Costs of Neutral/Unmanaged IP Networks

Richard N. Clarke[*]
Director of Economic Analysis, AT&T
Phone: +1.202.457.2308
Email: rnclarke@att.com

Executive Summary

Debate over "network neutrality" has tended to focus on issues such as who "manages" the Internet and how to ensure its maximum utility. Notably absent is any examination of the costs to consumers of the various management models and whether consumers can afford unmanaged completely neutral IP networks.

This paper examines the cost of an unmanaged network and demonstrates that as Internet usage patterns evolve and become both more bandwidth-intensive and real-time oriented, a model of pure neutrality would be extremely expensive for the typical consumer. Right now, the Internet activities of most residential consumers center around e-mail, web browsing and modest file downloads. Thus, the shared elements of current IP networks are scaled to accommodate the aggregate of this moderate use. Other applications, such as real-time, television-quality video services and business services typically have been delivered over managed networks.

However, usage is beginning to explode. As high bandwidth, real-time services such as streaming music and video proliferate on the Internet, existing IP network capacity is being exhausted and large upgrades to the local access, regional and core backbone portions of the Internet are required. The extent of these required upgrades depends on the degree of management network operators are allowed to exercise. If a completely unmanaged/neutral network is required, the required upgrade investments appear to be massive. Indeed, to provide sufficient network capacity to accommodate expected growth in traditional Internet data services as well as use of Internet connections for bandwidth-intensive applications equivalent to just two simultaneous SDTV channels per home, customers may need to pay roughly $140 per month for their Internet service only – before they additionally pay for the video content itself. But if customer data use or viewing habits become more intense – say, the equivalent of viewing two simultaneous HDTV channels, Internet service bills may reach $466 per month.

It seems unlikely that enough customers would be willing to pay the fees (between $140 and $466 per month) that fully unmanaged IP networks would require. Such fees are so high as to make such networks commercially nonviable.

---

[*] 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.

1. Overview

The words "network neutrality" connote many things to many people. To some proponents, network neutrality connotes 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. Certain network neutrality proponents suggest that the use of any network management controls would violate principles of network neutrality and would reduce the social and technological benefits offered by the Internet. These proponents generally argue that investment in network management capabilities can be avoided if network providers simply focus their investment on building "bigger and 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]

---

[1] The following articles offer a variety of arguments opposing IP carrier application of network management controls: http://www.thenation.com/docprint.mhtml?i=20060213&s=chester; http://www.netparadox.com/netparadox.html; http://judiciary.house.gov/media/pdfs/misener042506.pdf; http://www.techreview.com/read_article.aspx?id=13234&ch=infotech; http://www.pcmag.com/article2/0,1895,1937952,00.asp; http://www.democraticmedia.org/PDFs/LessigNN.pdf; http://www.commoncause.org/atf/cf/{FB3C17E2-CDD1-4DF6-92BE-BD4429893665}/NETWORK_NEUTRALITY_FACTSHEET.PDF.
While some of the positions expressed by these neutrality proponents suggest that certain nondiscriminatory 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 (nondiscriminatory) half slice of the available bandwidth would likely result in neither application executing 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 without the efficiencies offered by network management. The cost model developed in this analysis suggests that customers would be required to pay between $140 and $466 per month for adequate unmanaged transmission capability – depending on whether customers are using applications that involve Internet throughput roughly equivalent to viewing only "standard definition" video – or are able to consume Internet usage equivalent to watching modern "high definition" video. 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.

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.[2] 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

---

[2] See, Internet Society, *All About the Internet: History of the Internet,* version 3.32, last revised 10 December 2003 ("… while file transfer and remote login (Telnet) were very important applications, electronic mail has probably

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.[3] 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.[4] The result has been an increasing consumer shift from dial-up connections to broadband connections.[5]

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 customer loops that is higher in frequency than the voice spectrum. But since signals in this spectrum will not pass through voice switches, before the loop meets the Class 5 switch, the data signals in this high frequency spectrum are peeled off by a digital subscriber line access multiplexer ("DSLAM"). 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.

---

had the most significant impact of the innovations from that era."). Available at: http://www.isoc.org/internet/history/brief.shtml.

[3] 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 kilobytes per second.

[4] See, "The Broadband Incentive Problem," white paper by the Broadband Working Group of the MIT Communications Futures Program, September 2005 for a discussion of evolving Internet use. Available at: http://cfp.mit.edu/groups/broadband/docs/2005/Incentive_Whitepaper_09-28-05.pdf.

[5] 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 that are scattered around the country and world – and are not "local" to the customer.

IP Network Schematic

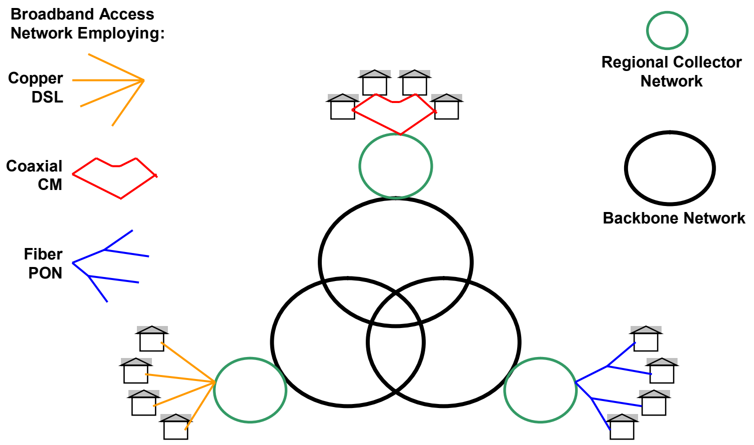

Cable modem ("CM") systems utilize the abundant spectrum available in coaxial cable (up to 860 MHz) to provide data transmission service to customers. Data signals are brought from the cable system's headend into a neighborhood riding (segregated from the system's video broadcast signals) on optical fiber cables. 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 VHF television channels in their native RF spectrum, only portions of the coax spectrum between 5 and 42 MHz are usable for upstream links.[6]

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

---

[6] See, http://www.iec.org/online/tutorials/hfc_dwdm/.

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 back into the Internet backbone.

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 DSLAM (greater than 4000 feet) generally will get less throughput.[7] Typically, upstream bandwidth is 1.5 Mbps or less, but the division of total available 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.[8] 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.[9]

The PON being deployed by Verizon and marketed under the FiOS^{SM} brand 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

---

[7] See, Piyush Sevalia, "Swimming Upstream: The Case for Higher Speeds." Available at: http://www.convergedigest.com/whitepapers/documents/Ikanos-VDSL.pdf.

[8] See, http://www.cablemodem.com/ or http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/cable.pdf.

[9] Cable system fiber nodes rarely serve less than 125 to 500 customers.

subscriber.[10] 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.[11]

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 after today's most advanced video compression ("MPEG-4/H.264"), standard definition television ("SDTV") channels require roughly 2 Mbps of throughput – and high definition television ("HDTV") channels require roughly 10 Mbps of throughput.[12] But at least as 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.[13] 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

---

[10] See, 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.

[11] 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.

[12] 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, http://www.bigbandnet.com/technology/tech_whitepaper_hdtv_roll.php and http://broadcastengineering.com/mag/broadcasting_compression_evolution_not/index.html.

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

than one program simultaneously, required capacities scale up directly with the number of programs being viewed (e.g., three SDTV channels require 6 Mbps; two HDTV channels require 20 Mbps).

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.[14] 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 truly residential users – or they are business customers that demand (and pay extra for) IP network connections that provide dedicated guaranteed throughput.[15] 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 connection – and even use during the network busy-hour may not exceed 3% of maximum access bandwidth.

---

[14] See, "Everything on the Net" presentation by BellSouth Chief Architect Hank Kafka. Available at: 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.

[15] 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). Because these connections are engineered to guarantee these data transmission rates, their prices are commensurately higher than corresponding speed residential broadband access connections.

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.[16] 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 if today's shared IP collection and backbone networks are sized to offer less than 50 Kbps of capacity per residential customer.[17]

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.[18]

---

[16] This is because certain of the resources in these access networks are "oversubscribed." See, http://www.xchangemag.com/articles/041feat1.html for a discussion of oversubscription (also called contention or overbooking) in broadband access networks.

[17] See "The Broadband Incentive Problem" MIT white paper for a discussion of how evolving usage patterns are exhausting current IP network capacity. This white paper also notes that even in Korea which leads the world in broadband penetration, per-subscriber traffic barely exceeds 50 Kbps.

[18] 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.

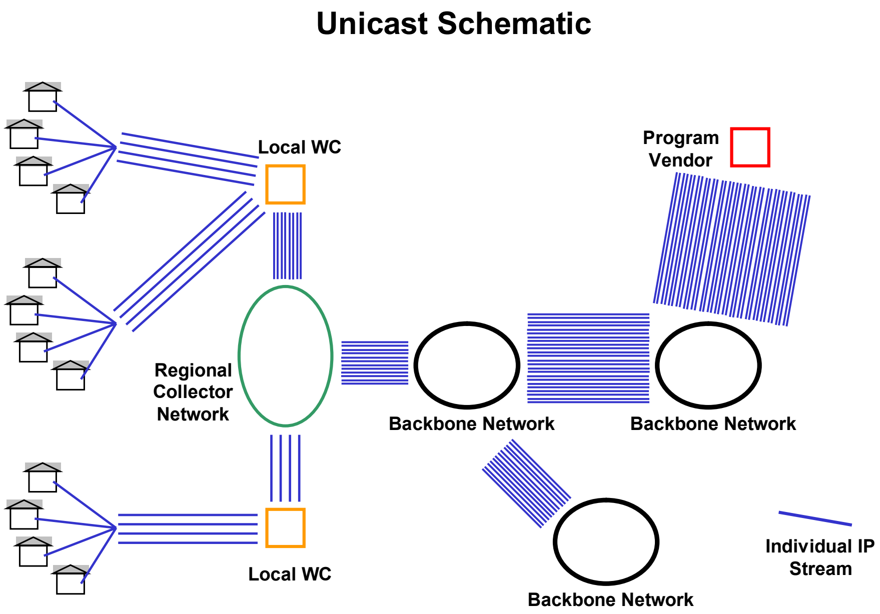

Therefore, for every customer that wishes to view an independently-streamed HDTV program during the network busy-hour (which is typically early evening – the same as prime viewing time), backbone and regional IP links need to be enlarged by 10 Mbps of capacity. And because other customers seeking to view streaming video will most likely 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 less than 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.[19]

---

[19] 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. 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 – prohibited under this unmanaged model.

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 (e.g., up to 100 Mbps), these customers could run whatever applications they desire (e.g., 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.[20] Furthermore, they suggest that thick unmanaged networks will cost less than overlaying management on thinner networks.[21]

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.[22] Instead, these advocates argue that a managed network

---

[20] Acronym VoIP is "Voice over Internet Protocol" service and acronym P2P means "peer-to-peer." The latter type of 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.

[21] See, Gary R. Bachula, Vice President, Internet2, testimony presented at United States Senate Committee on Commerce, Science and Transportation Hearing on Net Neutrality (February 7, 2006). Available at: http://commerce.senate.gov/pdf/bachula-020706.pdf.

[22] See, Cisco Systems White Paper, "Cisco Service Control: A Guide to Sustained Broadband Profitability," November 2005. Available at: http://www.democraticmedia.org/PDFs/CiscoBroadbandProfit.pdf.

holds better promise to provide all of the applications desired by customers at a much smaller cost than an unmanaged network.

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.

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.[23] 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 signal packets from all other customers served by that DSLAM (typically several hundred) and transmitted back

---

[23] See, ENGINEERING AND OPERATIONS IN THE BELL SYSTEM, 2ND EDITION, R.F. Rey, ed., AT&T Bell Laboratories, 1983, pp. 147-191.

into a high speed IP collection network for 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 (e.g., 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.

DSL Schematic

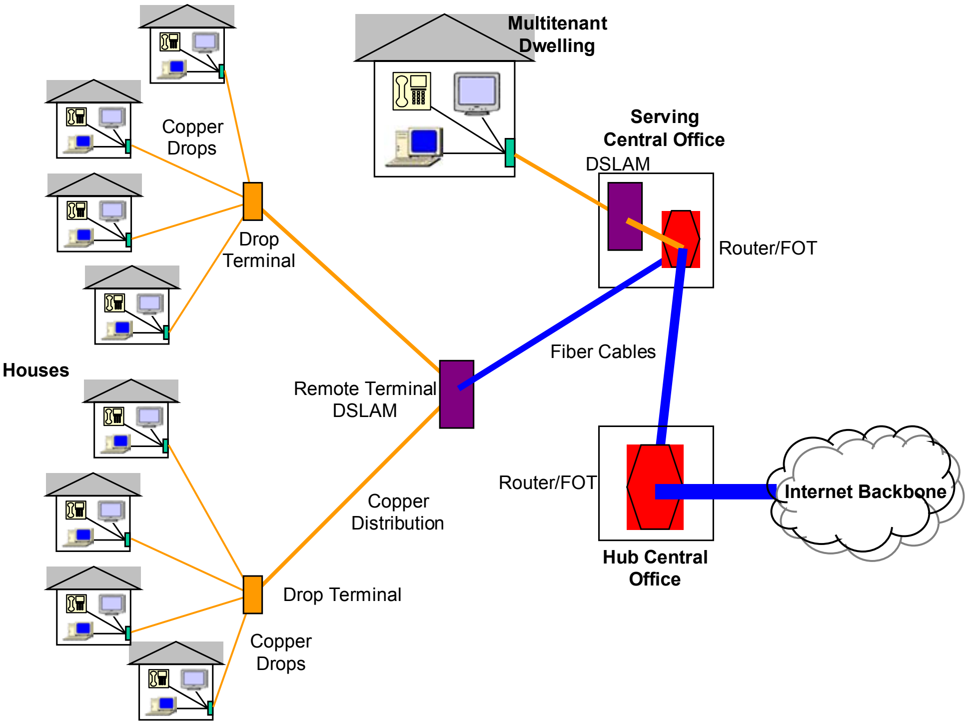

In a cable broadband network, resource sharing begins even sooner – at the customer's cable modem. This modem shares a given amount of bandwidth (e.g., 38 Mbps in basic DOCSIS networks) with typically several hundred other customers in the neighborhood. While customers may have their personal maximum 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.

CM Schematic

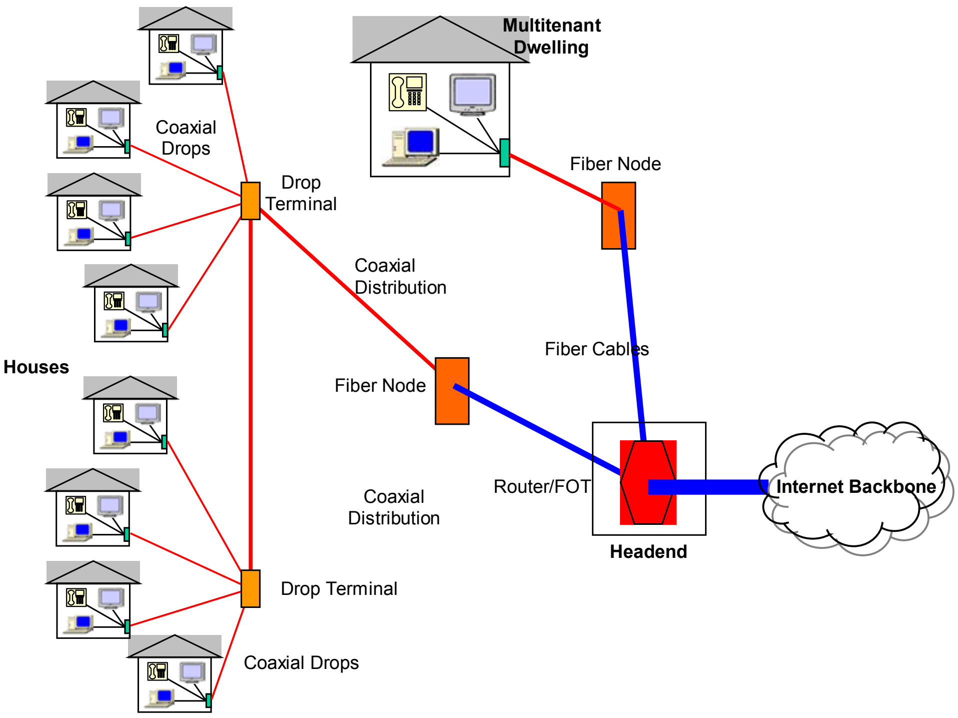

PONs are similar in architecture to the neighborhood distribution segment of CM networks, but use optical fiber rather than copper coax and offer much higher capacities (e.g., 622 Mbps and higher). Because the number of customers sharing neighborhood bandwidth is much more rigidly limited and controlled (e.g., 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.

PON Schematic

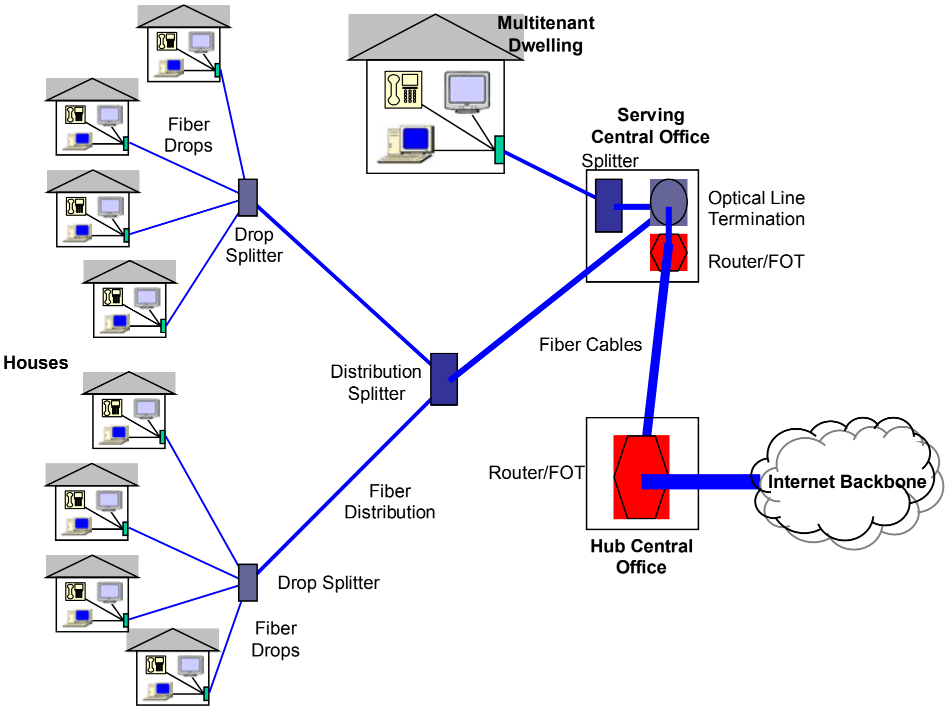

The degree to which today's broadband access, collector and backbone networks need to be expanded to accommodate tomorrow's demands for high bandwidth real-time services depends on existing amounts of excess capacity and the service delivery technology.

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 (e.g., 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.

Because today's shared IP collection and backbone networks are sized to offer less than 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 nearly similar factor. This is because these shared resource collection and backbone networks are already extremely high capacity networks that carry many services (e.g., 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.[24]

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

---

[24] Another way of stating this is that most scale economies in the construction and operation of high speed IP collection and backbone networks have already 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.

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 an order of magnitude.

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.[25] 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 very 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.[26] This fiber feeder strand carries an extremely high bandwidth signal (e.g., 622 Mbps, 1244 Mbps, 2488 Mbps, etc.) out to a fiber distribution terminal ("FDT") located in the neighborhood. At the FDT, a passive optical

---

[25] See, Credit Suisse First Boston Equity Research, "VZ Launches FiOS TV; Who's Most Exposed?" 22 September 2005 for a description of the architecture and costs associated with Verizon's FiOS PON. See, also, Anupam Banerjee and Marvin Sirbu, "Towards Technologically and Competitively Neutral Fiber to the Home (FTTH) Infrastructure," for a more general discussion of PON architectures and costs – available at: http://www.andrew.cmu.edu/user/sirbu/pubs/Banerjee_Sirbu.pdf.

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 a separate 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.[27]

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 (e.g., the 1/32nd or the 1/64th 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 (i.e., up on poles), as buried (i.e., placed in a trench directly in the soil), or as underground (i.e., 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 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 are 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

---

[26] The equipment on which this fiber strand terminated at the central office is called an OLT for "optical line termination."

[27] 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 × 4).

location is rural, 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 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.[28]

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").[29] 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

---

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

[29] 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.

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.

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 (e.g., power or surveillance), customer operations costs (e.g., service creation, ordering or billing) and corporate overheads (e.g., 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.

Of most significance is how costs vary based on different assumptions as to the services used by subscribers. The following table presents the traffic usage levels under several alternative use profiles. The first is today's typical user profile – data use only. The second profile represents the use characteristics of a current "power user" – again, data use only.[30] The third profile assumes modest future demand for video services (i.e., two simultaneous SDTV channels – or their bandwidth equivalents) plus expanded data use that is three times higher than current "power" use. The fourth profile assumes video demand levels that are likely to be typical of tomorrow's viewers (i.e., 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.

---

[30] 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, http://times.hankooki.com/lpage/tech/200503/kt2005032817531511810.htm .

	Busy-Hour Download Traffic Rates		
Usage profile	**Data**	**Video**	**Total**
Current typical user	45 Kbps	--	45 Kbps
Current power user	450 Kbps	--	450 Kbps
Future modest video user (2 SDTV channels)	1.5 Mbps	4.0 Mbps	5.5 Mbps
Future typical video user (1 HDTV + 3 SDTV channels or 2 HDTV channels)	1.5 Mbps	20.0 Mbps	21.5 Mbps

Based on these usage profiles, it is possible to determine the costs of satisfying each of these usage demands over a "thick" unmanaged PON.[31]

Alternative user profiles - unmanaged network

BH Capacity Utilization	Outside Plant Cost	WC+Cluster Cost	Operating Cost	Backbone Cost	Total Cost
Current typical usage	$30.64	$2.77	$12.00	$1.30	$46.71
Current power usage	$30.64	$3.07	$12.00	$8.84	$54.55
Modest future video usage	$30.64	$11.32	$12.00	$86.14	$140.09
Typical future video usage	$31.62	$86.49	$12.00	$336.15	$466.26

---

[31] 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.

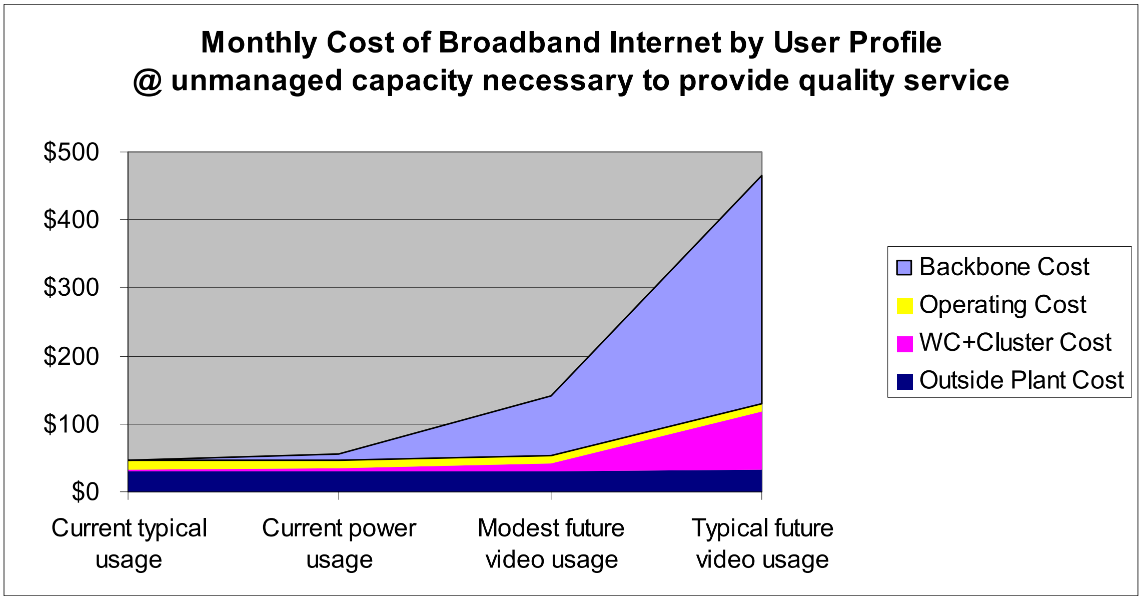

These costs are quite consistent with other estimates of the cost of providing streaming video services over broadband.[32] Thus it is clear that the cost of sustaining IP access, regional collection and backbone networks robust enough to satisfy tomorrow's demands for high bandwidth services such as video on an unmanaged basis are extremely high.[33] 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 roughly 45 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 and backbones as well, the largest percentage and dollar increases come from this required expansion of the collector and backbone networks.

---

[32] 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.networkcomputing.com/showitem.jhtml?articleID=183702346 or http://telephonyonline.com/voip/news/VON_BellSouth_tiers_031506/index.html.

[33] 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

5. Conclusions

At today's typical or power usage levels, customers generally pay between $15 and $45 per month for their Internet access service.[34] 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 of tomorrow may need to pay $140 per month for their Internet access service. A more typical user of tomorrow, one who expects to be able to receive two HDTV channels simultaneously in his home, may have to pay roughly $466 per month. But even these prices purchase only the ability 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.[35]

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.[36] Thus, it seems unlikely that unmanaged PONs with capacity adequate to stream unicast video services will gain commercial traction.[37] A network operator that builds such a

---

tomorrow's technologies are not capable of scaling to these large capacities without larger cost increases than contemplated today.

[34] See, Goldman Sachs, "2006 outlook – stuck in neutral," January 13, 2006, Exhibit 19.

[35] 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, "Twelfth Annual Report on Competition in the Market for Delivery of Video Programming, ¶¶ 43-44. Available at: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-06-11A1.pdf 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.

[36] See, Paul Rappoport and Lester D. Taylor, "Willingness-to-Pay and the Demand for Telco Video Services: Video Only or Video as part of a Triple-Play Bundle," Temple University Department of Economics working paper, April 2006. Available at: http://www.mbc-thebridge.com/iptv.pdf.

[37] 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.

"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, fully unmanaged IP networks may not be a feasible goal.

May 2006

Appendix

This appendix contains key input values and output calculations of 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
Maximum fiber splits:	32
Fiber splits at drop terminal:	4
Fiber splits at FDT:	8
Sharing factor for WC-FDT runs:	50%
Sharing factor for FDT-terminal runs:	50%
Average WC to WC distance:	30,000 feet
Sharing factor for WC-WC runs:	33%
Fibers per WC-WC route:	min(12,#OC768)
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:	1	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:	60%	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%

Capital category	Recovery period	Maintenance Factor
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:	Investment	Inv. per line	Monthly cost	per line
Terminal/premises:	$75,991,964	$759.92	$1,402,744	$14.03
Aerial plant:	$23,563,642	$714.05	$420,837	$12.75
Buried plant:	$48,794,331	$995.80	$848,568	$17.32
Underground plant:	$22,321,091	$1,240.06	$391,970	$21.78
Rural outside plant total:	**$31,637,422**	**$3,163.74**	**$562,112**	**$56.21**
Suburban outside plant total:	**$89,032,010**	**$1,483.87**	**$1,603,332**	**$26.72**
Urban outside plant total:	**$50,001,596**	**$1,666.72**	**$898,675**	**$29.96**
Outside plant total:	**$170,671,028**	**$1,706.71**	**$3,064,120**	**$30.64**
Wire center/cluster hub:	$12,803,118	$128.03	$276,988	$2.77
Overall access network total:	**$183,474,146**	**$1,834.74**	**$3,341,108**	**$33.41**
		Operating expense total:	**$1,200,000**	**$12.00**
		Backbone total:	**$129,801**	**$1.30**
		Rural total:	**$602,791**	**$72.28**
		Suburban total:	**$1,847,406**	**$42.79**
		Urban total:	**$1,020,712**	**$46.02**
		Grand total:	**$3,470,909**	**$46.71**

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:	$75,991,964	$759.92	$1,402,744	$14.03
Aerial plant:	$23,563,642	$714.05	$420,837	$12.75
Buried plant:	$48,794,331	$995.80	$848,568	$17.32
Underground plant:	$22,321,091	$1,240.06	$391,970	$21.78
Rural outside plant total:	**$31,637,422**	**$3,163.74**	**$562,112**	**$56.21**
Suburban outside plant total:	**$89,032,010**	**$1,483.87**	**$1,603,332**	**$26.72**
Urban outside plant total:	**$50,001,596**	**$1,666.72**	**$898,675**	**$29.96**
Outside plant total:	**$170,671,028**	**$1,706.71**	**$3,064,120**	**$30.64**
Wire center/cluster hub:	$14,153,118	$141.53	$307,201	$3.07
Overall access network total:	**$184,824,146**	**$1,848.24**	**$3,371,321**	**$33.71**
		Operating expense total:	**$1,200,000**	**$12.00**
		Backbone total:	**$884,006**	**$8.84**
		Rural total:	**$681,233**	**$80.12**
		Suburban total:	**$2,318,057**	**$50.63**
		Urban total:	**$1,256,037**	**$53.87**
		Grand total:	**$4,255,327**	**$54.55**

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:	$75,991,964	$759.92	$1,402,744	$14.03
Aerial plant:	$23,563,642	$714.05	$420,837	$12.75
Buried plant:	$48,794,331	$995.80	$848,568	$17.32
Underground plant:	$22,321,091	$1,240.06	$391,970	$21.78
Rural outside plant total:	**$31,637,422**	**$3,163.74**	**$562,112**	**$56.21**
Suburban outside plant total:	**$89,032,010**	**$1,483.87**	**$1,603,332**	**$26.72**
Urban outside plant total:	**$50,001,596**	**$1,666.72**	**$898,675**	**$29.96**
Outside plant total:	**$170,671,028**	**$1,706.71**	**$3,064,120**	**$30.64**
Wire center/cluster hub:	$50,993,118	$509.93	$1,131,684	$11.32
Overall access network total:	**$221,664,146**	**$2,216.64**	**$4,195,804**	**$41.96**
		Operating expense total:	**$1,200,000**	**$12.00**
		Backbone total:	**$8,613,536**	**$86.14**
		Rural total:	**$1,536,634**	**$165.66**
		Suburban total:	**$7,450,464**	**$136.17**
		Urban total:	**$3,822,241**	**$139.41**
		Grand total:	**$12,809,340**	**$140.09**

Modeling results: future typical video usage

Subscriber usage characteristics	Download	Upload	
Subscriber port capacity:	38.9	19.4	Mbps
Subscriber port average utilization:	18.00%	18.00%	
Average offered traffic per subscriber:	6.998	3.499	Mbps
Implied average overbooking:	6	6	
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:	$80,991,964	$809.92	$1,500,791	$15.01
Aerial plant:	$23,563,642	$714.05	$420,837	$12.75
Buried plant:	$48,794,331	$995.80	$848,568	$17.32
Underground plant:	$22,321,091	$1,240.06	$391,970	$21.78
Rural outside plant total:	**$32,137,422**	**$3,213.74**	**$571,917**	**$57.19**
Suburban outside plant total:	**$92,032,010**	**$1,533.87**	**$1,662,161**	**$27.70**
Urban outside plant total:	**$51,501,596**	**$1,716.72**	**$928,089**	**$30.94**
Outside plant total:	**$175,671,028**	**$1,756.71**	**$3,162,167**	**$31.62**
Wire center/cluster hub:	$386,877,918	$3,868.78	$8,648,819	$86.49
Overall access network total:	**$562,548,946**	**$5,625.49**	**$11,810,985**	**$118.11**
Operating expense total:			**$1,200,000**	**$12.00**
Backbone total:			**$33,614,747**	**$336.15**
Rural total:			**$4,798,273**	**$491.83**
Suburban total:			**$27,020,300**	**$462.34**
Urban total:			**$13,607,159**	**$465.57**
Grand total:			**$45,425,733**	**$466.26**