to consumers, and has not limited the ability of access ISPs to expand into on-line or other new media entertainment offerings when those seem attractive. From a cost perspective, being able to offer all of the services over a common, integrated IP platform may offer important economies. It would simplify network provisioning and operations and could allow the ISP to avail itself of global economies associated with commodity IP hardware and software solutions. However, as already noted, these cost economies could be realized by sharing the IP platform across multiple IP networks, each of which might be logically separated and dedicated to a different service.

⁸⁴ See FCC (2015), "Report and Order on Remand, Declaratory Ruling, and Order," In the Matter of Protecting and Promoting the Open Internet, Federal Communications Commission, GN Docket No. 14-28, Adopted February 26, 2015, available at https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-24A1_Rcd.pdf (hereafter, OIO).

⁸⁵ The debate over choice and consumer surplus is a long one. Traditionally, economists argued that more choice has to be Pareto improving because consumers could simply ignore options they were not interested in. However, because making decisions is costly, this is not obviously true. Barry Schwartz, a psychologist, explained how excess choice can actually result in consumers being less happy (see Schwartz, B. (2004), "The paradox of choice: why more is less," Harper Collins: New York, 2004)).

⁸⁶ The video services are typically provided as a broadcast service using RF channels that are separate from the RF channels that are used for broadband, although with the transition to all-IP cable platforms (and FIOS/ADSL for other access ISPs), the technical flexibility to dynamically manage resources across services is greatly enhanced.

Thus, while service bundling offers important advantages for many access network providers and consumers, and while it seems clear that the ability to share the cost recovery burden associated with deploying broadband infrastructure with entertainment video services, this does not require that the services be delivered into the home over a single converged IP network called “the Internet” via the regulated Broadband Internet Access Service (BIAS). Although, if the video is not delivered via the BIAS service, then we might have to create a new regulated service (e.g., VIAS) if we want to ensure competitive access to last-mile bottleneck facilities.

4.2. Entertainment Industry Economics and Broadband

In this section, we focus on how the economics of entertainment media are distinct from the economics that motivate the “other 20%” of Internet traffic. Our goal here is to sketch a picture in stereotypes, which over-simplifies the actual situation, but proves useful in highlighting important lessons. We begin by characterizing the demand drivers that motivate the economic activity that broadband is intended to address. Next, we contemplate what insights may be revealed regarding the societal value and willingness-to-pay for broadband services. We then consider the relative importance of different network capabilities. We conclude with speculations about how this impacts the behavior and structure of markets for broadband services.

We will use “Smart-X” to represent the demand that results in the traffic that is included in the “Other 20%”.⁸⁷ By Smart-X, we mean all of the ways in which an optimistic vision of “pervasive computing” might allow us to realize the full promise of information and communications technologies (ICTs) for enhancing the performance of complex systems across the economy.⁸⁸ The “X” refers to the many different complex systems that ICTs are expected to enhance. The “X” may be replaced with healthcare, public safety, energy grids (the three to keep in mind for our discussion here), but also eCommerce, eGovernment, etcetera.⁸⁹ The “Smart” refers to enabling ICTs in the X, and we assume

⁸⁷ This is not all of the “other 20%” of the traffic because that includes some entertainment video as well as everything else that is not Smart-X.

⁸⁸ The vision of “pervasive computing” is one in which networked processors and software applications to take advantage of them are ubiquitously available in the real-world. This extreme vision is one in which everyone/everything is always/everywhere connectable, and Big Data-analytics and AI/robot automation empower real-time, granular (local) decision-making. This has the potential to deliver better dynamic, flexible, and customizable resource allocation decisions to improve the operation of complex systems. Of course, we recognize that the extreme version of this optimistic vision will confront numerous challenges and is unlikely to be realized.

⁸⁹ We focus on these three because we believe the economics of healthcare, public safety, energy grids, and entertainment as “economic goods” are sufficiently distinct that readers will be able to follow our stereotypes without too much resistance. While we recognize that the first three may be as different in their needs as they are from entertainment, we want to focus on what it may mean to consider entertainment separately, and provide a potentially common infrastructure for everything else.

that this requires Internet access, but are interested in exploring whether the access that is needed to support Smart-X would be different (and how different) if the access were not also called upon to support entertainment video.

In thinking about how entertainment economics and smart-X economics may differ, it is useful to note that consumer expenditures for entertainment come out of discretionary leisure expenses that are generally a relatively small share of income (around 5%).⁹⁰ Over time, the budget shares allocated (in time and dollars) to different categories of activities have remained far more stable than the shares of specific items within categories. This suggests that households and individuals first allocate their income and time to various activity categories, and then within those categories, substitute among activities. In recent decades, the average working person spends about 40 hours per week working, and close to 70 hours per week on personal care (which includes sleeping), which leaves about 58 hours per week for everything else. Over the past hundred years, shifts have occurred with a significant decrease in work time and expansion in leisure time; and of course, the young, elderly, and unemployed have more time for leisure activities.⁹¹ The American Time Use survey reports that the average citizen over 15 years of age spent 5.3 hours per day on leisure and sports activities.⁹² According to Nielsen data, the average household is streaming video entertainment 6 hours per day, with 5.1 of that going to TVs and the rest to other connected devices (tablets, PCs, smartphones, etc.).⁹³ It is likely that for some of that time no one is actually paying attention.

While the leisure shares of consumer expenditures and how consumers spend their time are significant, the budget shares are less than what consumers spend on healthcare, housing, or transportation.⁹⁴ Entertainment video is competing for a relatively small share of consumers' time and discretionary budget dollars. Not surprisingly, competition for

⁹⁰ For example, the BLS reported that U.S. consumers spent about 5% of their budgets on entertainment in 2003, slightly up from what they spent in 1950, but significantly more than they spent in 1900 (see BLS (2006), "100 Years of U.S. Consumer Spending Data for the Nation, New York City, and Boston," U.S. Bureau of Labor Statistics (BLS), Report 991, May 2006, available at <http://www.bls.gov/opub/uscs/report991.pdf>).

⁹¹ One source shows that since 1870, the hours worked per week in many countries has fallen by slightly less than half to reach the 35 to 40 hour work-week we have today (see <http://www.business2community.com/tech-gadgets/humans-spend-time-changed-01108897#sLiYcuFujDDTjvLh.97>).

⁹² See American Time Use study data from U.S. Bureau of Labor Statistics for 2014, available at <http://www.bls.gov/careeroutlook/2015/data-on-display/work-and-play.htm>.

⁹³ See Nielsen (2016), "Television is still top brass, but viewing differences vary with age," July 18, 2016, available at <http://www.nielsen.com/us/en/insights/news/2016/television-is-still-top-brass-but-viewing-differences-vary-with-age.html>.

⁹⁴ In 2014, the average expenditures per consumer unit were \$53,495 in the U.S., with 8% going to healthcare, 33% to housing, and 17% to transportation (see BLS (2015), "Consumer Expenditures (Annual) News Release," U.S. Bureau of Labor Statistics (BLS), September 3, 2015, available at <http://www.bls.gov/news.release/cesan.htm>).

consumer attention is intense, and the quality and choice of entertainment options are key features that drive consumer demand. Allowing consumers to watch what they want to watch, where and when they want to watch it, and improving the resolution quality have all proved important. Furthermore, consumers are directly involved in making the choices of when and how to consume entertainment video.

Another key feature of video entertainment (and many other information goods) is that there is a high and irreversible (sunk) first-copy cost to produce the content, with low incremental costs to make additional copies available for consumption. This is what provides a key economic justification for copyright and other intellectual property rights law, and creates the need for digital rights management (DRM) technologies. Because the costs of producing content are incurred up front and sunk once the content is created, and often before its value is really known,⁹⁵ policymakers have created a tradable property right to the content that allows content creators to require payment for legal distribution of the content.

This helps ensure that content creators will have an incentive to invest in creating good content since it protects their ability to appropriate some of the value for their content if, once produced, it actually finds a market.⁹⁶ The desire to maximize the revenue to be realized leads to the practice of windowing, or market segmentation, discussed earlier. DRM helps facilitate the segmentation of markets. In many cases, the DRM is embedded in the contracts and distribution agreements, but sometimes it is implemented using encryption or video-coding technologies.⁹⁷ Content rights holders seek to segment the market along all possible dimensions of the consumption experience (the resolution quality, bundled vs. a la carte pricing, the platform delivered on, the location where viewing takes place, etcetera). Additionally, because digital media, once created and once on-line, can be shared at low incremental cost, the potential for losses due to piracy (i.e., illegal free goods competing with legitimate sales) are significant. Also, when attempting to capture revenues in different markets, content rights holders need to protect against self-cannibalization, which occurs when low revenue alternatives compete directly with higher-priced alternatives.

The proliferation of distribution channels, and increasingly of over-the-top video service providers complicates efforts to price discriminate and segment markets. For example, content providers typically received higher prices from cable television distributors (like

⁹⁵ It is hard to predict audience interest in entertainment content until audiences can experience the content, but by then the production costs have already been incurred.

⁹⁶ Like with venture capital, many of the programs produced fail to recover the costs associated with their production. To make up for this, program producers often invest in portfolios of programs with the few hits making up for the many programs that fail to be successful. This is another form of bundling.

⁹⁷ For example, encrypted content can be sold with different usage rights that are only enabled with the appropriate key. Or, video distributors may offer different quality resolution programming at different prices; or with and without commercials; etcetera.

Comcast) than from over-the-top distributors (like Netflix), in part because when Netflix initially negotiated its content deals, it was not expected to be as successful as it has become and because much of the content is not premium content (e.g., first-run movies and sports).⁹⁸ Although the movie libraries differ there is significant overlap in the movies available via Netflix and cable television, leading a growing number of subscribers to ask whether it makes sense to continue to pay for a cable television service (at close to \$100 per household per month⁹⁹) that includes many more programming options to select from, but which may include much that is not of interest to particular consumers. For a growing number of subscribers, a skinnier bundle of programming that offers fewer options but is comprised of Netflix (about \$9 per month), free over-the-air broadcast TV (available in digital format that is of equivalent or better quality than may be available from cable TV), and one or two other pay-per-view channels of special interest may be regarded as a superior option. Of course, the cable programming company bundles the cost of the programming and the digital delivery of the video via its broadband platform, whereas Netflix requires its streaming customers to pay for the broadband Internet access service used to receive its programming. Not surprisingly, Netflix would like the cost of that broadband Internet access to be as low as possible.

Most economists may agree that entertainment consumers benefit from the expansion of video programming options and competition that lowers pricing, so long as the competition does not adversely impact the ability of content creators to continue to create good quality content,¹⁰⁰ but it is unclear what contribution all of this competition makes to the overall economy if most of the effect is to shift revenue from one set of providers or rights holders to another. If the budget shares of entertainment expenditures are relatively constant and not increasing in aggregate as much as they are shifting from one type of media to another (e.g., from print to video media, from legacy channels to Internet), then the potential for this to generate jobs and economic growth is also limited.

In contrast, Smart X Internet services are expected to enhance the efficiency of transportation systems, household HVAC systems, improve healthcare, and facilitate

⁹⁸ One report claims that subscribers pay only \$0.20 per hour for Netflix content compared to \$0.61 per hour for cable-television content (see SlashDot, July 27, 2016, available at <https://news.slashdot.org/story/16/07/27/0046200/subscribers-pay-61-cents-per-hour-of-cable-but-only-20-cents-per-hour-of-netflix>).

⁹⁹ Leichtman Research Group reported that the average household that subscribed to pay TV services in 2015 paid \$99.10 per month (see Leichtman Research Group (2015), "83% of U.S. Households Subscribe to a Pay-TV service," Press Release, September 3, 2015, available at <http://www.leichtmanresearch.com/press/090315release.pdf>).

¹⁰⁰ A number of musical artists have looked with dismay to the rise of streaming music sites like Spotify that pay per song actually listened to at rates that make it difficult for artists to earn a living wage (see Resnikoff, P. (2015), "My song was streamed 178 million times. I was paid \$5,679," Digital Music News, September 24, 2015, available at <http://www.digitalmusicnews.com/2015/09/24/my-song-was-played-178-million-times-on-spotify-i-was-paid-5679/>).

telecommuting. The potential for those sorts of innovations to yield significant economic benefits seems very promising. The expenditures and time budgets for these activities come out of the rest of – and by far the larger share of – consumer time and budget expenditure shares. But whereas the consumer expenditures on entertainment services are direct, many of the expenditures associated with Smart X may be indirect and may be mediated by a third-party (e.g., a doctor for healthcare, the power company for energy-efficiency in the home, or an employer for telecommuting) and be contingent on other decisions (e.g., durable goods purchase of an electric car, home improvement investment in renewable power, or a change in employment or how work is organized).

The economics driving these other Smart-X related activities are diverse, and in many cases, may not be viewed as discretionary in the same way as consumer decisions about what entertainment consumption to engage in. For example, a consumer's investment in smart-home technologies may be motivated by the desire to save money on heating and electricity and to improve the quality of the consumer's home which may involve a long-lived capital investment and imply a lifestyle change. Alternatively, the expenditure to act on a doctor's recommendation for using Smart-X technology for at-home monitoring for an elderly relative will be viewed as healthcare related, and not as a leisure expense.

Moreover, for many of the goods and services that are associated with Smart X, there may be only limited need for DRM or price discrimination capabilities associated with the underlying digital delivery service. That is, the bundling of the digital distribution and the service provides less value for Smart X than for entertainment goods, where such bundling of the content and distribution channel is often intrinsic to how the rights holders hope to capture value. With Smart X, the need is for a basic communications capability. Additionally, whereas there may be strong reasons to protect the security and privacy of information flows associated with Smart X (e.g., to protect the integrity and reliability of the electric power grid or secure the privacy of healthcare data), the security/privacy concerns associated with entertainment media seem far less pressing. Indeed, with much entertainment media, the popularity of the media may depend on fads and shared social interest. Knowing what other people are watching can result in positive feedback loops and bandwagon equilibria that make the content even more popular; whereas the same is often not the case with Smart X information flows.¹⁰¹

Also, it is worth noting that the intellectual property regime associated with copyright is not troubled by the creation of monopolies. A rights holder may set the prices for access to content at whatever the market will bear, and collectively, we rely on robust competition for entertainment content to discipline the pricing of entertainment media. It is only with a few special types of entertainment media that significant concerns regarding pricing and market power typically arise, with sports programming being the most notable example. Taken together, these features make price discrimination and

¹⁰¹ Entertainment content may be a search good so hearing that lots of others like a show can signal that the show is higher quality and increase the audience. Also, entertainment is a social experience and the shared experience can contribute to its value. Th

“content” management key attributes of entertainment media products. A video-bit is not just a bit – it depends on where/how it is viewed and by whom. All of those attributes need to be managed to extract consumer surplus and the network can help support that fine-grained management.

In contrast, the Smart-X services may have need for basic communications functionality, but with the much richer support for expanded abstractions for what constitutes electronic communications (any-to-any, real-time and asynchronous, mixed media, planned and ad hoc) that the Internet provides.

With Smart-X, it is unclear what the right model is for bundling services, but it may make more sense generically to have consumers pay for the broadband Internet access as a general service that may be used by other businesses. For example, the smart lightbulb company or refrigerator company, may wish to make use of the Internet access to the home, but would be unlikely to bundle that service with its offering in the same way as the New York Times or Netflix may choose to bundle (or not) the cost of delivery in its basic service (e.g., both bundle delivery costs into home delivery of newspapers and DVD disks, but have the consumer self-provision broadband access for on-line access).

Before 2000, when broadband home access began to be generally available, it was relatively rare for employees to have high-speed digital connections that allowed them to work from home. In the cases where employees did have such access, it was typically provided as a business-grade service that was paid for by the employer (e.g., a T1 line). Similarly, many of the employees with personal computers or cellular phones were provided with those devices by their employers. As broadband, computers, and cell phones have become mass-market consumer appliances, it is less common for users to separate business and personal use across their devices. With the high-penetration of broadband into homes, a growing number of broadband subscribers are telecommuting and it is harder to identify what share of the value of having home broadband should be attributable to personal home use and to business use.¹⁰² Moreover, video-conferencing and other rich, interactive (2-way) multimedia work applications are making telecommuting more productive. This blurring of the boundaries between what is work-related versus personal or entertainment-related infrastructure makes it more challenging to separate what broadband Internet capabilities and consumer expenditures may be needed for Smart-X versus for entertainment media.

Summing up, entertainment media and Smart X economics are sufficiently different that it is reasonable to anticipate that they might give rise to distinctly different sets of network requirements if each were the sole motivator for investing in Internet infrastructure; however, separating out these motivations may be increasingly impractical

¹⁰² In 2015, 24% of all workers did some or all of their work from home, up from 19% in 2003; and, for managers, the share who worked partly from home was 38% in 2015 (see BLS (2016), "American Time Use Survey – 2015 Results," Press Release, Bureau of Labor Statistics (BLS), June 24, 2016, available at <http://www.bls.gov/news.release/atus.nr0.htm>).

as personal and business uses blur and as the rationale for bundling services on a common all-IP platform converge. However, the distinct differences in the economics for the ultimate goods and services that are giving rise to the traffic are likely to cause tension for how best to design and price different Internet capabilities such as price discrimination/DRM, security, basic transport, mobile access, and reliability.

5. Policy

Even if one concludes that technical and market forces are compelling us toward convergence and that that is desirable, we would still need to confront the difficult challenge of how to bring about the necessary policy reforms. Historically, in keeping with the silo-nature of legacy networks and the markets they served, and the public policy issues of greatest concern that naturally arose in each of those silos, separate regulatory frameworks evolved for each: PSTN regulation for the telephone networks and other telecommunication services; Broadcasting, cable television and media regulations for television services; and a mix of PSTN and computer industry regulation for data services and the Internet).

In the following, we briefly characterize some of the different concerns that arise in each of these domains that will pose problems for convergence at the policy-level, we discuss one model for how regulatory policy might deal with separating most entertainment video traffic from other broadband Internet traffic, and speculate about what this might mean for several of the policy issues the FCC is grappling with currently.

5.1. Separate policy domains and the challenge of Convergence

In the following three sub-sections we briefly characterize the traditional regulatory constructs that emerged to govern telecommunications (PSTN), TV (broadcasting/cable), and data (Internet), and today's "policy patch" for addressing convergence issues today.

5.1.1. PSTN Regulation

PSTN regulation evolved over the course of a century of building what originally was the analog, circuit-switched telephone network that became the AT&T Bell System, and which subsequently has been restructured several times to enable competition to proceed where feasible (first in equipment markets, then long distance services, then local services) and in response to changing technology, regulatory, and market dynamics.

The telephone companies owned the network facilities that provided the bidirectional, real-time connectivity that eventually made universal access telephone service feasible: any telephone number could establish a call to any other telephone number. These networks were multiplexed to facilitate sharing of backhaul transmission capabilities and support hierarchical switched routing. Over time, these legacy telephone networks evolved into the modern nation-spanning, global-reach telecommunications networks of today. This network of interconnected telephone networks became the PSTN, which was subject to strong regulations to ensure universal (affordable) accessibility to basic

telephone services and end-to-end connectivity (interconnection). The former concern gave rise to Universal Service mandates;¹⁰³ while the latter gave rise to a range of interconnection regulations impacting how telephone calls were interconnected across local, long-distance, and international networks and services.

In addition to providing end-to-end telephone (and other electronic communication services) for both mobile and fixed telephones, these networks provided the basic network infrastructure that supported bidirectional data services, that prior to the rise of the Internet as a mass market platform in the 1990s, was almost exclusively used by commercial, government and other non-consumer entities (e.g., universities). Private data lines and other business services were regulated as telecommunications services under the same general framework as the PSTN.

The telecommunication service providers were principally regulated under a two-tiered model in which state-level Public Utility Commissions (PUCs) and the Federal Communications Commission (FCC) set rules for how these providers offered services to both residential and commercial customers. The bifurcation into intrastate/local and interstate jurisdictions was due in part to (1) the local concerns associated with building out the last-mile infrastructure; (2) the fact that long distance services were open to competition whereas local services were still regulated as monopoly franchises; and (3) because the hierarchical nature of telephone networking facilitated differentiating between intrastate and interstate traffic.¹⁰⁴ With the rise of one-rate plans, mobile telephony (and later VoIP), and industry restructuring, the rationale and feasibility of separate jurisdictions for interstate and intrastate has eroded, with the FCC assuming principal responsibility for PSTN regulations. Over time, the FCC has evolved separate wireline and cellular bureaus to manage the extensive fabric of rules and regulations governing the technologies, services, and pricing of the various services provided over the PSTN.

Today, many of these legacy telephone networks are well-on-the-way to morphing into all-IP broadband networks capable of supporting the full spectrum of wireless and wired electronic communication services.¹⁰⁵ Many have already substantially replaced legacy

¹⁰³ In the U.S., universal service policies included federal and state subsidy programs, as well as "carrier of last resort" and "duty to serve" obligations imposed on local telephone service providers. There were also service quality, reporting, and pricing regulations that applied to different services.

¹⁰⁴ For further discussion of why the intrastate/interstate separation of jurisdictions is increasingly nonsensical, see Sicker, Douglas C. "End of Federalism in Telecommunication Regulations, The." *Northwestern Journal of Technology and Intellectual Property* 3.2 (2004-2005): 130-159; or Cooper, C. and B. Koukoutchos (2008), "Federalism and the Telephone: The Case for Preemptive Federal Deregulation in the New World of Intermodal Competition," *Journal on Telecommunications & High Technology Law*, 6 (2007-2008) 293.

¹⁰⁵ For further discussion of the challenges of morphing today's telecommunications regulatory framework for the broadband Internet future, see Lehr, W., D. Clark, and S. Bauer (2013),

copper wire distribution plant and switching systems, with fiber optic cables and software switches, which poses a challenge for regulatory authorities seeking to manage the retirement of the legacy PSTN infrastructure.¹⁰⁶ Increasingly, telephone services that were embedded in the fabric of the technology from whence the PSTN was constructed have become just another application that can ride on top of IP (i.e., Voice-over-IP or VoIP) that can be supported over diverse network infrastructures (wired or wireless, fixed or mobile, etc.). Some of the VoIP traffic today rides over the Internet, making use of broadband access services to provide connectivity to traditional telephones or softphone applications that may be running on multiple types of customer premise equipment. However, most of the VoIP services operate over managed IP networks that share many of the same physical and other network resources that support the Internet, but are managed separately for business and technical reasons. When cellular services are added to the mix, these divergent ways to deliver services that consumers regard as (imperfect) substitutes for each other are regulated with a patchwork of divergent regulatory permissions and obligations. For example, when cable television-based providers (like Comcast or New Charter) offer voice services, they provide those services over managed IP networks that interconnect with the PSTN and support emergency calling (e911) and other services that are not supported by pure Voice-over-the-Internet providers.

Moving voice services into the IP domain makes it possible to redefine the communications abstraction that characterized legacy telephone services. What used to be a real-time analog voice to analog voice call, can now be any-time (with voicemail), any-format (audio, text, or video), any-to-any (persons-to-persons for conferencing, or persons-to-machines) communications by integrating with other IP-enabled services that take advantage of the computer processing and storage capabilities increasingly available both in the core, and at the edge of modern telecommunication networks.

The core regulatory concern that arises in telecommunications is how to separate the regulation of the network services that provide the basic electronic transmission and connectivity functionality ("the conduit") from rules and regulations that may govern the information that is conveyed via the electronic communications ("the content"). The basic idea is that communications are about sharing information between end-points that control what that information is, while telecommunications is about providing the supporting network infrastructures to make such communication feasible. In the U.S., the

"Measuring Performance when Broadband is the New PSTN," *Journal of Information Policy*, Vol. 3 (2013), pg. 411-441 (available at: <http://jip.vmhost.psu.edu/ojs/index.php/jip/article/view/94>).

¹⁰⁶ The Telecommunications Act of 1996 required local telephone companies to unbundle core elements of their network infrastructure. This included unbundling access to the copper loops which were used by Competitive Local Exchange Carriers (CLECs) to offer DSL-based broadband and telephone services in competition with the Incumbents (ILECs). We are rapidly approaching the point – some would argue are past the point – where turning off the copper networks makes sense, but that raises the question of what happens to the CLECs and others who are still using the legacy equipment and the plethora of rules and regulations that were specifically crafted with the legacy PSTN and its embedded technologies in mind.

principle regulatory framework for managing the provision of telecommunications services has been Title II of the Communications Act that treats the regulation of such services under a common carrier framework.

How to separate content and conduit has proved to be an enduring challenge that has manifested in different ways as technologies, markets, and regulatory policies have evolved. One key challenge has been to determine which services should be regarded as Title II telecommunication services, which are subject to much more stringent regulatory controls, from information services, which are more lightly regulated under the Communications Act. Another challenge is to draw the line between the computer processing and telecommunications industries. A series of consent decrees between the U.S. Department of Justice (DoJ) and industry leaders in the sector helped steer the early development of the computing and telecommunications industries along separate paths.¹⁰⁷

In a world of all-IP broadband networks, the underlying technology does not limit the services that can be offered or how they might be classified from a regulatory perspective. This poses challenges for how to transition legacy regulations for telephony to the new world of any-to-any communications capabilities supported on modern all-IP broadband capable networks. How should today's separate regulatory frameworks which apply to fixed versus mobile voice, legacy TDM versus VoIP, Voice-over-the-Internet versus Voice-over-other-IP be harmonized and reconciled? What about when voice becomes video or text?

Ensuring continued access to (universal service) and connectivity (interconnection, interoperability) for essential electronic communication services are core policy mandates that relate to "conduit" network services. Determining which of these require regulatory interventions, and potentially new rules, and which may be left to market forces to be sorted out is both contentious and unclear. Managing the transition from the legacy copper-wire-based PSTN to the new all-IP broadband-capable PSTN presents a daunting challenge across a host of policy issues both near and longer-term.

Of special note are the need for such essential telecommunication services as public safety (e911 emergency services) and lawful wiretaps and investigations of electronic communications services (CALEA). The design of modern communications systems for public safety and law enforcement in today's post 9/11 world raises numerous technical and management issues that are far removed from the concerns that arise in the context of designing network services for the electronic delivery of entertainment media. Addressing issues where the principal users of the services are government employees (not for-profit businesses and their customers), where life and death situations are

¹⁰⁷ In 1956 and 1982, the DoJ reached consent decrees with IBM and AT&T respectively, that in part, limited the ability of the former to enter the telecommunications industry and the latter to enter the computer processing industry. The lines between these two industries has continued to blur as processing as become more distributed and telecommunications has become more intelligent.

common, and which may engage national security concerns (including access to classified information), presents a fundamentally different decision-making context than is confronted in the case of entertainment media.

The people who need to be "in the room" for the discussions about designing new e911 or CALEA rules and systems are unlikely to be the same as those who need to be there to discuss new video codec protocols or program licensing regimes for entertainment media. These discussions will, for the most part, engage different commercial and public entities, representing different stakeholder interests and requiring different types of expertise and knowledge about the legacy context that is being reformed. Of course, some of the key participants will be common across the different decision domains (e.g., the broadband platform providers); but others will not be (e.g., the producers and rights holders for video entertainment media).

5.1.2. Broadcast, Cable and Media Regulation

First radio and then television began as over-the-air services delivered via broadcast networks that provided channels of programming that could be received by customer-owned radios and televisions. The original idea was for manufacturers of the radios and televisions to subsidize the programming in order to stimulate demand for the consumer appliances. Relatively quickly, however, it became clear that the programming provided a vehicle for consumer advertising, and advertising-support became the principal funding vehicle to subsidize the delivery of broadcast media services. Consumer product companies paid for the programming in return for having access to consumer leisure attention to advertise their products. Other forms of pay-per-view entertainment media such as live performances, movie theaters, and print media used non-electronic distribution networks to reach their audiences and capture leisure expenditures.

In the world of over-the-air broadcasting, and with the technologies in use at the time, commercial broadcaster access to scarce spectrum (below 1GHz) was restricted and regulated by the FCC. In return for compliance with a range of public interest commitments (e.g., providing news services), over-the-air broadcasters were granted exclusive licenses to use the airwaves. As a consequence, the choice of networks or channels and the programming selections that could be provided in each local market, were severely limited.

Beyond spectrum regulation, the focus of media regulation was on ensuring that diverse programming choices were available. Programs promoting public interest and children's programming, rules restricting pornography, access for persons with disabilities (e.g., closed captioning requirements), and support for emergency public notification services (e.g., to make broadly accessible public announcements in the event of a national emergency), needed to be balanced with First Amendment protections for free speech and media access. In contrast to telecommunications regulations that focused on the "conduit," broadcasting and media regulations focused on the "content."

With the rise of cable television systems, over-the-air RF signals were increasingly moved to higher capacity wired (coaxial) cable RF systems for local delivery. Cable

television networks were granted local monopoly franchises in return for providing public access services (e.g., network services for the municipal government) and meeting community coverage commitments. The new cable networks enabled a large expansion in the number of programming channels, including redistribution of higher-quality signals for the legacy over-the-air TV stations. Rather than separating the network services (the conduit) and the programming (the content), a new framework was crafted for regulating cable television service providers as a new category of providers. This became Title VI – Cable Communications, which was added to the Communications Act in 1984.

Just as legacy telephone network providers were updating their networks, so too were cable television providers. They were motivated in part by the promise of being able to offer enhanced pay-per-view (PPV) television services and to respond to competition from Direct Broadcast Satellite (DBS) services, the impending threat of telephone companies offering TV services, and changing consumer habits for consuming entertainment media (including the use of VCRs to time-shift programs and bypass advertisements). The cable providers added two-way capabilities to their legacy one-way broadcast networks in order to be in a better position to offer video-on-demand and interactive programming options, while also positioning themselves better to compete for services that previously had been provided solely by telephone operators. This positioned them well to take an early lead in the market for broadband Internet access services that began to emerge after 1996.

Today, the largest cable programming providers are Multiple System Operators (MSOs) with all-IP broadband last-mile networks in markets across the U.S.. Their video services are regulated as Multichannel Video Program Distributors (MVPDs). The MVPD rules, which also apply to DBS and telecommunications-based providers of video programming, include a range of rules that seek to promote competition and protect access to programming. These include program access rules, must-carry/retransmission rights, and special copyright rules. The program access rules are designed to ensure that scarce programming rights are not monopolized by individual operators in ways that might threaten to foreclose other distributors from being able to offer potential subscribers attractive bundles of programming. The must-carry/retransmission framework was adopted to ensure local over-the-air broadcasters would be able to continue to reach the homes of cable subscribers in their coverage areas. Broadcasters have a right to require cable providers to retransmit local station broadcasts via the cable plant for stations that elect to invoke the "must-carry" right. Alternatively, stations can elect to negotiate for retransmission rights, which may result in cable networks having to pay for the right to retransmit local stations (or, depending on the outcome of the negotiation, for over-the-air stations to have to pay for redistribution). In recent years, payments of retransmission fees from cable companies to over-the-air broadcasters have been an important source of revenue for the broadcasters and a source of increasing program-related costs for the MVPDs. Another key element of the MVPD framework is a special licensing arrangement for copyrighted material that allows MVPDs to distribute the copyrighted material under a single blanket licensing agreement, rather than being forced to negotiate separate usage rights arrangements with each copyright holder.

Although many of the largest MVPDs are also providers of voice and broadband Internet access services (e.g., Comcast, New Charter, Verizon, or AT&T), the video entertainment business they are engaged in as MVPDs includes many content providers who do not own, operate, or sell broadband network services. This includes media companies like Viacom, Disney, HBO, Netflix, and Paramount. Moreover, the "content" considerations that dominate their business decision-making are quite distinct from the "conduit" considerations that dominate telecommunications industry decision-making.

5.1.3. Data Communications and Internet Regulation

Prior to the emergence of the Internet as a mass-market platform for data communication services, these were almost exclusively targeted at enterprise customers (businesses, governments, universities, etc.). The telecommunication services were provided using facilities largely owned and operated by the same telephone companies that supported the PSTN and were used to interconnect the computing and data networks of the enterprise customers. Private lines and a range of other data communication services were regulated as Title II services, subject to the same sort of strong regulatory oversight as other PSTN services.

Meanwhile, the computer processing and services businesses were largely unregulated, or at least free from the sort of sector-specific regulatory oversight that characterized telecommunications and the media industries discussed previously. Industry standardization (to define interoperability standards to enable components provided by different providers to interconnect), trade agreements (to manage international commerce in computing equipment and services), and general competition regulations were the principal mechanisms for ensuring that market forces could keep the computing industry on track. When IBM was dominant (and later when Microsoft became dominant), antitrust suits brought by the U.S. Department of Justice (DoJ) helped regulate competitive dynamics in the computing sector.¹⁰⁸

In contrast to the legacy of telecommunications regulation associated with the PSTN or the media regulations associated with broadcasting or cable network (MVPD) regulation, the Internet was largely unregulated. It first emerged as an application that was supported over dial-up telephone lines on the PSTN, and since 2000, has emerged as broadband Internet access services provided by today's broadband service providers that have evolved from the convergence of telephone and cable television networks. From its legacy as an unregulated application that rode on top of the PSTN, the Internet is emerging as the platform for future communications and media services.

¹⁰⁸ In 1956, the DoJ reached a consent decree with IBM limiting their ability to abuse their market power over computing equipment, and also restricted their ability to compete in telecommunications sector (see Note 107 *supra*). In 1998, the DoJ brought an antitrust case against Microsoft alleging, in part, abuse of market power in the markets for web browsers. That case was ultimately settled in 2001.

As broadband and the Internet have come to be seen by many as *the* global platforms for all electronic communication services, a natural question emerges as to how to morph or converge the legacies of PSTN and broadcast/cable/media regulations with the mostly unregulated Internet. Combining the delivery of all of these services into the same "conduit" may make it more difficult to separate and manage needed policy transitions that address quite distinct and separable issues, requiring participation from different sets of stakeholders, with very different realms of expertise. Much of the legacy regulation may be discarded as no longer necessary or relevant in an all-IP world; while other legacy regulations will require new frameworks to reconcile with the continuing public interest mandates (e.g., to ensure universal service, interoperability, and competition where feasible). Addressing issues like program access, e911, CALEA, or Universal Service raise special challenges in the context of the Internet.

Today, we have a patchwork of regulatory constructs that have sought to address the challenges of convergence piecemeal. For example, the FCC, after initially seeking to exempt broadband access services from Title II regulatory oversight (by classifying broadband as an information service), has opted to reverse itself and classify a new Broadband Internet Access Service (BIAS) as a Title II service. It was motivated in part by earlier challenges to the FCC's assertion of authority to impose regulatory obligations on broadband service providers under Title I and other provisions of the Communications Act. Meanwhile, the FCC has chosen not to define VoIP as a Title II service, and has sought to regulate voice services through a series of targeted proceedings focused on addressing specific issues, such as how VoIP providers support e911 services and their obligations to contribute to universal service subsidy programs. A similar patchwork approach applies in the case of video distribution services. The MVPD framework and various provisions govern the distribution of video services over the same wires that are used to deliver broadband Internet access and voice services (whether VoIP or otherwise), but with different sets of rules applying to the different ways in which the traffic is delivered.

As noted earlier, most of the entertainment video delivered into homes is delivered over non-IP RF transmission services or non-Internet IP networks, but a rapidly growing volume of entertainment video traffic is moving into the Internet and coming into homes via the BIAS service. Once again, this is stressing the ability to separate content/conduit policy regulations, and blurring the boundaries between separate technical, business, and regulatory systems.

5.2. Relevance to Current Policy Issues

There are numerous current policy issues that underscore the challenges that arise as video traffic over BIAS grows. While we are not stating that any of these issues necessarily suggest that video should be subtracted from BIAS, they do highlight the policy and technical issues relating to entertainment video on the Internet.

In the following sub-sections we discuss a number of examples of policy issues that may be addressed and/or complicated by whether entertainment video remains integrated with the BIAS traffic, or is carried in a separately-managed IP network.

5.2.1. Performance measurement and managing consumer expectations

Delivering high-resolution entertainment video via the BIAS will require high data rates downstream to the consumer's home; and if the goal is to deliver multiple such streams, still higher data rates will be required. Moreover, if the content to be viewed is not cached relatively close to the edge of the network and has to travel across ISP interconnections, then there is a risk that congestion of those interconnection links might degrade performance.

In Bauer et al. (2015), we discussed the challenges for performance measurement that arise as the speeds for the typical BIAS connection rise to rates exceeding 100Mbps.¹⁰⁹ With such superfast broadband access services, it will not be reasonable to sustain those rates as average speeds to arbitrary Internet destinations, but consumers will expect to realize those rates for particular services (e.g., perhaps for Netflix) and to some set of destinations.

As discussed earlier, if the entertainment video were carried via a separate IP service (VIAS), then the data rates required for the BIAS might be lower and the measurement challenges and expectations would certainly be different. At the same time, knowing that the VIAS service was supporting entertainment video might allow for better performance measurements that are closely tailored to the requirements of entertainment video.¹¹⁰

5.2.2. Universal service

One of the key justifications for PSTN regulation was the recognition that telephone service is essential basic infrastructure for society and the economy, and thus, the government confronted a policy mandate to ensure universal and affordable access to telephone service for all citizens and businesses. This has resulted in the creation of regulatory obligations like "duty to serve" and "carrier of last resort" obligations, justified rate regulations that embedded implicit price subsidies (e.g., from long distance to local, from urban to rural, and from business to residential services, and justified the creation of a \$4 to \$8 billion dollar per year set of subsidy programs, funded by telecommunication service revenues, directed at supporting universal access to affordable telephone service.

It has long been recognized, and since the National Broadband Plan, has been national policy to determine how to transition a program focused on supporting telephone service access to one focused on supporting broadband access. The motivation for this is the growing recognition that broadband is the essential infrastructure that is needed to

¹⁰⁹ See Bauer, Steve, William Lehr, and Shirley Hung (2015), "Gigabit Broadband, Interconnection Propositions, and the Challenge of Managing Expectations," TPRC2015, Alexandria, VA, September 2015, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2586805.

¹¹⁰ Knowing what the traffic is and what quality of service is needed to ensure a good user experience simplifies the measurement challenge (design, interpretation, etc.).

support the full spectrum of electronic communication services, of which, voice telephony is just one application.

As noted earlier, the justification for why broadband is essential infrastructure makes frequent reference to the capabilities required to enable Smart-X solutions, but seldom or in much more muted tones, argues in favor of broadband so everyone can have access to more channels of high definition entertainment video.

If the Internet were not the principle service for delivering video, this would have a number of interesting implications. First, the capacity needed would be reduced, so potentially the overall investment required to ensure ubiquitous Internet access and subsidies desired to ensure affordable access might be less. Second, the standard for what might constitute adequate Internet access might be more easily met by cellular-based services, facilitating more symmetric regulation of mobile and fixed broadband Internet access service.

It is unclear how national support for Internet as an essential service, versus broadband as an essential service, might shift. Some might argue that an Internet largely stripped of entertainment traffic would be more easily justified as necessary for economically productive Smart-X activities, and so would be even more essential to preserve as an open, ubiquitously available platform for innovation. Some may view access to Internet enabled electronic communications as essential human rights, while viewing the ability to access entertainment media as merely something that may be nice to have but is not worth expending public subsidies on. Others may disagree and argue that a sufficient share of the populace wants access to high-quality entertainment video such that providing ubiquitous access ought to be a national priority.

These divergent views might be addressed by separating universal service support for broadband and universal support for Internet access. However, separating universal support for Internet and broadband access will not be easy in practice. For example, it is unclear to what extent video revenues currently subsidize Internet infrastructure. As noted earlier, the broadband networks are shared across triple-play offers and the total revenues captured by the network operators need to be sufficient to recover network operators' economic costs for their businesses to remain viable – but there is no single best way to allocate the shared costs back to the individual services. To the extent public funds may be made available to subsidize Internet infrastructure (instead of broadband infrastructure that may also support the delivery of entertainment video), ISPs may have a perverse incentive to shift network costs to the Internet service. On the other hand, to the extent video service revenues are contributing to covering network costs, the loss of those revenues (potentially due to competition from over-the-top video services) may force operators to shift the cost burden to the remaining services, including the Internet access service.

5.2.3. Internet Interconnections

Historically, Internet interconnection was not regulated. A tiered structure of peering and transit relationships emerged with so-called Tier 1 ISPs at the top, exchanging traffic via

revenue-neutral peering agreements. Other providers, lower down the hierarchy, purchased transit services to ensure delivery of traffic to all public Internet addresses. The rise of video and other developments in the Internet ecosystem have resulted in significant changes in how traffic is managed in the Internet and the types of interconnection arrangements that govern the exchange of traffic among ISPs. There has been a proliferation of new types of arrangements with peering occurring lower down in the hierarchy, and involving payment flows and special traffic management requirements.¹¹¹ The Internet topology has flattened and we have seen the rise of hypergiants that include the largest ISPs that provide broadband access services (e.g., Comcast) and the largest content providers and content provider networks (e.g., Netflix, Google).¹¹²

Much of the traffic that has driven the changes in the Internet interconnection ecosystem is entertainment video, as already noted. This traffic has been highly asymmetric (from content providers into access ISP networks), causing stresses for the legacy model of revenue-neutral peering and resulting in congestion that deteriorates the end-users' quality of experience when using broadband Internet access services. Figuring out who should pay and how payment should be made for expanding these interconnection links has led to well-publicized disputes between providers, and to calls for increased regulation of interconnection.¹¹³

5.2.4. MVPD Regulatory Reform

With the rise of Other Video Distributors (OVDs) like Hulu, Netflix, and YouTube, the legacy framework for managing video services is increasingly under stress. The OVDs do not confront the must-carry/retransmission, special access, or rate regulations that apply to MVPDs, but also do not benefit from the program access and copyright licensing provisions that benefit many MVPDs.

Extending the MVPD obligations to OVDs risks imposing media regulations on the Internet, while retaining the two-tiered regulatory structure risks distorting competition and investment incentives in the (mostly) highly competitive video entertainment industry. The differences in regulatory treatment are understandable in light of the history of how these businesses have evolved, but become increasingly untenable as the services become closer substitutes that are competing directly for the same pool of consumer dollars. At the end of 2014, the FCC launched a proceeding to address whether certain

¹¹¹ See Clark, D., W. Lehr, and S. Bauer (2011) "Interconnection in the Internet: the policy challenge," 39th Research Conference on Communications, Information and Internet Policy (www.tprcweb.com), Alexandria, VA, September 2011 (available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1992641).

¹¹² See "ATLAS Internet Observatory," Arbor Networks, November 2009, available at <https://wiki.tools.isoc.org/@api/deki/files/2426/=isoc-researchers-lunch.pdf>.

¹¹³ See Bauer, S., W. Lehr, and D. Clark (2012), "A Data Driven Exploration of Broadband Traffic Issues: Growth, Management, and Policy," TPRC2012, Alexandria, VA, September 2012 (available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2029058).

OVDs should be classified as MVPDs.¹¹⁴ Industry stakeholders differed widely in their support for the proposal, with Verizon and programmers favoring it, while cable operators (via the NCTA) and Amazon opposing it.¹¹⁵

To the extent the principal concerns related to the entertainment video programming industries, it might be easier if the traffic were not carried in the BIAS but via a separate service that might be more easily integrated into a common framework that applied to video distribution services in the last-mile (i.e., a VIAS)

5.2.5. Set-top boxes

Set-top boxes are another hot point of contention being addressed by the FCC. In February 2016, the FCC proposed new rules that would require MVPDs to unbundle the set-top boxes that are used to control access to the programming that MVPDs deliver.¹¹⁶ Today, most subscribers rent their set-top boxes from their MVPD providers. The FCC proposal is to establish rules that would require MVPDs to provide three streams of information that could be accessed by third-parties and would allow consumers to use their own set-top boxes. The three streams of information would include (1) service discovery information which would include information about the channel line-up and the available selection of programming; (2) rights information about what permissions were associated with the programming restricting its use by consumers; and (3) the video programming stream. Proponents of the rules argue that they would increase competition in video services, including allowing OVDs to better integrate programming with other new Internet services. Opponents argue that the rules impose additional regulatory costs and obligations on an industry that are unnecessary or worse. They point to the vibrancy of competition and the potential risks that such rules may pose for Internet technology, consumer privacy and the video distribution marketplace.¹¹⁷ Industry opponents have proposed abandoning the set-top boxes as gateway devices that control access to video programming distributed via last-mile networks in favor of an approach that would rely

¹¹⁴ See FCC (2014), "Notice of Proposed Rulemaking," In the Matter of Promoting Innovation and Competition in the Provision of Multichannel Video Programming Distribution Services , MB Docket No. 14-261, released: December 19, 2014, available at https://apps.fcc.gov/edocs_public/attachmatch/FCC-14-210A1.pdf.

¹¹⁵ See Silbey, M. (2015) "To be or not to be an MVPD," LightReading, April 6, 2015, available at <http://www.lightreading.com/video/ott/to-be-or-not-to-be-an-mvpd/d/d-id/714880>.

¹¹⁶ See FCC (2016), Notice of Proposed Rulemaking and Memorandum Opinion and Order," In the matter of Expanding Consumers' Video Navigation Choices (MB Docket No. 16-42) and Commercial Availability of Navigation Devices (CS Docket No. 97-80), Before the Federal Communications Commission, Released February 18, 2016, available at https://apps.fcc.gov/edocs_public/attachmatch/FCC-16-18A1.pdf.

¹¹⁷ For example, opponents argue that the new proposal would allow third-parties access to consumer viewing data, posing additional threats to consumer privacy and allowing those parties to gain free access to programming they did not pay for (see Lenkov, O. (2016), "The FCC Hoists the Jolly Roger on Your Cable Box," Wall Street Journal, June 8, 2016).

on open-standards-based software applications that could run on TVs and other devices to control access.¹¹⁸

While it is unclear what will happen, the debate over the set-top boxes highlights the potential risks to Internet openness when multiple 800-lb gorillas are tussling over the future of the way digital entertainment media gets delivered to households over last-mile networks. These issues would not go away, but might be rendered more tractable from a policy perspective if the BIAS traffic and the entertainment video traffic were delivered via separate channels.

5.2.6. Summing up the policy challenges

In the preceding five sections we briefly reviewed current policy issues that are impacted by the fact that entertainment video is integrated with Internet traffic in a hybrid model of patchwork policies. This is helping drive discussions across the range of communications policy issues – from measurement to interconnection, from universal service to media regulation. Although we have only skimmed the surface of how entertainment video concerns may complicate crafting good policy for either the Internet or markets for entertainment media, we believe engaging in the thought experiment of asking "what if the entertainment video traffic and its proponents were not in the room?" is a useful thought experiment. As we have suggested, there are enough arguments on both sides of the debate to suggest that we are not ready ourselves to conclude whether things would be better or worse if the video traffic were carried in a separately managed network. However, we believe that it may still be possible and there may be benefits in keeping the regulatory challenges distinct, at least in the near to medium term.

5.3. A framework for separating Broadband and Video

In this section, we take up the question of how might a separate video service (from the Internet) be regulated and what some of the challenges would be. The most obvious model would be to mirror the approach that the FCC used in crafting a regulatory framework for BIAS.

Today, cable operators' use of RF for video is governed by MVPD regulations; and, use of broadband service to make content like Xfinity available is covered by the FCC's Open Internet Order.¹¹⁹ The regulatory posture is that a Comcast subscriber should have no better experience when accessing affiliated Xfinity content than when accessing content from other providers (e.g., HBO, Netflix) by virtue of how packets are managed in the BIAS flow. Of course, if Netflix under-provisions its servers or the interconnection links, then the QoE might be different (worse) for Netflix relative to Xfinity for that subscriber.

¹¹⁸ See Eggerton, J. (2016), "NCTA Pitches 'Ditch the Box' Set-Top Proposal," Multichannel News, June 16, 2016, available at <http://www.multichannel.com/ncta-pitches-ditch-box-set-top-proposal/405730>.

¹¹⁹ See FCC (2015), Note 84 *supra*.

At the same time, the BIAS-delivered content is being delivered (in many cases) over the same wires but over different technologies (often not IP) or over separately managed IP services (e.g., IPTV) that are not part of the BIAS flows. This gives rise to additional QoE differences, and business relationships that are difficult to regulate without distorting market competition and investment incentives across the entertainment media value chain.

In principle at least, the FCC could consider defining a new Video-over-IP Access Service (VIAS) that would be a new BIAS service. As with BIAS, the FCC could preempt state regulation, asserting plausibly that the basic service was inseparably interstate, and thereby limit the risk that excessive use of legacy Title II authority might be applied.

Postponing for a minute the sizable challenges confronting such an approach, it is worthwhile considering what this might enable. First, were VIAS a viable regulatory model, then this would provide an obvious trajectory for what to do with OVD reclassification and the MVPD regulatory framework. This might help isolate the Internet (BIAS) from the effect of spillover of entertainment industry regulatory issues. Likewise, the FCC may be able to define a special requirement for VIAS interconnection to isolate pressure to regulate Internet interconnection arising from the incredible growth of asymmetric video traffic over BIAS services. Third, a separate VIAS service might be better positioned to add-on features and capabilities for service and content management that are video-specific. This could include such things as better in-network support for Digital Rights Management (DRM) and video-specific encoding support (variable bitrate encoding to address different video standards). This might spur the development of better and more capable IPTV network support that would be optimized for the needs of entertainment businesses.

Of course, the above regulatory model would confront numerous challenges, so that it might not even be politically/legally feasible, even if a case could be made for the economic merits of such an approach. First, unlike with the BIAS service, where the regulatory framework was imposed on an existing marketplace with well-developed broadband services already being provided by ISPs; the marketplace for VIAS services is not yet developed. It is uncertain that either the ISPs or content-providers who would be expected to make use of the service would prefer it to the alternative of either unregulated IPTV delivery, or strong protections for video delivered over BIAS. The FCC would lack an existing service model to point to in order to specify the new VIAS framework. Moreover, ISPs would likely resist any attempt to extend FCC regulatory authority over additional services, and their opposition would find significant political support in the U.S. Congress from politicians strongly opposed to almost any sort of government regulation. If, in addition, the content providers thought they could ensure a better regulatory deal for themselves by pushing for stronger BIAS protections, then the prospects for advancing such an idea would be even more daunting.

Second, even if a VIAS framework were created, it would not eliminate the need to balance the sharing of cost recovery burdens and resource allocations of the broadband

platform infrastructure between BIAS, VIAS, and potentially other IP-based services delivered over the broadband wired infrastructure. Earlier we discussed the technical options for allocating capacity between services. These arise in any case, so they would not be a *result* of creating VIAS, but might be more easily addressed under the VIAS framework.

Third, it is unreasonable to expect that all video content, or more narrowly, entertainment video content, would migrate to VIAS, which some might use as a basis for arguing that the framework would not provide a complete solution. Indeed, the VIAS approach is really just a different framing of the current patchwork situation, in which the largest share of entertainment video traffic is regulated under a different framework than Internet traffic. Our goal in suggesting that regulatory options exist for sustaining a world in which traffic is not fully converged into a single BIAS service is that it may offer advantages either in the long term, or at a minimum, in establishing a good glide path for regulatory reform if the longer-term horizon goal is for full convergence.

6. Summary Conclusions

In this paper, we took up the question of whether the convergence of entertainment media traffic and the Internet is desirable or inevitable. To investigate this question, we explored the different ways in which this convergence may occur at the technical, business/market, or policy level. This leads us to conclude that the jury is still out.

Although most of the traffic on the Internet is already entertainment media and that traffic is growing rapidly, most of the entertainment media delivered into homes is not being delivered via the Internet. We still have a long way to go before all entertainment video is delivered via IP ("everything over IP"), or even stronger, over the Internet ("everything over a single IP network"). It is far from clear whether getting to either of those points is a good thing for either the Internet or for video entertainment services, and it is certainly not inevitable. While we feel comfortable in those assertions, figuring out how much convergence is appropriate and at what level is a much more difficult question, and one we do not have an answer to.

Instead, we propose that it will be helpful to continue to examine the counter-examples to convergence. Exactly when folks say "everything over IP" or "everything over the Internet" is the right answer, pause and ask yourself how that may be wrong. While convergence brings many benefits, it also poses unavoidable challenges. Some of those may be fundamental in the bringing together of stakeholders with divergent interests, capabilities, and constraints; and some of those may be the pains of transitioning from yesterday's world to whatever comes next. By continually investigating the alternatives to convergence, we may be better positioned to address the challenges that will continue to arise.

7. Exhibits (landscape)

7.1. Figure 1: Qualitative Characterization of Video Traffic Types

	Communications	Sensing	Entertainment Video
Traffic direction (predominant)	Symmetric	Asymmetric (upstream)	Asymmetric (downstream)
Real-time	Yes, but may also have demand for cacheable)	Variable, lots may be interruptible	No (mostly, sports/breaking news an exception; new media sometimes)
Bandwidth			
Total (MB for duration)	Mixed (moderate). High-resolution requires a lot of bandwidth, but communication applications may be more tolerant generally of lower resolution quality	Low, but variable depends on resolution which depends on purpose for sensing. Lower resolution may be acceptable for most sensing applications.	Mixed (larger). Video files may be big, but resolution is mixed and short-form files may be small. Trend is toward higher quality and high-resolution is more important for entertainment video.
Upstream data rate	High (since inherently 2-way)	High (since inherently one-way upstream)	Low (since inherently one-way downstream)
Downstream data rate	High (since inherently 2-way)	Low (since inherently one-way downstream)	High (since inherently one-way upstream)
Interconnection	High (since any-to-any)	Depends where cached	Depends where cached
Cacheable	No (but recording/playback capability likely desirable)	Variable. A lot is cacheable, but some is real-time. Edge caching for latency intolerant apps may be important.	Yes. Mostly cacheable, but some may also require real-time resources.
Communication Model (Routing)	Any-to-Any (including many-to-many)	Mixed, but classic may be many-to-one (lots of sensors talk to aggregate collector) or mesh (IoT)	Broadcast. (Sources are fewer than destinations – viewers).
Wireless (5G) support	Not necessary, but nice to have	Yes. Ubiquitous coverage and mobility likely to be important.	Less important than for communications

7.2. Figure 2: Entertainment Video versus “Other 20%” traffic characterized

	Entertainment Video	“Other 20%”
Traffic direction (predominant)	Asymmetric (mostly downstream, limited upstream capacity needed)	Symmetric (2-way needed, but mixed traffic. Heavier upstream capacity required)
Data rate (Average, Peak)	High peak (High-resolution) and high average because video is transmission resource intensive generally, and especially so for high-resolution.	Mixed. “Other 20%” includes all types of traffic (all non-video and non-entertainment video traffic)
Burstiness (ratio peak to average data rate)	Less. Entertainment media viewing tends to be relatively predictable with flash crowds possible.	More. Mixed traffic that may be event driven (require unpredictable bursts of data of variable size) or require low latency guarantees even during aggregate peak, requiring extra capacity headroom to manage
Real-time/Cache-ability	Mostly not realtime since most traffic is cacheable (even “live” can often be near-real-time, and time-shifting is common).	Mixed and in-network storage may be needed solely to support disruption tolerant networking.
Routing	“Broadcast”-like (but not always multicast since that does not support VoD). Fewer video sources (content providers) than viewers (destinations downstream). Hierarchical tree.	Any-to-any mesh, especially to support IoT.
Wireless	Increasingly desired as mobile video fastest growing segment, but typically much lower resolution. Large displays typically wired.	Necessary since ubiquitous coverage and mobility required for many key applications (e.g., consider public safety, IoT).
Upstream	High (since inherently 2-way)	High (since inherently one-way upstream)