

Costs of Neutral/Unmanaged IP Networks

RICHARD N. CLARKE *

AT&T

Abstract

An important issue in “network neutrality” is the degree to which networks that comprise the Internet may be “managed” by their operators. This paper examines the current costs of unmanaged/neutral network technology and demonstrates that as customer Internet usage patterns evolve and become both more bandwidth-intensive and real-time oriented, a model of pure neutrality would be extremely expensive for consumers. Unmanaged network capacity sufficient to accommodate expected consumer demand growth currently would result in bills between \$143 and \$416 per month – depending on the customer’s level of TV viewing. Such fees would likely make unmanaged/neutral Internet service commercially prohibitive.

1 Introduction

The words “network neutrality” connote many things to many people. To some proponents, network neutrality requires a uniform flat rate structure where user charges are based strictly on maximum connection speeds. Others interpret network neutrality to mean that broadband access providers receive reimbursement only from end user customers attached to their networks and not from the content providers that send traffic to these end users. Still others argue that network neutrality requires that all Internet Protocol (“IP”) networks interconnect freely with one another and without payment of any explicit compensation (“peering”).

While the above are just a few of network neutrality’s recurrent flavors, there is another frequently mentioned flavor: network management. Briefly stated, IP networks may differ in the extent to which they are “managed”. Such management by IP network operators can range from simple maximum flow limits on links, to prioritization of packet transport and delivery by service type or payment class. Such management may be used to relieve network congestion, or also to improve capacity utilization through the network. Certain network neutrality proponents suggest that the use of any network management controls beyond simple byte-count or flow limits would violate principles of network neutrality and would reduce the social and technological benefits offered by the Internet.

* Mailing address. AT&T, 1133 21st Street, NW, Suite 900, Washington, DC 20036. Email: mclarke@att.com The analyses and conclusions presented in this paper are solely those of the author and should not be construed as reflecting any official analysis or position of AT&T. I am grateful for discussions with Richard Bennett, David Burstein, William Norton, Andrew Odlyzko, Marvin Sirbu, with many colleagues at AT&T and with the editors of this issue, Philip Weiser and Marius Schwartz.

These proponents generally argue that investment in network management capabilities can be avoided if network providers simply focus their investment on building “fatter” pipes. They reason that if networks were just built with enough capacity to handle any level of offered traffic load, there would be no need to utilize network management controls to maintain good quality of service.^{1,2}

While the current debate over network neutrality has been largely theoretic, with supporters of neutrality generally unable to point to any current widespread violations of basic principles and network operators not yet needing to engage in too much network management, this may be changing. Supporters of “no network management or prioritization” point to Comcast’s application of network management controls against BitTorrent traffic as demonstrating the immediacy of the issue.³ Further, these interests have introduced bills in Congress to codify a “dumb pipe” engineering model of the Internet into U.S. law.⁴ At the same time, growing demands for extremely high bandwidth peer-to-peer (“P2P”) or quality-sensitive Internet-delivered applications and services such as streaming video are beginning to overwhelm the capability of current networks to service these demands efficiently using “best effort” traffic engineering.⁵

¹ Many articles offer a variety of arguments for opposing IP carrier application of network management controls. Examples include Chester (2006), Dvorak (2006), Lessig (2002), Garfinkel (2003) and Wu (2003). In addition, see, D. Isenberg and D. Weinberger, “The Paradox of the Best Network”.

(<http://www.netparadox.com/netparadox.html>) as well as information sheets issued by various interest groups such as Common Cause (http://www.commoncause.org/atf/cf/{FB3C17E2-CDD1-4DF6-92BE-BD4429893665}/NETWORK_NEUTRALITY_FACTSHEET.PDF), Save the Internet (<http://www.savetheinternet.com/=faq>) and Media Access Project (<http://www.mediaaccess.org/programs/broadband/index.html>).

² While positions expressed more recently by neutrality proponents, such as Public Knowledge *et al* (2007, pp.13-14), suggest that certain non-discriminatory management controls might be permissible, it is difficult to see how such controls could operate efficiently. If network congestion exists because two real-time applications are contending for the necessary bandwidth, providing each application with a (non-discriminatory) half slice of the available bandwidth would likely result in neither application executing successfully. Similarly, alternately allocating a full slice of bandwidth to contending applications may also result in unsatisfactory performance (for example, on even-numbered days football game video reception is clear and news reception is poor, while on odd-numbered days news reception is good and football reception is poor).

³ See, Svensson (2007).

⁴ See, “Internet Freedom Preservation Act” introduced by Senators Dorgan and Snowe on January 9, 2007. (http://dorgan.senate.gov/documents/newsroom/net_neutrality.pdf). Section 2(a) of this bill states that “each broadband service provider shall – (1) not block, interfere with, discriminate against, impair, or degrade the ability of any person to use a broadband service to access, use, send, post, receive, or offer any lawful content, application, or service made available via the Internet; ... (4) enable any content, application, or service made available via the Internet to be offered, provided, or posted on a basis that – (A) is reasonable and non-discriminatory, including with respect to quality of service, access, speed, and bandwidth; (B) is at least equivalent to the access, speed, quality of service, and bandwidth that such broadband service provider offers to affiliated content, applications, or services made available via the public Internet into the network of such broadband service provider; and (C) does not impose a charge on the basis of the type of content, applications, or services made available via the Internet into the network of such broadband service provider; (5) only prioritize content, applications, or services accessed by a user that is made available via the Internet within the network of such broadband service provider based on the type of content, applications, or services and the level of service purchased by the user, without charge for such prioritization; and (6) not install or utilize network features, functions, or capabilities that impede or hinder compliance with this section.”

⁵ Under “best effort” engineering, all packets receive transport based simply upon whether capacity exists at the instant they request it. If no capacity exists, the packet is either buffered and delayed, or dropped. The network operator provides users with no guarantees that any particular packet will be delivered successfully

The analysis presented in this paper attempts to provide a crucial element that has been largely missing from this debate over network management: the resource costs of building a broadband network that is able to meet evolving capacity demands for delivery of high quality Internet television services without the efficiencies offered by network management.⁶ The cost model developed in this analysis suggests that customers seeking Internet-delivered television service may need to pay between \$143 and \$416 per month for adequate unmanaged transmission capability – depending on whether these customers are using applications that involve Internet throughput roughly equivalent to viewing only “standard definition” video – or are consuming Internet usage equivalent to watching modern “high definition” video.⁷ These cost figures seem so high that it appears unlikely that a product offering of unmanaged IP video access service could entice enough customers to be a commercial success – and demonstrate some possibly unintended consequences of regulatory limitations on permitted network management.

The following sections first provide some background on the evolving nature of demand for Internet access services and the importance of increasing capacity in both “last mile” broadband access networks and in the “middle mile” and core backbone links of the Internet to accommodate these increasing demands. The paper then develops a simplified quantitative model of unmanaged broadband network costs and demonstrates the costs that would be incurred to provide advanced video services or their bandwidth equivalent under this configuration. Concluding remarks follow.

2 Backdrop to the debate

Internet usage patterns are evolving rapidly. When the Internet and its predecessor the ARPANET were first developed in the 1970s, its intended uses were principally for file transfers between mainframe computers and remote login to time-sharing systems. But its key innovation turned out to be electronic mail.⁸ With the development of the World Wide Web in the early 1990s, the preponderance of Internet traffic moved to web browsing: first of relatively static monochromatic text pages and later of more colorful pages combining text with graphics images. While dial-up access at speeds of up to 56 Kbps accommodated effectively the demands of electronic mail and basic web browsing, as web pages became

or within any particular quality-of-service parameters. While such engineering may provide adequate service capability to users when network capacities exceed greatly typical user demand or demanded applications are not particularly quality-sensitive, as demand rises huge new investments and expansions in network capacity may be required in order to continue to serve this demand successfully on a best effort basis. See, McClure (2007). See also, Working Group on Net Neutrality (2007, pp.15-32) for a discussion of the challenges that Japan’s rapid demand growth has presented for network capacity expansion and management.

⁶ The network management at question here is to improve capacity utilization. Such management differs substantially from the congestion-relief network management that may have been employed by Comcast to address peer-to-peer (“P2P”) traffic loads from BitTorrent. P2P file transfer occurs when the computers exchanging files are end user peers. This differs from traditional file transfer that takes place between a web server computer and an end user computer.

⁷ These cost levels assume today’s technology. To the extent that customers may only gradually migrate to Internet-delivered video services, it is possible that technology may evolve so that customer demand may be served more cheaply some years hence. See, Odlyzko (2009) for implications of this alternative view of demand and cost evolution. But see also notes 19 and 49, below.

⁸ See, Leiner *et al* (2003) stating that “... while file transfer and remote login (Telnet) were very important applications, electronic mail has probably had the most significant impact of the innovations from that era.”

more elaborate and animated, and as file transfer shifted from the exchange of multi-kilobyte text and data files to multi-megabyte music files, customer satisfaction with dial-up began to wane.⁹ When web pages comprised only several tens of kilobytes of data, dial-up access generally provided adequate speed to paint the computer screen within a few seconds. But as web pages have become more complex and require hundreds of kilobytes to render and demanded music files hover in the megabytes, waiting tens of seconds to tens of minutes for a download over a dial-up connection has become less satisfactory.¹⁰ The result has been an increasing consumer shift from dial-up connections to broadband connections that exhibit access speeds ranging from several hundred Kbps up to tens of Mbps.¹¹

Dial-up connections' data throughput is limited because these links exploit only the 3000 Hz of analog voice-frequency bandwidth that is processed by Class 5 circuit switches and carried by voice-grade interoffice channels. Thus, data transmission at speeds greater than 56 Kbps requires bandwidth beyond that available in voice-grade channels. Broadband connections employ several different technologies to provide this expanded bandwidth.

Digital subscriber line ("DSL") technologies meet this challenge by using spectrum available in copper pair customer loops that is higher in frequency than the voice spectrum. But since signals in this high frequency spectrum will not pass through voice switches, the data signals must be peeled off by a digital subscriber line access multiplexer ("DSLAM") before the copper loop meets the Class 5 switch. The DSLAM then aggregates these digital data signals with digital data signals from other subscriber lines. These aggregated signals are then sent on special high capacity data lines back into a regional collector network, and from there back into the Internet backbone.

Cable modem ("CM") systems utilize the abundant spectrum available in coaxial cable (up to 860 MHz) to provide data transmission service to customers. Data signals are brought from the cable system's headend into a neighborhood on optical fiber cables in circuits segregated from the system's video broadcast signals. Once these optical data signals reach a fiber node in the customer's neighborhood they are converted into electrical data signals and carried on a 6.4 MHz channel in the upper radio frequency ("RF") spectrum of the coax to provide a dynamically shared download channel to the various CM customers in that neighborhood. Upload capacity in CM systems is more constrained. Because of the transmission characteristics of the coax and the need to maintain analog

⁹ Note that while data transfer speeds are typically measured in *bits* per second, file sizes are measured in *bytes*. Because a byte is comprised of eight bits, this means that at maximum throughput a 56 Kbps dial-up connection can transfer only 7 kilobytes per second. In practice, because of coding overheads and error correction, actual byte-per-second throughput is at best a tenth of the bit-per-second transfer rate. Thus, the effective throughput of a perfectly functioning 56 Kbps connection is roughly 5 to 6 kilobytes per second.

¹⁰ See, Broadband Working Group (2005) for a discussion of evolving Internet use.

¹¹ Data download and upload throughput rates depend not only on the speed of the customer's access connection, but also on the capacities of the regional interoffice and backbone IP networks. This is because most information content demanded by customers is sourced from web servers or "peers" that are scattered around the country and world – and are not "local" to the customer. Horrigan (2007) reports that by March 2007, 71% of U.S. adults reported using the Internet and 47% of U.S. adults accessed the Internet through a home broadband connection – suggesting that over 66% of home Internet access is via broadband and rising rapidly.

VHF television channels in their native RF spectrum, only portions of the coax spectrum between 5 and 42 MHz are usable for upstream links.¹²

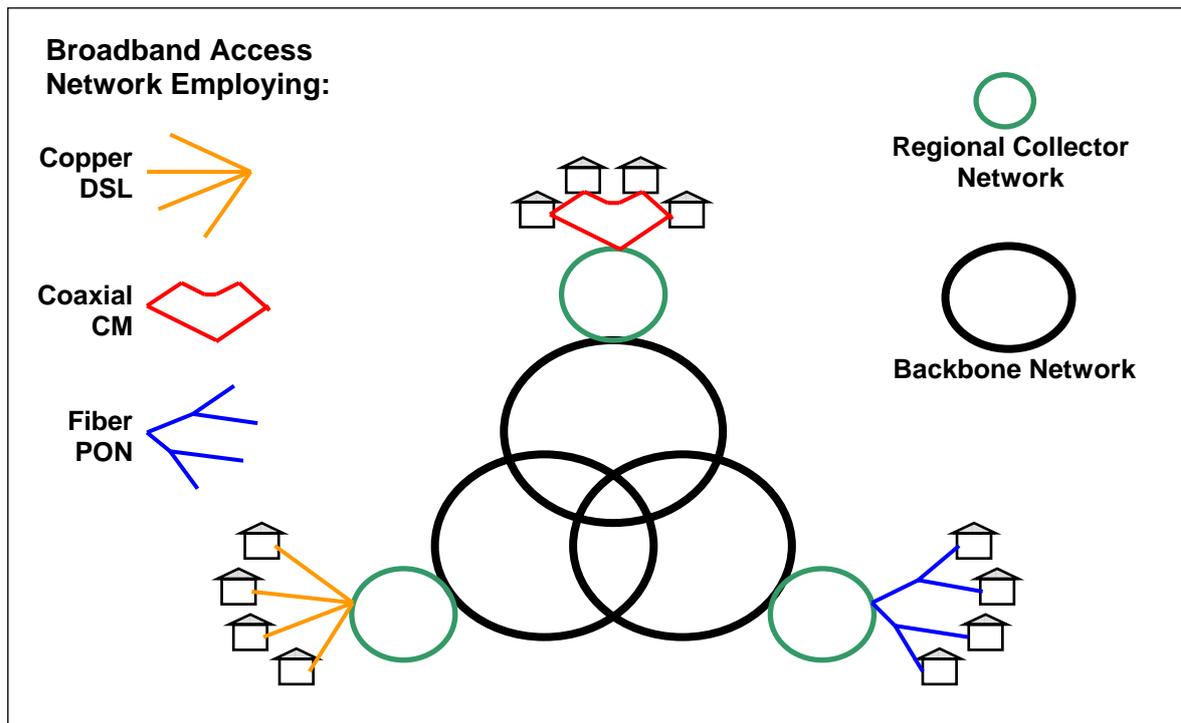


Figure 1: IP network schematic

Passive optical networks (“PONs”) face fewer limits. Although they distribute digital signals among customers in a neighborhood in the same general fashion as CM systems, their available bandwidth is more abundant in the last mile; and it is shared among customers in fixed guaranteed slices whose bandwidth does not depend on the amount of simultaneous use demanded by other customers in the neighborhood. If desired, PON upload links may have nearly the same capacity as download links. Once the customers’ IP packets from a neighborhood reach the optical line termination equipment at the central office however, these packets are aggregated with packets from PONs serving other neighborhoods and these aggregated signals are then sent on special high capacity data lines back into a regional collector network and from there into the Internet backbone.¹³

In terms of relative access line capacity, DSL links using ADSL technology generally have a maximum downstream throughput of 6 Mbps for customers that are relatively close to the DSLAM. Using VDSL technology, downstream speeds may increase to 25-30 Mbps – again for customers located close to the DSLAM. Customers located further from the

¹² See, International Engineering Consortium Web ProForum on “Hybrid/Fiber Coax (HFC) and Dense Wavelength Division Multiplexing (DWDM) Networks” (http://www.iec.org/online/tutorials/hfc_dwdm/).

¹³ Another optical fiber option is active Ethernet. This architecture has many proponents (see, B. Kanter, “The FTTP battlefield: Active Ethernet vs. PON,” *Lightwave Online* http://lw.pennnet.com/Articles/Article_Display.cfm?Section=OnlineArticles&SubSection=Display&PUBLICATION_ID=13&ARTICLE_ID=207692). However, the largest fiber deployments in the U.S. have been PONs.

DSLAM (greater than 4000 feet) generally will get less throughput.¹⁴ Typically, upstream bandwidth is 1.5 Mbps or less, but the division of total available DSL bandwidth between upstream and downstream use is fairly arbitrary.

Cable modem links employing current versions of the Data Over Cable Service Interface Specification (“DOCSIS”) generally provision downstream channels that can carry a maximum of 38 Mbps and upstream channels carrying between 10 and 30 Mbps – depending on the DOCSIS version.¹⁵ But because this bandwidth is shared across all cable modem customers served by the same fiber node, the actual bandwidth available to each customer will depend both on the number of other customers in the neighborhood demanding bandwidth at the same time and the maximum throughput limit per customer specified by the CM operator.¹⁶

The PON being deployed by Verizon and marketed under the FiOSSM brand generally provides downstream bandwidth of 622 Mbps divided among a maximum of 32 subscribers connected to the same fiber strand, or an average downstream capacity of 19.4 Mbps per subscriber.¹⁷ But because many customers subscribe only to Verizon’s 5 Mbps or 15 Mbps downspeed options, Verizon is able to offer individual customers the option of receiving up to 30 Mbps of downstream throughput. Upstream FiOS bandwidth of 155.5 Mbps is similarly shared by no more than 32 customers – yielding an available upstream throughput of 4.9 Mbps per customer. Other generally available PONs have capacities of 1244 or 2488 Mbps downstream and 622 or 1244 Mbps upstream.¹⁸

The reason maximum capacities of broadband access networks are key is that use of IP delivery to provide real-time video services or their bandwidth equivalents would swamp currently available access capacities. This is because even today’s most advanced video compression protocols (“MPEG-4/H.264”), standard definition television (“SDTV”) channels currently require roughly 2 Mbps of throughput – and high definition television (“HDTV”) channels currently require close to 10 Mbps of throughput.¹⁹ But at least as

¹⁴ This is discussed in P. Sevalia, “Swimming Upstream: The Case for Higher Speeds” (<http://www.convergedigest.com/whitepapers/documents/Ikanos-VDSL.pdf>).

¹⁵ DOCSIS standards are maintained by CableLabs (<http://www.cablemodem.com/>). Also see, Cisco Systems “Cable Access Technologies” (http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/cable.pdf) for a further description of DOCSIS specifications. The most advanced DOCSIS specification currently on the drawing board is DOCSIS 3.0. This specification employs channel bonding to provide up to 160 Mbps of shared download capacity.

¹⁶ Cable system fiber nodes rarely serve less than 125 to 500 customers.

¹⁷ Further information about FiOS is available at <http://en.wikipedia.org/wiki/FiOS>. Note that Verizon distributes its FiOS multichannel video signal over a separate wavelength in the fiber from that used to supply its high speed Internet access service.

¹⁸ The newest FiOS installations may be employing this gigabit (1244 or 2488 Mbps) PON equipment, but it is not yet common in Verizon’s network. Specifications even exist for PONs with up to 4976 or 9952 Mbps of downstream capacity and 2488 or 4976 Mbps of upstream capacity. Divided across 32 users, such PONs could offer average per customer capacities of up to 311 Mbps downstream and 155.5 Mbps upstream. Given the immense capacities of these PONs, it is also possible that the network operator would choose to split the PON across 64 customers rather than 32 customers. If that is the case, maximum throughput per customer would be 77.8 or 155.5 Mbps downstream and 38.9 or 77.8 Mbps upstream.

¹⁹ These compression estimates for MPEG-4 will become more firm as IP networks gain greater experience in handling this traffic. The MPEG-2 standard that MPEG-4 is expected to replace has required at least 4 Mbps for the transmission of SDTV and 19 Mbps for HDTV. See, Ruel (2004) and Motorola Solutions Paper, “MPEG-4 Network Migration: Solutions and Benefits” (<http://www.motorola.com/staticfiles/Business/Solutions/Industry%20Solutions/Service%20Providers/Telcos/Digital%20Headend%20Solutions/MPEG-4%20Encoding%20for%20Telcos/Documents/MPEG->

important as the level of these throughput rates is the fact that they must be virtually continuous – without more than a few milliseconds of jitter while a person is watching.²⁰ Thus, a connection that varies between 1 and 3 Mbps (depending, perhaps, on neighboring customers' use of the network) is wholly inadequate for streaming SDTV. The resulting video dropouts or pixelation would make the image unacceptable. The available throughput must be a continuous minimum of roughly 2 Mbps. Similarly, streaming HDTV cannot be accommodated by a broadband link that averages 10 Mbps in speed but achieves this average by vacillating between 5 and 15 Mbps of throughput. It must have a continuous minimum throughput of 10 Mbps. And of course, if there is more than one television in the house and its occupants wish to view more than one program simultaneously, required capacities scale up directly with the number of programs being viewed (for example, three SDTV channels require 6 Mbps; two HDTV channels require 20 Mbps).²¹

But adequate capacity in the customer access links is only one of several keys to the delivery of real-time video-type services to customers. Whether or not today's IP networks can accommodate tomorrow's demands for high bandwidth real-time services like SDTV or HDTV also depends on the amount of excess capacity currently available in regional collector and backbone networks that are used by every broadband access network to convey data packets to/from and through the Internet core. If excess capacity is abundant everywhere in these shared resources, incremental investment and cost may be small. If it is nonexistent or constrained, greater investments may be required.

Current use-intensity of broadband access connections is quite modest, less than 50 Kbps per customer during the network busy-hour.²² Most customers do not spend the entire evening (typically the IP network busy-hour) communicating at full bore with the Internet and saturating the full bandwidth of their access connection. While certainly some customers may engage in such hyper-intensive use, these constitute a small minority of current residential users – or they are business customers that demand (and pay extra for) IP network connections that provide dedicated guaranteed throughput.²³ Rather, most residential customers use the Internet intermittently and largely for low bandwidth applications such as email and web browsing. As a result, a customer's total daily bandwidth use is typically less than 1% of the potential maximum throughput of his access

[4 Solutions Paper 553893-001-a.pdf](#)). As compression technologies develop further, it is possible that SDTV throughput could require as little as 1 Mbps and HDTV throughput could require as little as 6-8 Mbps. But it is also possible that customer demands for increased picture quality could limit or reverse these bandwidth savings. Indeed, customers are beginning to demand HDTV in 1080p format – which requires up to 50% more data throughput than the 720p and 1080i HDTV formats that are currently in popular use. See, Penston (2007) and Luthra (2006).

²⁰ A signal experiences jitter when the packets containing it do not arrive at a constant rate.

²¹ Rapid channel-changing capability may also increase the subscriber line bandwidth requirements for broadband-delivered television.

²² See, BellSouth Chief Architect H. Kafka's presentation "Everything on the Net" <http://www.ofcnfoec.org/materials/2006KafkaPlenary.pdf>. Because network capacity is not fungible over time, capacities must be adequate to handle periods of peak load – known as the network busy-hour. See also, usage statistics stated by UK ISP PlusNet reported in Ofcom (2007, p.32).

²³ Broadband providers typically offer a separate portfolio of guaranteed end-to-end throughput access connections. These may be called T1 or DS1 lines (1.5 Mbps of continuous symmetric throughput); T45 or DS3 lines (45 Mbps of continuous symmetric throughput) – or even higher capacity optical lines. Because these connections are engineered to guarantee these data transmission rates, their prices are commensurately higher than corresponding speed residential broadband access connections.

connection – and even use during the network busy-hour may not exceed 3% of maximum access connection bandwidth.²⁴

Whether or not these usage patterns suggest that current access connections exhibit substantial excess capacity depends on whether all customers simultaneously are capable of using the maximum advertised bandwidth of their broadband access connections. As will be discussed in the following section, generally they are not.²⁵ Hence, significant capacity upgrades of access connections will be necessary for them to handle any but the most rudimentary video signals. But even if there is some equivocation about the need to expand existing broadband access capacity, there is no question but that regional IP collection and backbone links will need to be expanded to accommodate higher bandwidth demands. This is because today's shared IP collection and backbone networks are sized to offer only 50 Kbps of simultaneous capacity per residential customer.²⁶

The amount by which current broadband access and regional and backbone IP networks will need to be expanded to handle new demands for video will depend on how video and other high-bandwidth program material is sourced, and how its delivery over the Internet may be managed or controlled. If customers are to be capable of choosing any independent vendor (located anywhere on the Internet) to provide their video content, this distribution will likely follow the unicast format. Unicast means that each program watched by each customer is provided via a separate stream of packets through the Internet – all the way from video content provider to the viewing customer. As a result, Internet backbone and regional collection network links need to have the capacity to carry the simultaneous packet stream demands of all customers.²⁷

Therefore, for every customer that wishes to view an independently-streamed HDTV program during the network busy-hour (which is typically early evening – the same as prime viewing time), backbone and regional IP links need to be enlarged by 10 Mbps of capacity.²⁸ And because other customers seeking to view streaming video will most likely

²⁴ This is not to say that all current customers match this profile. A growing number of residential customers are heavy users of P2P services such as BitTorrent. Such services are both high bandwidth and intrinsically symmetric. They require as many packets to be uploaded by subscribers as are downloaded. Because current broadband architectures are largely asymmetric with more download capacity than upload capacity, proliferation of P2P services is likely to congest residential uplinks before it congests downlinks. Indeed, Comcast's attempts to relieve uplink congestion appear to be at the root of its recently highlighted network management practices targeting BitTorrent traffic. See, Ou (2008).

²⁵ This is because certain of the resources in these access networks are "oversubscribed". See, Mitchell (2000) for a discussion of oversubscription (also called contention or overbooking) in broadband access networks. See Ofcom (2007, p.21) for more information about how contention ratios affect the capabilities and pricing of broadband access lines.

²⁶ See Broadband Working Group (2005) for a discussion of how evolving usage patterns are exhausting current IP network capacity. This white paper also notes that even in Korea which is a world leader in broadband penetration and speeds, per-subscriber traffic barely exceeds 50 Kbps.

²⁷ This is because there is no assurance that independent video content providers will cache their content "close" to every customer. To the extent that content providers do create local caches, this may mitigate a portion of the load on the Internet backbone – but it is unlikely that independent providers will find it feasible to cache their content at more than a handful of distributed locations. To eliminate excess backbone load, each content provider would need to maintain thousands of caches. Norton (2007) examines the economics of video downloads under unicast transit and content distribution network caching. In general, he finds that the economics of caching are no better than those for unicast transit.

²⁸ P2P traffic generates similar profiles of "across the Internet" traffic loads. The nature of P2P file-distribution applications is that every subscriber downloading a file receives his own unique copy of the file's packets – sourced and assembled from various cooperating subscribers to the P2P application scattered

demand it during the busy-hour, an equal 10 Mbps of capacity expansion in IP backbone and regional links is needed for each and every customer seeking to view this video. Given that regional and backbone IP networks currently are engineered to handle only the roughly 50 Kbps of busy-hour throughput utilized by each broadband customer, it is clear that under unicast, regional and backbone IP links would need a capacity expansion that is close to two orders of magnitude larger than the capacity expansion required in broadband access networks.²⁹

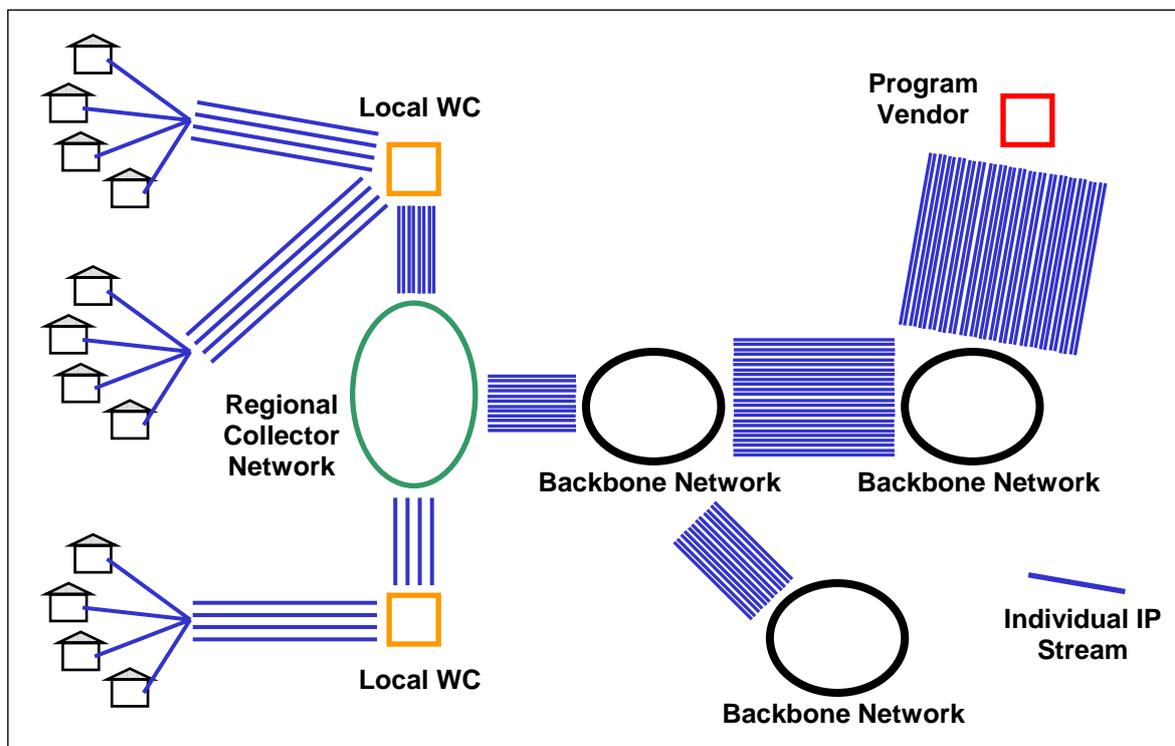


Figure 2: Unicast schematic

around the Internet. In addition, because P2P applications are generally not capable of favoring content sources that are “close” to the requesting client, P2P packets will typically need to traverse both the sender’s and receiver’s access, regional collection and backbone networks.

²⁹ It is possible to provide equivalent video services over regional and backbone networks with capacities much smaller than those discussed above if multicast rather than unicast distribution technology is employed. See, Cisco Systems “IP Multicast Technical Overview,” August 2007 (http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6552/prod_white_paper0900aecd804d5fe6.pdf). Under a multicast architecture, bandwidth across the Internet is preserved by sending only a single copy of the program stream across any particular network link. But a multicast network is one that requires active network management that may be prohibited (either technically or economically) under this unmanaged/neutral model. See, Cisco Systems “IPTV Demands a Real-Time Control System that Assures the Quality of Experience,” 2006 (http://www.cisco.com/en/US/prod/collateral/netmgtsw/ps5689/ps6478/prod_white_paper0900aecd804de983.pdf) and “IP Next-Generation Network Requirements for Scalable and Reliable Broadcast IPTV Services,” September 2006 (http://www.cisco.com/en/US/prod/collateral/routers/ps5763/prod_white_paper0900aecd8051a9a7.pdf) for a description of the suite of network management controls needed to provide high quality IPTV services over efficient multicast networks. Because these network managements are costly and would treat IPTV traffic differently from other traffic, it is likely that they would run afoul of network neutrality regulations such as proposed in the “Internet Freedom Preservation Act,” described in note 4, above. IPTV services offered by major broadband networks such as AT&T’s *U-verse* or PCCW’s *now TV* employ managed multicast technologies. See, Ortiz (2006).

The following section discusses in greater detail the engineering and costs of broadband access networks, regional collection networks and Internet backbones and the import of these costs for “network neutrality.”

3 Network expansion: managed and unmanaged solutions

As discussed above, there are two basic ways to construct an IP network capable of handling all of the diverse uses demanded by end user and information service provider customers: build an extremely “thick” unmanaged network or build a “thinner” managed network.

Advocates of the first option suggest that if broadband IP networks were constructed “thick” enough to furnish each end user customer with extremely large amounts of bandwidth (for example, up to 100 Mbps), these customers could run whatever applications they desire (for example, browsing, email, VoIP, streaming video, P2P file transfer, etc.), sourced from whatever independent information service provider they desire, at adequately high levels of service quality without any significant network management intervention by IP network operators.³⁰ Furthermore, they suggest that thick unmanaged networks will cost less than overlaying management on thinner networks.³¹

Advocates of the second option suggest that building an extremely thick unmanaged network is unwise for two reasons. The first is that it would be uneconomic – most end users would be unwilling to pay its full cost. The second is that the inevitable spread of bandwidth use-indifferent applications (such as “push” applications, certain P2P, uncompressed video, etc.) will likely cause even the thickest unmanaged network to fail to supply the consistently high levels of service quality needed to enable many of the applications that customers desire.³² Instead, these advocates argue that a managed network holds better promise to provide all of the applications desired by customers at a much smaller cost than an unmanaged network.³³

The following analysis focuses on just one aspect of this debate – that of the costs associated with building a very thick unmanaged IP network capable of providing customers the diverse and bandwidth-intensive applications they desire. To begin, we provide some background on the general cost structure of IP networks and customers’ usage and service quality requirements.

³⁰ Acronym VoIP is “Voice over Internet Protocol” service.

³¹ See, Bachula (2006). Note, however, that the potential relevance of Internet2’s cost and operational experience for commercial IP networks has been downplayed by Internet2’s President and CEO, Douglas Van Houweling who has said, “Commercial networks don’t want a lot of excess capacity. ... They try to run at 75 percent. But we want to run about 25 to 30 percent full so there’s always headroom to try something new. We don’t think of it as excess capacity. We think of it as room to innovate.” See, Van (2006).

³² See, Cisco Systems “Cisco Service Control: A Guide to Sustained Broadband Profitability,” November 2005 (<http://www.democraticmedia.org/files/pdf/CiscoBroadbandProfit.pdf>).

³³ Researchers have attempted to quantify the extra network capacity that may be required if it is impermissible to manage IP networks in a way that provides differentiated classes of service to different packets. One such study finds that in order to guarantee a given level of service quality (for example, latency and packet loss) to premium class traffic as achieved in a differentiated network, a network without class-of-service differentiation generally will require at least 60% more capacity than a differentiated network – and possibly as much as two or more times the amount of capacity required by a differentiated network – depending on the utilization level of bottleneck links in the network or whether the network is operating under a transient failure situation, etc. See, Yuksel *et al* (2007).

As noted earlier, all IP networks are largely shared resources. That is, only fractions of their infrastructures are wholly dedicated to particular individual customers' use. The rest of the infrastructure is shared among multiple customers' contending uses. But basic traffic engineering principles demonstrate that the more widely shared a resource, the more feasible it is to load it to capacity.³⁴ Thus, given throughput levels may be maintained over core Internet backbone links with the least amount of excess capacity. Regional IP collector links are likely to require a bit more excess capacity. And the shared portions of broadband access lines are likely to need the greatest amount of unused capacity to maintain adequate throughput. These two principles, coupled with the fact that most residential customers of broadband service use only a tiny portion of the potential capacity of their access connection and transfer a relatively small number of bytes to and from the Internet over the course of a month, drive the costs associated with upgrading today's IP networks to be video-capable. These principles have similar import for the capabilities and costs of each of the three major wireline broadband access architectures.

When a customer purchases a DSL connection with a speed advertised as 1.5 Mbps, this means only that the customer has a dedicated 1.5 Mbps of bandwidth between his home and the DSLAM serving his DSL line. In heavily populated areas, this DSLAM is likely located at the central office serving the customer. In less densely populated areas, the DSLAM may be located at a remote terminal in the subscriber's neighborhood. Once the customer's digital signals reach the DSLAM, they are mixed with the data signal packets from all other customers served by that DSLAM (typically several hundred) and transmitted back into a high speed IP packet collection network for routing and delivery to the Internet backbone. In particular, neither this shared collection network nor the shared Internet backbone are sized to have a dedicated 1.5 Mbps of bandwidth available for each DSL customer they serve. Rather, because typical customer usage is much, much less than the full capacity of their DSL access connections, signals from the individual DSL lines may be statistically-multiplexed (combined) into a signal that requires much less transmission bandwidth than the simple sum of the many DSL lines' capacity that they serve (for example, if a DSLAM serves 600 1.5 Mbps DSL lines – or 900 Mbps of total digital loop capacity, it likely needs less than 45 Mbps of capacity back to the Internet to adequately serve this total demand). And because this is the case, competitive economics demands that these shared resources be provisioned no more lavishly than this efficient capacity.

In a cable broadband network, resource sharing begins even sooner – at the customer's cable modem. This modem shares a given amount of bandwidth (for example, 38 Mbps in basic DOCSIS networks) with typically several hundred other customers in the neighborhood. While customers may have their personal maximum connection speed limited (say, to 3 Mbps), because there are many more than 12 customers sharing this bandwidth, the actual bandwidth available to the customer at any point in time may be much less than the advertised maximum connection speed. A customer's actual effective throughput depends on how many other customers in the neighborhood are trying to use their connections at the same time. And once cable modem signals leave this shared neighborhood network, they also are sent through a collection network back to the Internet backbone that typically has less than 38 Mbps of bandwidth available for each neighborhood CM network that it serves.

³⁴ See, Members of the Technical Staff (1983, pp.147-191).

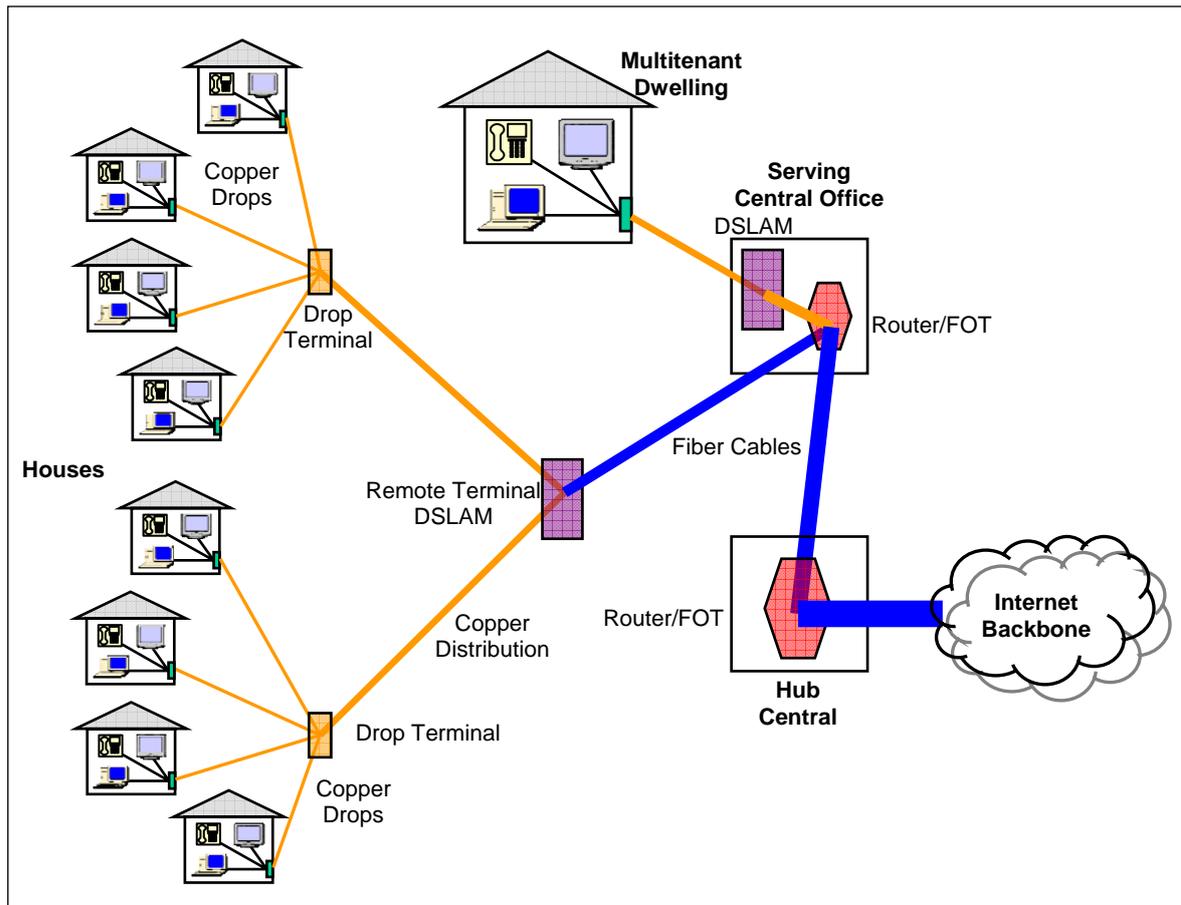


Figure 3: DSL schematic

PONs are similar in architecture to the neighborhood distribution segment of CM networks, but use optical fiber rather than copper coax and offer much higher capacities (for example, 622 Mbps and higher). Because the number of customers sharing neighborhood bandwidth is much more rigidly limited and controlled (for example, 32 or 64 customers per PON based on the type of optical signal splitters deployed), the continuous bandwidth available to each customer may be guaranteed. But once these customers' IP signals reach the PON's optical line termination equipment at the central office, they too must be sent over shared collection networks to the Internet backbone – and economics and typical customer usage dictate that these shared networks be sized at much less than full multiples of the capacities of the individual PONs served.

The degree to which today's broadband access, collector and backbone networks need to be expanded to accommodate tomorrow's demands for high bandwidth real-time services depends on existing amounts of excess capacity and the service delivery technology.

As indicated earlier, access connections may currently offer substantial excess capacity. This is because the low-bandwidth and bursty nature of today's predominant Internet applications results in typical customer use levels that are much less than the maximum advertised bandwidth of most broadband access connections. Thus, if all access connections are, say, 1.5 Mbps, shared network resources need only be sized to no more than 3% of total access bandwidth (for example, 45 Kbps per customer) in order to provide adequate capacity and quality to serve today's customers. This provides great economy

over an alternative shared network sized to handle continuous simultaneous use of the full capacity of every customer's access connection.

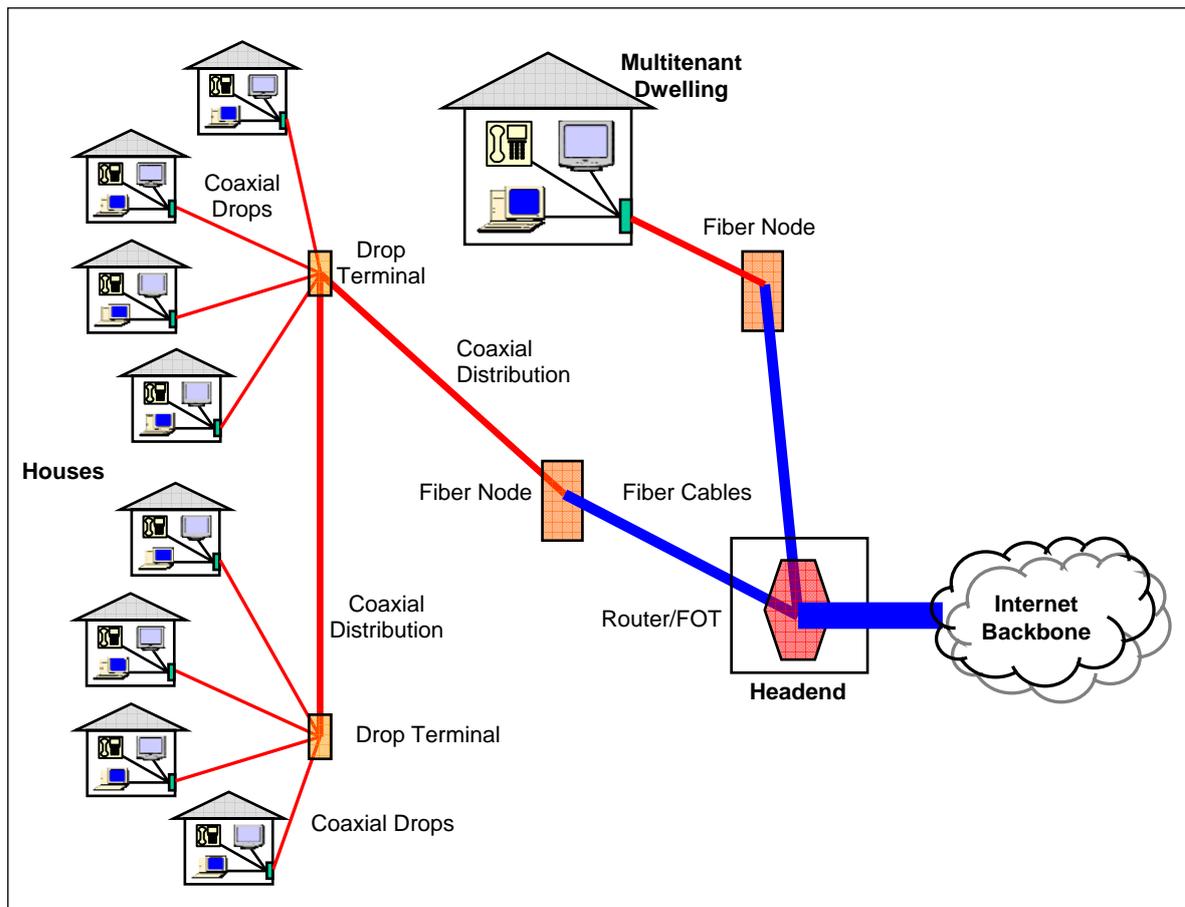


Figure 4: CM schematic

Because today's shared IP collection and backbone networks are sized to offer only 50 Kbps of capacity per residential customer, their per-customer costs are modest. But if these shared resources were re-sized to permit all residential customers to use the full capacities of their current access connections during the busy hour, shared resource capacities would need to increase at least 30-fold from today's levels, and shared resource costs would rise by a similar order of magnitude. This is because these shared resource collection and backbone networks are already extremely high capacity networks that carry many services (for example, traditional switched voice and dedicated data) in addition to IP traffic. Thus, expansion of their capacities is unlikely to yield greatly lowered costs per unit of capacity.³⁵

³⁵ Another way of stating this is that most currently-available scale economies in the construction and operation of high speed IP collection and backbone networks may already have been exploited. Thus, further expansions are likely to require substantial replication of facilities rather than the simple substitution of higher-capacity, lower unit-cost facilities. See Norton (2007) for a discussion of these issues.

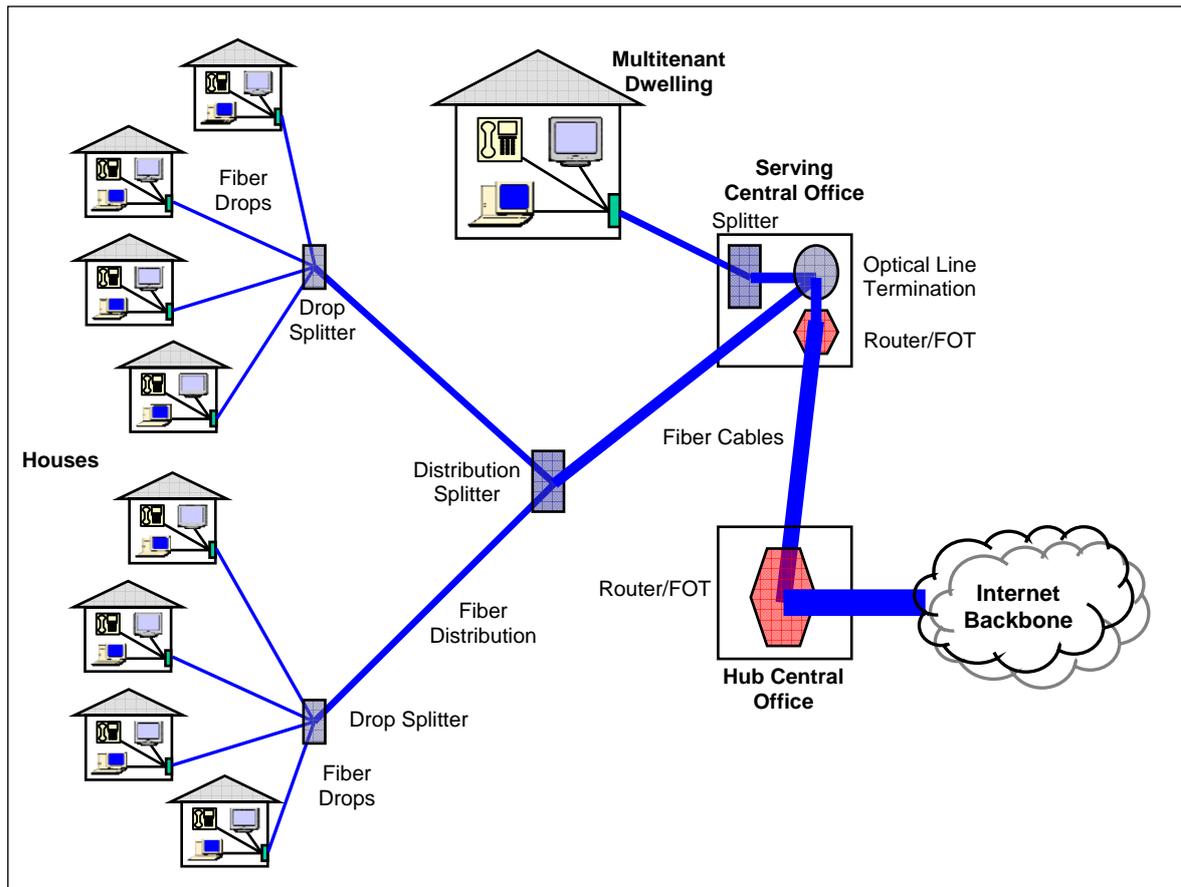


Figure 5: PON schematic

But expanding collector and backbone networks to accommodate increased customer busy-hour use of their existing broadband access lines (with capacities of only 1.5 or 3 Mbps) is almost certain to be inadequate to serve tomorrow's customer demands. High quality VoIP services require up to 100 Kbps for the duration of the conversation. And as stated earlier, each SDTV channel provided through IP streaming requires roughly 2 Mbps of continuous throughput and HDTV channels require roughly 10 Mbps of continuous throughput. Furthermore, households generally contain several televisions and viewers – who may wish to watch different programs during the busy-hour. Other members of the household may at the same time be wishing to share multi-megabyte music files or multi-gigabyte video files. While file-sharing applications and web browsing may be tolerant of less-than-perfect service quality or transmission delays, real-time applications like VoIP or streaming video are completely intolerant. Conversations that drop out or video scenes that pixelate or pause likely render such applications unacceptable to customers. Nearly all current broadband access connections are incapable of handling the 20+ Mbps of load that will be generated by these burgeoning demands. Thus, to accommodate the quality requirements for such services, not only will regional collector and backbone networks require massive expansions, but access networks will also require upgrades.

Meeting this challenge without imposing network management requires engineering broadband access, collector networks and backbones to be so massively “thick” that IP traffic delays or bottlenecks never arise. The next section of this analysis develops the costs associated with unmanaged network capacity expansion. Note that of necessity, the

modeling used to establish these costs is somewhat rudimentary. Because this type of extremely high capacity network is not in long-standing or widespread operation, data surrounding its costs are sparse. But in any event, the accuracy of these calculations is sufficient to establish that the estimated cost of the thick unmanaged network exceeds today's costs by a substantial sum.

4 Costs of unmanaged network expansion

For purposes of this analysis, the PON network architecture is used to determine the cost of the unmanaged network alternative.³⁶ PONs link customers with their serving central office by optical fibers – without any active electronics located between the residence and central office. Although this could be achieved by running a separate fiber strand between each residence and the central office, this is expensive and fails to take advantage of the virtually unlimited transmission capacity within each fiber strand. Thus to economize on fiber placement costs, PONs generally serve their neighborhoods through a single fiber strand terminating at the central office.³⁷ This fiber feeder strand carries an extremely high bandwidth signal (for example, 622 Mbps, 1244 Mbps, 2488 Mbps, etc.) out to a fiber distribution terminal (“FDT”) located in the neighborhood. At the FDT, a passive optical splitter divides the high bandwidth signal in the fiber into several replicas with each replica being channeled onto an individual fiber distribution strand leaving the FDT and heading to the houses in the neighborhood. Once a fiber distribution strand reaches a small cluster of houses, it enters a drop terminal (“DT”) containing another passive optical splitter. This splitter again divides the extremely high bandwidth signal in the distribution strand into replicas to send to each house in the cluster on a drop fiber. In this fashion, a single fiber strand terminating at the central office can serve 32 or 64 houses.³⁸

At each house, the drop fiber terminates on an optical network terminal (“ONT”), which is located in a small weatherproof box placed on the side of the house. The ONT then collects the extremely high bandwidth signal from the drop fiber and pulls out the portion of this data signal (for example, the $1/32^{\text{nd}}$ or the $1/64^{\text{th}}$ of it) that is assigned to this specific residence.

The costs associated with a PON generally divide into several pieces. First, there are the costs specific to each house. These include the cost of the ONT, and the costs of the drop fiber and terminal. Drop cables may be placed as aerial (that is, up on poles), as buried (that is, placed in a trench directly in the soil), or as underground (that is, placed in a conduit that is placed in a trench in the soil). Whether the drop is aerial, buried or underground generally depends on how the distribution fiber cable reaches the house

³⁶ See, Credit Suisse First Boston Equity Research “VZ Launches FiOS TV: Who’s Most Exposed?” September 22, 2005 and Verizon’s FiOS Briefing Session investor conference held on September 27, 2006 (<http://investor.verizon.com/news/20060927/>) for a description of the architecture and costs associated with Verizon’s FiOS PON. See also, Banerjee and Sirbu (2005) for a more general discussion of PON architectures and costs.

³⁷ The equipment on which this fiber strand terminated at the central office is called an OLT for “optical line termination”.

³⁸ For example, if the fiber feeder signal is split among eight distribution strands at the FDT; and if the distribution strand is split among four houses at the DT; the total number of splits in the PON is 32 ($= 8 \times 4$). In rural areas, PON splits may be fewer than in suburban and urban areas to lessen dissipation of the strength of the optical signal and permit it to travel greater distances.

cluster. If the distribution cable is aerial, then generally so is the drop. If it is buried, then generally so is the drop. And if the distribution fiber is underground, then generally, so is the drop.

The next category of costs is those of plant linking DTs with the central office. This consists of fiber distribution cables going to a neighborhood FDT, and then fiber feeder cables from the FDT to the central office. This plant may be aerial, buried or underground, and costs vary accordingly. In particular, aerial plant is generally the least expensive, buried plant the next most expensive and underground plant the most expensive. In addition, plant costs are extremely sensitive to the surrounding topography. If the location is rural, per-foot costs may be relatively low. This is because pole spans may be long, trenching is easy and rights-of-way are plentiful and inexpensive. In suburban areas, costs increase. Pole spans are shorter, trenching is more frequently interrupted by driveways and other pavement; and rights-of-way are more constrained and expensive. Urban plant is the most expensive. Pole spans are short, trenching is made extremely difficult by pervasive pavement; and rights-of-way are tight and their use may be extremely expensive. The prevalence of aerial versus buried versus underground plant may also vary by location. In general, rural and suburban plant is more likely to be aerial or buried, while urban plant is more likely to be underground.³⁹

Additional PON costs are incurred once feeder cables reach the serving central office or wire center. Here, each neighborhood's fiber strands are terminated on an OLT. Aggregated signals are then passed from the OLT to a router and then to a fiber optic transmission terminal ("FOT").⁴⁰ The FOT at the serving wire center then sends these signals back into a fiber optic regional collection network to a FOT located at a hub wire center. Once at this hub wire center, these signals are passed through another router to aggregate them with signals from other wire centers that subtend this hub. They are then ready to be transferred over to an Internet backbone. This is commonly done via a high speed dedicated access line leased from a local or interexchange telecommunications company.

The cost of the high speed dedicated access line from the hub wire center to an Internet backbone depends on the capacity, location and length of this line. In general, higher capacity and longer length increase its cost, as does rural location. Further costs associated with the provision of broadband Internet service to residence customers are for the "transit" service that the customer's Internet service provider ("ISP") must purchase from the Internet backbone operator. Transit service is the to-anywhere/from-anywhere delivery and receipt of the customer's data packets over the Internet. This service is generally priced based on the maximum rate at which the ISP sends or receives data packets to/from the Internet backbone during the network busy-hour.⁴¹

In addition to the above investment and maintenance costs of the access PON and the costs related to connecting to the backbone, there are costs of operating the broadband access network. These operating expenses include network operations costs (for example,

³⁹ Facilities located at the customer premises and connecting the premises to the central office will be called access outside plant in our cost reporting.

⁴⁰ A router is a type of packet switch that examines the header of the packet that contains its destination information. Based on this information and routing algorithms within the router, a router "routes" the packet to the next network node on its journey to its ultimate destination.

⁴¹ For convenience, the cost of the dedicated access line from the hub wire center to the Internet backbone plus the transit cost over the backbone will be denoted as "backbone" costs in the quantitative analysis.

power or surveillance), customer operations costs (for example, service creation, ordering or billing) and corporate overheads (for example, senior management and support services).

A simple quantitative model of the cost of an unmanaged PON-based IP network is presented in the Appendix. Using reasonable (but user-adjustable) input values for the costs of the different elements of the network, monthly per-customer costs may be derived. These costs, of course, depend critically upon the throughput capacity and assumed traffic quantities handled by the network.⁴²

Of most significance is how costs vary based on different assumptions as to the services used by subscribers. The following table presents traffic usage levels under several alternative use profiles. The first is today's typical user profile – data use only.⁴³ The second profile represents the use characteristics of a current “power user” – again, data use only. The third profile assumes modest future demand for video services (that is, two simultaneous SDTV channels – or their bandwidth equivalents) plus expanded data use that is three times higher than current “power” use. The fifth profile assumes video demand levels that are likely to be typical of tomorrow's viewers (that is, one HDTV channel viewed simultaneously with three SDTV channels – or two simultaneously viewed HDTV channels – or their bandwidth equivalents) plus the expanded data use already mentioned. Finally, the fourth profile assumes that future users split 50/50 into the categories of modest video use and typical video use.⁴⁴

⁴² Presentation of this quantitative analysis is intended to advance usefully the debate over the costs of advanced broadband networks beyond that of duelling anecdotes.

⁴³ A current “power user” is assumed to use roughly ten times as much bandwidth as an average user. This is consistent with Korea Telecom reports that the top 5% of their users consume almost half of their network's bandwidth. See, Kim (2005).

⁴⁴ The purpose of these customer categorizations is just to provide some intuition into what type of services customers in each class might be consuming. See, Motorola White Paper, “Anticipating the Bandwidth Bottleneck,” February 2007

(http://www.motorola.com/staticfiles/Business/Solutions/Industry%20Solutions/Service%20Providers/Telcos/Documents/static%20files/Anticipating_the_Bandwidth_Bottleneck_Copy.pdf). For the purposes of unicast network cost development, the key consideration is the total bandwidth demand emanating from an average customer in the usage class. It matters less exactly what set of applications generates the demand. Note that although half of the customers in this scenario may consume only a modest amount of video services, the identity of these customers will not be evident to the broadband carrier when its plant is installed. Thus, these customers will likely be provisioned with access outside plant, ONTs and OLTs that are also capable of serving typical video use customers.

	Usage profile	Busy-Hour Download Traffic Rates		
		Data	Video	Total ⁴⁵
1	Current typical user	45 Kbps	--	45 Kbps
2	Current power user	450 Kbps	--	450 Kbps
3	Future modest video user (2 SDTV channels)	1.5 Mbps	4.0 Mbps	5.5 Mbps
4	50/50 split between future modest video use and typical video use	1.5 Mbps	12.0 Mbps	13.5 Mbps
5	Future typical video user (1 HDTV + 3 SDTV channels or 2 HDTV channels)	1.5 Mbps	20.0 Mbps	21.5 Mbps

Table 1: Bandwidth usage breakdown across user profiles

Based on these usage profiles, it is possible to determine the costs of satisfying each of these usage demands over a “thick” unmanaged PON deployed using currently available and priced technology.^{46 47}

BH Capacity Utilization	Outside Plant	WC+Cluster Cost	Operating Cost	Backbone Cost	Total Cost
Current typical usage	\$32.93	\$3.52	\$12.00	\$1.30	\$49.75
Current power usage	\$32.93	\$3.98	\$12.40	\$8.84	\$58.15
Modest future video usage	\$32.93	\$10.57	\$13.60	\$86.14	\$143.24
50/50 split modest/typical	\$33.91	\$19.12	\$15.60	\$211.14	\$279.78
Typical future video usage	\$33.91	\$29.14	\$16.60	\$336.15	\$415.80

Table 2: Cost to service alternative user profiles – unmanaged network

Note: “Outside Plant” includes equipment located between the customer’s house and serving wire center.
“WC+Cluster” includes equipment between the serving wire center and cluster hub wire center.

⁴⁵ While the monthly bandwidth usages corresponding to these busy-hour traffic rates depend on the ratio of busy-hour to average throughput, if we assume this ratio is 3:1, 45 Kbps of busy-hour use corresponds to monthly bandwidth use of about 5 GB, 450 Kbps corresponds to 49 GB/month, 5.5 Mbps corresponds to 602 GB/month, 13.5 Mbps corresponds to 1477 GB/month and 21.5 Mbps corresponds to 2353 GB/month.

⁴⁶ The cost figures presented here differ slightly from those presented in an earlier May 2006 version of this paper. These updated figures reflect an improved understanding of PON network structures as well as fuller and more recent information concerning Verizon’s PON implementation costs.

⁴⁷ Note that the costs presented here assume that this unmanaged PON is built completely from scratch. While it is possible that some portions of pre-existing copper access network facilities may be reused to support a PON infrastructure, the fraction of these facilities available for reuse may be rather small. In any event, such reuse would reduce principally the access network portion of PON costs, and would do little to reduce the most significant drivers of unmanaged video costs: the regional collector and backbone networks.

“Backbone cost” includes dedicated transport from the cluster hub to the Internet backbone plus IP transit over the backbone.

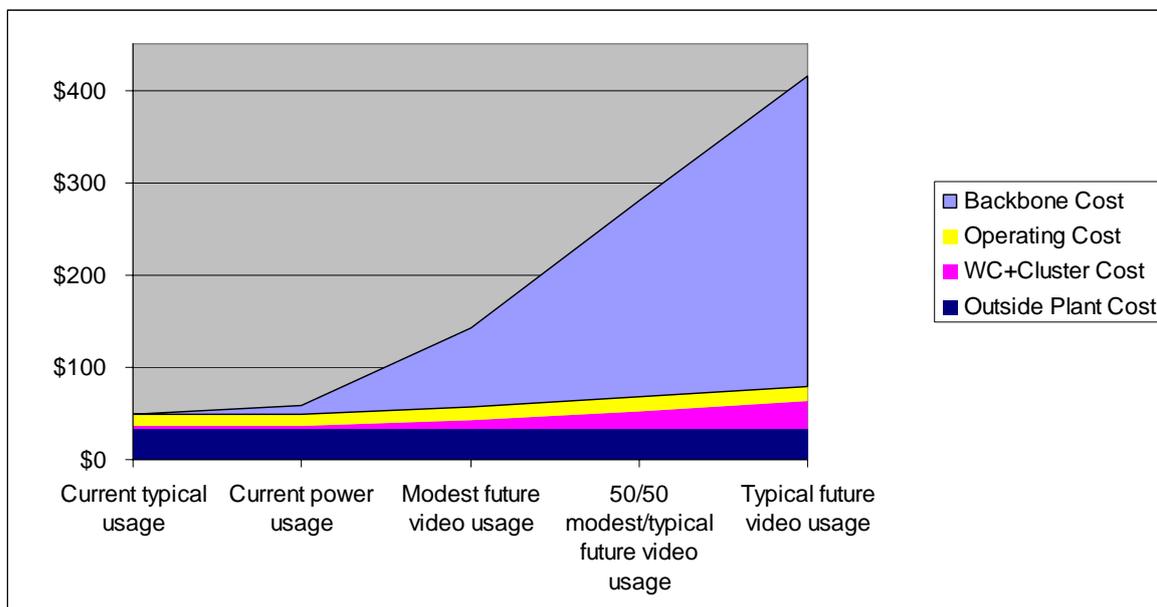


Figure 6: Monthly cost of broadband Internet by user profile assuming unmanaged network is sized to provide quality service

These costs are quite consistent with other estimates of the current cost of providing unicast streaming video services over broadband.⁴⁸ Thus, these studies suggest that the cost of sustaining IP access, regional collection and backbone networks robust enough to satisfy emerging demands for high bandwidth services such as video on an unmanaged basis are extremely high.⁴⁹ Indeed, the major cost increments are not in the neighborhood access portion of the network, but in the wire center cluster and backbone portions of the network. This is because current IP interoffice facilities and backbone cores are sized only to provide the 45 to 50 Kbps that each subscriber currently uses during the network busy hour. But since expected video usage demands the provision of at least 20 Mbps of throughput for each customer – not only in the access network, but through the IP collector

⁴⁸ BellSouth has noted that the cost of providing Internet backbone services sufficient for today’s typical users is roughly \$1.00 per month. For a power user this figure rises to \$4.50 per month. But for future SDTV users the figure mounts to \$112 – and as HDTV is included the figure rises still further to \$560 per month. See, <http://www.ofcfoec.org/materials/2006KafkaPlenary.pdf> or Wilson (2006). Similar expense figures have also been developed in Penston (2007).

⁴⁹ The figures calculated here and in other presentations are suggestive only. Large scale networks providing streaming high-quality video-type services to millions of customers have not yet been deployed. As technology develops, it may be possible to find some economies over today’s best technologies – or it may be that tomorrow’s technologies are not capable of scaling to these large capacities without larger cost increases than contemplated today. In particular, past improvements in communications throughputs have developed within an environment where no source of technical advance is eschewed. But if net neutrality regulations limit further throughput advances to ones deriving from raw data speed, and not, say, due to improvements in data compression or routing algorithms, there is no assurance that Moore’s Law-type advances in communications throughputs should continue.

and backbones as well, the largest percentage and dollar increases come from this required expansion of the collector and backbone networks.⁵⁰

5 Conclusions

At today's typical or power usage levels, customers generally pay between \$15 and \$45 per month for their Internet access service.⁵¹ But if usage intensity continues to rise and IP networks are permitted only to expand on an unmanaged basis, even a modest IP video viewer may need to pay \$143 per month for their Internet access service. A more typical user, one who expects to be able to receive two HDTV channels simultaneously in his home, may have to pay roughly \$416 per month. But even these prices purchase only the transmission services necessary to receive the given amount of video service – and not the actual video programming. Thus, in addition to these network costs, the customer for unmanaged streaming video will still need to make additional payments to independent content providers for the actual video programming they purchase.⁵²

While it is possible that some customers so value the possible extra freedom and diversity they may enjoy from obtaining services over an unmanaged network that they may choose to pay these lofty prices, these are daunting figures for most customers. Fewer than 5% of all households are willing to pay as much as \$150 per month for a “triple play” bundle of local telephone, long distance telephone and video services that includes programming costs.⁵³ Thus, it seems unlikely that unmanaged PONs with capacity adequate to stream unicast video services will gain commercial traction at today's cost levels.⁵⁴ A network operator that builds such a “thick” broadband network would likely have to charge so much for it that only a very few customers would buy it – and their payments would be insufficient to cover its cost because high capacity wireline networks typically are only financially successful if they achieve high penetration rates.

Customer willingness-to-pay and engineering costs invoke stern constraints. Simple quantitative modeling of the cost of unmanaged IP networks capable of satisfying tomorrow's video or other high bandwidth demands suggests cost levels that exceed today's levels by an order of magnitude. Thus, limits on permitted network management

⁵⁰ Note that because most video traffic may be inbound to the subscriber's broadband network, it is possible that certain of the subscriber's calculated backbone costs will be paid by the distant video provider. But this will simply result in the video provider passing these backbone costs through into its video service fees, they will ultimately be borne by the subscriber.

⁵¹ See, Goldman Sachs “2006 Outlook – Stuck in Neutral,” January 13, 2006, Exhibit 19.

⁵² Programming costs paid currently by multi-channel video programming distributors (“MVPDs”) are estimated to be between \$16 and \$20 per month. See, Federal Communications Commission (2006, pp.43-44) and Citigroup Equity Research “Telco Video Entry: Good for Some,” April 21, 2006. It seems unlikely that individual customers or small groups of customers will be able to negotiate as effectively for low programming prices as today's large MVPDs. Indeed, Citigroup estimates that today, a zero-size MVPD pays 30% more for its video programming than a large MVPD.

⁵³ See, Rappoport and Taylor (2006).

⁵⁴ The IP network expenses developed in this analysis are “best estimates” from the cost model. Because of the simplified nature of this modeling, it is reasonable to expect that more advanced modeling may show these cost estimates not to be perfectly accurate. But in order to reverse their qualitative implication that unmanaged IP networks are not economically sensible, these current modeling results would have to be overstated by a factor of between five and ten. Estimation errors of this magnitude seem unlikely.

techniques may have unintended implications for Internet-delivered video or large file transfer services – and fully unmanaged/neutral IP networks may not be a feasible goal.

6 References

Bachula, G. (2006) Testimony before the United States Senate Committee on Commerce, Science and Transportation, at its hearing on “Net Neutrality,” February 7. <http://commerce.senate.gov/pdf/bachula-020706.pdf>

Banerjee, A. and M. Sirbu (2005) “Towards a Technologically and Competitively Neutral Fiber-to-the-Home (FTTH) Infrastructure,” in I. Chlamtac *et al* (ed.) *Broadband Service*, John Wiley: New Jersey.

Broadband Working Group (2005) “The Broadband Incentive Problem,” MIT Communications Futures Program, September. <http://cfp.mit.edu/docs/incentive-wp-sept2005.pdf>

Chester, J. (2006) “The End of the Internet?” *The Nation*, February 1.

Dvorak, J. (2006) “Hijacking the Internet,” *PCMAG.COM*, March 15. <http://www.pcmag.com/article2/0,1895,1937952,00.asp>

Federal Communications Commission (2006) “Twelfth Annual Report on Competition in the Market for Delivery of Video Programming” *In the Matter of Annual Assessment of the Status of Competition in the Market for the Delivery of Video Programming*, MB Docket No. 05-255. March 3.

Garfinkel, S. (2003) “The End of End-to-End?” *Technology Review*, MIT, July.

Horrigan, J. (2007) “Home Broadband Adoption 2007,” Pew Internet & American Life Project, June. http://www.pewinternet.org/pdfs/PIP_Broadband%202007.pdf

Kim, T. (2005) “KT Seeks Usage-Based Internet Pricing,” *The Korea Times*, March 29.

Leiner, B., V. Cerf, D. Clark, R. Kahn, L. Kleinrock, D. Lynch, J. Postel, L. Roberts and S. Wolff (2003) “A Brief History of the Internet,” Internet Society, December 10. <http://www.isoc.org/internet/history/brief.shtml>

Lessig, L. (2002) Testimony before the United States Senate Committee on Commerce, Science, and Transportation, at its hearing on “The Government’s Role in Promoting the Future of Telecommunications Industry and Broadband Deployment,” October 1. <http://commerce.senate.gov/pdf/lessig.pdf>

Luthra, A. (2006) “Digital Video Compression – Past, Present and Future,” IEEE International Symposium on Multimedia. <http://ism2006.eecs.uci.edu/KeynoteWedMor.pdf>

McClure, D. (2007) “The Exabyte Internet,” US Internet Industry Association, May1. <http://www.usiia.org/pubs/The%20Exabyte%20Internet.pdf>

Members of the Technical Staff (1983) “Chapter 5: Traffic” in R. Rey (ed.) *Engineering and Operations in the Bell System, 2nd Edition*, AT&T Bell Laboratories: Murray Hill, New Jersey.

Mitchell, T. (2000) “Avoiding the Pitfalls of Oversubscription in DSL Networks,” *xchange*, April 1. <http://www.xchangemag.com/articles/041feat1.html>

Norton, W. (2007) “Video Internet: The Next Wave of Massive Disruption to the U.S. Peering Ecosystem (v1.3),” February 12. http://www.apricot2007.net/presentation/conference/interprovider_stream/peering-Internet_Video_Next%20Wave_of_disruption_v1.3.pdf

Odlyzko, A. (2009) “Network Neutrality, Search Neutrality, and the Never-ending Conflict between Efficiency and Fairness in Markets,” *Review of Network Economics*, 8: 40-60.

Ofcom (2007) “The Communications Market: Broadband Digital Progress Report,” April 2.

Ortiz Jr., S. (2006) “Phone Companies Get into the TV Business,” *IEEE Computer*, 39: 12-15.

Ou, G. (2008) “Comments” *In the Matter of Petition for Declaratory Ruling by Free Press, et al, and Petition for Rulemaking by Vuze. Inc.*, Federal Communications Commission WC Docket No. 07-52, February 13.

Penston, J. (2007) “What’s the Real Cost of Internet Video?” *The Register*, April 20.

Public Knowledge, Consumer Federation of America, EDUCAUSE, Electronic Frontier Foundation, Electronic Privacy Information Center, Freeculture.org, Free Press, Knowledge Ecology International, Media Access Project, New America Foundation, U.S. Public Interest Research Group (2007) “Reply Comments” *In the Matter of Broadband Industry Practices*, Federal Communications Commission WC Docket No. 07-52, July 12.

Rappoport, P. and L. Taylor (2006) “Willingness-to-Pay and the Demand for Telco Video Services: Video Only or Video as Part of a Triple-Play Bundle,” Working Paper, Department of Economics, Temple University, April.

Ruel, B. (2004) “Compression: An evolution, not a revolution” *BroadcastEngineering*, April 1.

Svensson, P. (2007) “Comcast Blocks Some Internet Traffic,” *San Francisco Chronicle*, October 19.

Van, J. (2006) “On Internet2, innovating at higher speed,” *Chicago Tribune*, December 6.

Wilson, C. (2006) “VON: BellSouth exec pitches need for IP service tiers,” *TelephonyOnline*, March 15.

Working Group on Net Neutrality (2007) “Report on Network Neutrality,” Ministry of Internal Affairs and Communications of Japan, September 20.

Wu, T. (2003) “Network Neutrality, Broadband Discrimination,” *Journal of Telecommunications and High Technology Law*, 2: 141-178.

Yuksel, M., K. Ramakrishnan, S. Kalyanaraman, J. Houle and R. Sadhvani (2007) “Value of Supporting Class-of-Service in IP Backbones,” *Proceedings of IEEE International Workshop on Quality of Service*. June: 109-112.

7 Appendix

This appendix contains the key input values and output calculations from a simple model of unmanaged PON network costs. The parameter values used in this model should be considered suggestive only. While taken as a whole, this modeling is believed to generate reasonable estimates of underlying cost levels and engineering relationships, but specific elements should not be considered individually precise or reflective of the actual costs incurred by any particular operator of an unmanaged PON in any particular geographical location. To the extent that more accurate input values are available for any of these geographical cost or engineering parameters, these values may be easily substituted into the spreadsheet model and revised cost results obtained.

Network engineering and financial parameters common to all model runs

Network structure characteristics			
Total number of BB lines at WC:	12,500		
Number of WCs in cluster:	8		
Total BB lines modeled:	100,000		
PON capacity code:	1 or 3		Rural
Maximum fiber splits:	32		16
Fiber splits at drop terminal:	4		4
Fiber splits at FDT:	8		4
Sharing factor for WC-FDT runs:	50%		
Sharing factor for FDT-terminal runs:	40%		
Average WC to WC distance:	36,000	feet	
Sharing factor for WC-WC runs:	33%		
Fibers per WC-WC route:	min(12,4xOC768)		
Network router capacity sizing factor:	75%		
Fraction of traffic leaving cluster:	85%		
Geographical outside plant characteristics			
	Rural	Suburban	Urban
Percent of total lines:	10%	60%	30%
Number of lines:	10,000	60,000	30,000
OLTs serving each FDT location:	1	4	8
Cable tapers per distribution route:	1	2	3
Cable tapers per feeder route:	2	2	3
Average WC-FDT distance:	24,000	12,000	6,000
Average FDT-terminal distance:	6,000	3,000	1,000
Average drop distance:	150	100	50
Percent aerial structure:	30%	40%	20%
Percent buried structure:	70%	60%	20%
Percent underground structure:	0%	0%	60%
Drop terminal sizing factor:	50%	70%	80%
FDT sizing factor:	70%	80%	90%
Financial characteristics			
BB carrier WACC:	14.00%		
Tax rate on equity return:	39.25%		
Other taxes on investment:	1.00%		
<u>Capital category</u>	<u>Recovery period</u>	<u>Maintenance Factor</u>	
Buildings	40 years	3.50%	
General Support	7 years	5.00%	
CO circuit/switch	10 years	3.00%	
Premises circuit	15 years	2.00%	
Drop cables/terminals:	25 years	0.70%	
FDT:	20 years	1.00%	
Poles	25 years	1.00%	
Aerial fiber	20 years	0.50%	
Buried fiber	25 years	0.30%	
Underground fiber	20 years	0.15%	
Conduit	50 years	0.10%	

Modeling results: current typical usage

Subscriber usage characteristics	Download	Upload		
Subscriber port capacity:	19.4	4.9	Mbps	
Subscriber port average utilization:	0.08%	0.08%		
Average offered traffic per subscriber:	0.016	0.004	Mbps	
Implied average overbooking:	1,250	1,250		
Subscriber port busy-hour utilization:	0.23%	0.23%		
Offered BH traffic per subscriber:	0.045	0.011	Mbps	
Implied BH overbooking:	435	435		
Minimum required PON capacity:	622	156	Mbps	
Calculated costs:	<u>Investment</u>	<u>Inv. per line</u>	<u>Monthly cost</u>	<u>per line</u>
Terminal/premises:	\$76,104,464	\$761.04	\$1,404,737	\$14.05
Aerial plant:	\$26,482,417	\$802.50	\$473,218	\$14.34
Buried plant:	\$55,045,288	\$1,123.37	\$957,469	\$19.54
Underground plant:	\$26,077,091	\$1,448.73	\$457,943	\$25.44
Rural outside plant total:	\$35,603,155	\$3,560.32	\$632,072	\$63.21
Suburban outside plant total:	\$93,909,610	\$1,565.16	\$1,688,959	\$28.15
Urban outside plant total:	\$54,196,496	\$1,806.55	\$972,335	\$32.41
Outside plant total:	\$183,709,261	\$1,837.09	\$3,293,365	\$32.93
Wire center/cluster hub:	\$16,452,892	\$164.53	\$351,846	\$3.52
Overall access network total:	\$200,162,153	\$2,001.62	\$3,645,212	\$36.45
		Operating expense total:	\$1,200,000	\$12.00
		Backbone total:	\$129,801	\$1.30
		Rural total:	\$680,237	\$80.02
		Suburban total:	\$1,977,947	\$44.97
		Urban total:	\$1,116,829	\$49.23
		Grand total:	\$3,775,012	\$49.75

Modeling results: current power usage

Subscriber usage characteristics	Download	Upload
Subscriber port capacity:	19.4	4.9 Mbps
Subscriber port average utilization:	0.80%	0.80%
Average offered traffic per subscriber:	0.156	0.039 Mbps
Implied average overbooking:	125	125
Subscriber port busy-hour utilization:	2.30%	2.30%
Offered BH traffic per subscriber:	0.447	0.112 Mbps
Implied BH overbooking:	43	43
Minimum required PON capacity:	622	156 Mbps

Calculated costs:	Investment	Inv. per line	Monthly cost	per line
Terminal/premises:	\$76,104,464	\$761.04	\$1,404,737	\$14.05
Aerial plant:	\$26,482,417	\$802.50	\$473,218	\$14.34
Buried plant:	\$55,045,288	\$1,123.37	\$957,469	\$19.54
Underground plant:	\$26,077,091	\$1,448.73	\$457,943	\$25.44
Rural outside plant total:	\$35,603,155	\$3,560.32	\$632,072	\$63.21
Suburban outside plant total:	\$93,909,610	\$1,565.16	\$1,688,959	\$28.15
Urban outside plant total:	\$54,196,496	\$1,806.55	\$972,335	\$32.41
Outside plant total:	\$183,709,261	\$1,837.09	\$3,293,365	\$32.93
Wire center/cluster hub:	\$18,492,892	\$184.93	\$397,502	\$3.98
Overall access network total:	\$202,202,153	\$2,022.02	\$3,690,867	\$36.91
	Operating expense total:		\$1,240,000	\$12.40
		Backbone total:	\$884,006	\$8.84
		Rural total:	\$760,223	\$88.42
		Suburban total:	\$2,457,863	\$53.36
		Urban total:	\$1,356,787	\$57.63
		Grand total:	\$4,574,873	\$58.15

Modeling results: future modest video usage

Subscriber usage characteristics	Download	Upload		
Subscriber port capacity:	19.4	4.9	Mbps	
Subscriber port average utilization:	10.00%	10.00%		
Average offered traffic per subscriber:	1.944	0.486	Mbps	
Implied average overbooking:	10	10		
Subscriber port busy-hour utilization:	28.30%	28.30%		
Offered BH traffic per subscriber:	5.501	1.375	Mbps	
Implied BH overbooking:	4	4		
Minimum required PON capacity:	622	156	Mbps	

Calculated costs:	Investment	Inv. per line	Monthly cost	per line
Terminal/premises:	\$76,104,464	\$761.04	\$1,404,737	\$14.05
Aerial plant:	\$26,482,417	\$802.50	\$473,218	\$14.34
Buried plant:	\$55,045,288	\$1,123.37	\$957,469	\$19.54
Underground plant:	\$26,077,091	\$1,448.73	\$457,943	\$25.44
Rural outside plant total:	\$35,603,155	\$3,560.32	\$632,072	\$63.21
Suburban outside plant total:	\$93,909,610	\$1,565.16	\$1,688,959	\$28.15
Urban outside plant total:	\$54,196,496	\$1,806.55	\$972,335	\$32.41
Outside plant total:	\$183,709,261	\$1,837.09	\$3,293,365	\$32.93
Wire center/cluster hub:	\$47,982,892	\$479.83	\$1,057,491	\$10.57
Overall access network total:	\$231,692,153	\$2,316.92	\$4,350,856	\$43.51
		Operating expense total:	\$1,360,000	\$13.60
		Backbone total:	\$8,613,536	\$86.14
		Rural total:	\$1,599,175	\$173.52
		Suburban total:	\$7,491,575	\$138.46
		Urban total:	\$3,873,643	\$142.72
		Grand total:	\$12,964,392	\$143.24

Modeling results: 50/50 split between future modest and typical video usage

Subscriber usage characteristics	Download	Upload		
Subscriber port capacity:	38.9	19.4	Mbps	
Subscriber port average utilization:	12.00%	12.00%		
Average offered traffic per subscriber:	4.665	2.333	Mbps	
Implied average overbooking:	8	8		
Subscriber port busy-hour utilization:	34.73%	34.73%		
Offered BH traffic per subscriber:	13.501	6.751	Mbps	
Implied BH overbooking:	3	3		
Minimum required PON capacity:	1,244	622	Mbps	

Calculated costs:	Investment	Inv. per line	Monthly cost	per line
Terminal/premises:	\$81,104,464	\$811.04	\$1,502,783	\$15.03
Aerial plant:	\$26,482,417	\$802.50	\$473,218	\$14.34
Buried plant:	\$55,045,288	\$1,123.37	\$957,469	\$19.54
Underground plant:	\$26,077,091	\$1,448.73	\$457,943	\$25.44
Rural outside plant total:	\$36,103,155	\$3,610.32	\$641,877	\$64.19
Suburban outside plant total:	\$96,909,610	\$1,615.16	\$1,747,787	\$29.13
Urban outside plant total:	\$55,696,496	\$1,856.55	\$1,001,749	\$33.39
Outside plant total:	\$188,709,261	\$1,887.09	\$3,391,412	\$33.91
Wire center/cluster hub:	\$86,190,992	\$861.91	\$1,912,131	\$19.12
Overall access network total:	\$274,900,253	\$2,749.00	\$5,303,543	\$53.04
		Operating expense total:	\$1,560,000	\$15.60
		Backbone total:	\$21,114,142	\$211.14
		Rural total:	\$2,944,504	\$310.05
		Suburban total:	\$15,563,550	\$274.99
		Urban total:	\$7,909,630	\$279.25
		Grand total:	\$26,417,684	\$279.78

Modeling results: future typical video usage

Subscriber usage characteristics	Download	Upload		
Subscriber port capacity:	38.9	19.4	Mbps	
Subscriber port average utilization:	20.00%	20.00%		
Average offered traffic per subscriber:	7.775	3.888	Mbps	
Implied average overbooking:	5	5		
Subscriber port busy-hour utilization:	55.31%	55.31%		
Offered BH traffic per subscriber:	21.502	10.751	Mbps	
Implied BH overbooking:	2	2		
Minimum required PON capacity:	1,244	622	Mbps	

Calculated costs:	Investment	Inv. per line	Monthly cost	per line
Terminal/premises:	\$81,104,464	\$811.04	\$1,502,783	\$15.03
Aerial plant:	\$26,482,417	\$802.50	\$473,218	\$14.34
Buried plant:	\$55,045,288	\$1,123.37	\$957,469	\$19.54
Underground plant:	\$26,077,091	\$1,448.73	\$457,943	\$25.44
Rural outside plant total:	\$36,103,155	\$3,610.32	\$641,877	\$64.19
Suburban outside plant total:	\$96,909,610	\$1,615.16	\$1,747,787	\$29.13
Urban outside plant total:	\$55,696,496	\$1,856.55	\$1,001,749	\$33.39
Outside plant total:	\$188,709,261	\$1,887.09	\$3,391,412	\$33.91
Wire center/cluster hub:	\$130,985,492	\$1,309.85	\$2,914,175	\$29.14
Overall access network total:	\$319,694,753	\$3,196.95	\$6,305,587	\$63.06
		Operating expense total:	\$1,660,000	\$16.60
		Backbone total:	\$33,614,747	\$336.15
		Rural total:	\$4,294,769	\$446.08
		Suburban total:	\$23,665,140	\$411.02
		Urban total:	\$11,960,425	\$415.28
		Grand total:	\$39,920,334	\$415.80