

A simple cost model for broadband access: What will video cost?¹

David Clark
TPRC, September, 2008
Revision 1 of August 27, 2008

Abstract

This paper provides a very simple, “back of envelope” analysis of the costs of delivering Internet data to customers over residential broadband facilities. While the retail price of consumer broadband access is well-understood and widely advertised, the fraction of that cost that should be associated with usage is very hard to determine, since access providers typically treat their internal cost structure as highly proprietary. At the same time, some understanding of these costs is very important as we try to understand the true impact of increased usage due to high-bandwidth applications such as video. Using perhaps over-simple cost estimates, I derive some perhaps plausible insights about costs of consumer video over the Internet, the future of consumer pricing for broadband access, and the possible relevance of two-sided market models to describe the access ISP. I conclude that the bandwidth requirements of video will indeed shift the cost model of the access ISP to the point that some change in consumer pricing may result. However, the costs will not make broadband access impractical. Video may in fact raise issues that can be analyzed using the framework of two-sided markets.

Introduction

The landscape of consumer broadband Internet access has shifted a number of times in the last few years, with music sharing and the rise of peer-to-peer (P2P) content distribution schemes, and then the advent of Internet-based video, both P2P and hosted (e.g. Youtube). Most recently, commercial, high-quality video has begun to be distributed to consumers over the Internet.

The rise of these applications, and particularly video, has triggered a number of actions on the part of broadband providers, including the blocking or slowing of some P2P applications, which has in turn triggered calls for regulation of these providers to provide a consumer experience that is free of intervention or manipulation by the providers. The providers have argued that the rapid emergence of Internet video has overloaded parts of their system, caused a material degradation in some aspects of broadband delivery, and added to the costs of the ISP. There have been dire predictions about how much it might cost the consumer to download a movie, and once again predictions that the Internet will not manage to keep up with demand.

¹ This research supported in part by the National Science Foundation under Award No. 0519997, and by the industrial partners of the MIT Communications Futures Program.

One question that lurks behind many of these debates is just how much it actually costs to deliver content to the user. What will it actually cost an ISP to deliver a movie to a consumer over a broadband Internet service? It is not necessary to get a precise answer, but without some rough estimate of actual cost, it is very hard to balance the expressed concerns of the ISPs, the content providers, and the consumer advocates. The purpose of this paper is to do a rough but perhaps defensible cost estimate of the cost of usage, and then draw some conclusions about what these costs might imply for the future of Internet-based video delivery, and the future of pricing for residential broadband.

The paper is divided into two parts: the cost estimates and the analysis and conclusions.

Estimating the cost of Internet usage

Most broadband consumers today purchase access to the Internet on a fixed-cost basis: a high-end (e.g. cable) broadband customer in the U.S. might pay \$40/month for access, with no per-byte charge. Users that transfer “too much” traffic may be dismissed from the system, but other than that, at least in the U.S. there are few ISPs that impose explicit controls on usage. Some ISPs may attempt to shape usage by blocking or degrading certain applications. I will not consider that aspect further in this paper.

It is very hard to find out what fraction of the \$40/month revenue ISPs currently allocate to covering usage-based costs. Most ISP costs have nothing to do with usage: they cover billing, marketing and customer retention, customer service, depreciation on the outside plant and so on. So the first challenge is to estimate just how much traffic a typical consumer sends or receives, and what the cost is to the ISP to carry this traffic.

Ignoring the cost of the access network

In this paper, I propose to exclude from the cost estimates a very expensive component of the cost equation—the access network itself. It may seem odd to exclude from the discussion of costs for broadband access the access network itself, but the logic is as follows. For most sorts of access networks, the cost of construction is an up-front capital expense, a sunk cost in economic terms. Over the short run, very little of that cost relates to the level of usage. (This is not strictly true in some cases: for cable systems (hybrid fiber-coax or HFC), capacity can be increased with incremental expenditures.) None the less, the cost is usually viewed as a capital expenditure independent of incremental usage. Second, the cost structure for different sorts of access networks (DSL, HFC, wireless or FTTH) can be very different, which adds an industry-specific complexity to the model.

There is a critical debate that swirls around this issue—a debate I wish to inform but not to enter. The debate is as follows. The access network is often the point of constriction in access throughput. Congestion often happens in the access network, which can be modeled in economic terms as *congestion costs*. Systems such as HFC and FTTH have very high capacity in principle, but on HFC networks today, for example, almost all of the capacity is dedicated to television, and only a small fraction to Internet broadband access. One conclusion that can be drawn from this fact is that the perceived congestion is not an intrinsic result of the network but is an “artificial” scarcity that results from

business decisions. In this line of reasoning, if this artificial scarcity were to be removed, then users could send as much data as they pleased, and we would not have to debate the cost of congestion or how to control it. While the means proposed to remove this scarcity are sometimes extreme (e.g. making access networks publicly funded, like roads, rather than a private-sector activity), this line of reasoning also presumes that if the artificial scarcity were removed, there would be no other incremental costs of usage, so that usage could really be “free”. It is this hypothesis I want to explore in this paper: if we ignore the cost of the access network itself, what are the other real costs of usage. How close to “free” can usage get?

How should we measure usage?

Consumer broadband service is typically characterized in the market in terms of peak rate, in mb/s. This measure is easy to understand and to calibrate. But it may not match well to the cost of usage. In the interior of the network, where the traffic from many users is mixed and mingled, the peaks and dips of individual usage tend to smooth out, and it is more the average usage of the user over time that signals the cost to serve that user. So instead of an instantaneous measure of usage, in mb/s, ISPs need a measure of usage that better maps to real cost.

One possible measure would be average contribution to total usage during a peak-load period—say a busy hour. The system must be sized to carry the total aggregate load during periods of peak demand, and these are the engineering decisions that drive actual cost. However, at the present time, this measure of usage appears to be too complex to bring to the market, so ISPs often use a longer-term measure of usage, such as GB/month. This measure is perhaps less intuitive than peak rate—most users have no idea how many GB/month they send or receive—but it is probably a better measure of cost than instantaneous peak rate, and (for the context of this paper) a better way to get a handle on the cost of video. So we will use GB/month as the unit measure of consumer broadband.

As a rough starting point, Andrew Odlyzko² estimates total Internet traffic in the U.S. as 3 GB/month per capita. However, this does not really tell us much about residential usage, since first it is total traffic, including business to business, and second, it is per capita, not per on-line household. BellSouth’s Chief Architect Henry Kafka estimated (as of March 2006) that the typical user consumes about 2 GB/month³. Data from the Australian Bureau of Statistics⁴ (one of the few countries that seem to gather this information), shows that average broadband usage was around 3.6 GB/m as of December 2007. Other informal estimates, hard to pin down but oft-repeated, put the number somewhat higher, at 4 to 5 GB/month.

² See <http://www.dtc.umn.edu/mints/home.html>, visited May 19, 2008.

³ See http://telephonyonline.com/iptv/news/BellSouth_VOD_costs_030706/, visited May 19, 2008.

⁴ <http://www.abs.gov.au/AUSSTATS/abs@.nsf/Lookup/8153.0Main+Features1Dec%2007?OpenDocument>, visited July 15, 2008

Estimates of cost are equally vague. Kafka, quoted above, puts the cost to serve the residential customer at \$.50/GB. Wikipedia puts the cost at \$.10/GB⁵. The difference between the two is substantial, but can somewhat be explained by noting that there are many costs to an ISP that can either be viewed as fixed/sunk costs, or on the other hand allocated to usage. Depending on the context, and the point that the speaker wants to make, one can allocate these costs in different ways. BellSouth's Kafka was arguing in his presentation that video was going to generate substantial burdens to the ISP, including replacement of much of their infrastructure, so he might well have been motivated to roll many costs, including some costs associated with the outside plant, into a bucket that is allocated to usage. The author of the Wikipedia article (or David Burstein, as cited above) might well have been motivated to make the opposite point—that bandwidth is actually not that expensive, and might well have looked only at short-term incremental costs. So perhaps with a high of \$.50/GB and a low of \$.10/GB, we have bracketed the actual cost. We can further triangulate on this estimate by looking at some other cost numbers for network access.

What does bulk Internet service cost?

The term for Internet access is *transit*—a service in which an ISP provides connectivity for a customer to all of the Internet. This service is in contrast to *peering*, where two ISPs connect, often in a settlement-free arrangement, in order to exchange between them traffic destined for the other. To some extent, what the residential broadband customer purchases from his access ISP is a retail form of transit, while businesses and small ISPs purchase bulk transit from bigger ISPs.

There are a number of ways of estimating the cost of transit. Some ISPs do provide a price list, but most contracts include discounts and bundled related services. One source of posted prices is the web site for Invisible Hand Networks⁶, a bandwidth broker for both spot and long-term contracts for transit service in New York. According to their data, the price for transit service is around \$30/month for each mb/second, with a low end of \$10/month for each mb/s. These numbers (both average and low-end) are consistent with what one hears by asking actual customers today.

Transit is typically priced in terms of dollars per month for a megabit per second. How does 1 mb/s convert to GB/month? If the link were fully loaded 24 hours a day, 1 mb/s would work out to about 325 GB/month. But no network operator expects (or wants) to fully load transit links. Even during peak periods, one tries to over-provision to avoid persistent congestion, and not all hours are peak hours. If we make the simplistic assumption that the link is 50% loaded on average (and this may still over-estimate how much traffic will actually be sent), then 1 mb/s works out to 162 GB/month.

⁵ See http://en.wikipedia.org/wiki/Broadband_Internet_access, visited May 21, 2008. The Wikipedia article in turn cites an interview with David Burstein (editor of DSL Prime) who gives this number. See <http://bits.blogs.nytimes.com/2008/01/17/time-warner-download-too-much-and-you-might-pay-30-a-movie/?ref=technology>

⁶ See <http://www.invisiblehand.net/>, visited May 19, 2008.

Given our assumption that transit today has an average price of \$30/month for a mb/s, this works out to \$.18/GB. Our bottom price of \$10/month would imply \$.06/GB.

When thinking about these prices (in either formulation), remember that these are bi-directional links. So a provider can potentially have two business activities that exploit the two directions. A content provider is most likely unable to exploit the reverse direction on his access link, but two ISPs with a mix of sources and sink on each net can often use both directions equally.

What do content providers pay?

One way to deliver content into the Internet is by using the services of a Content Delivery Network, or CDN. Many CDNs today host streaming (e.g. real-time audio or video) content, and the prices they charge for service can provide some hints about the cost of data delivery. While actual pricing is almost always proprietary, there are folks who attempt to track the market. One web site run by Dan Rayburn gives the following information as of the end of 2007⁷, where the units are again monthly cost per GB delivered:

Monthly cap	High cost	Low cost
1 TB	\$2.00	\$1.50
10 TB	\$1.20	\$0.89
100 TB	\$0.19	\$0.08

(To get an estimate of what these volumes mean, 100TB/month, assuming a 1 Mb/sec stream for 16 hours a day, is a bit fewer than 500 simultaneous flows.)

Since these are payments from the customer to the CDN, these must represent an upper bound on what the CDN in turn pays the Internet Service Providers to which they attach.

For non-streaming content delivery, a source of pricing is Amazon, which offers a service called S3 (Secure Stable Storage). They charge a monthly storage fee of \$.15/GB and a delivery charge that ranges from \$.17 for the first GB/month down to \$.10 for amounts over 100 TB. (For comparison, the retail price of disk storage today is running around \$.20/GB.) Again, we should ask what Amazon pays to delivery a GB into the Internet at its points of connection. If they charge their own customers \$.10/GB for delivery, their own connection charges have to be less than that.

So as a rough estimate, let us assume that content providers pays \$.05/GB on average to attach to an access network, which is less than transit.

We are now in a position to draw a first rough picture of money flow (not packet flow) for an access ISP.

⁷ See http://blog.streamingmedia.com/the_business_of_online_vi/2007/11/cdn-pricing-dat.html, visited May 16, 2008.

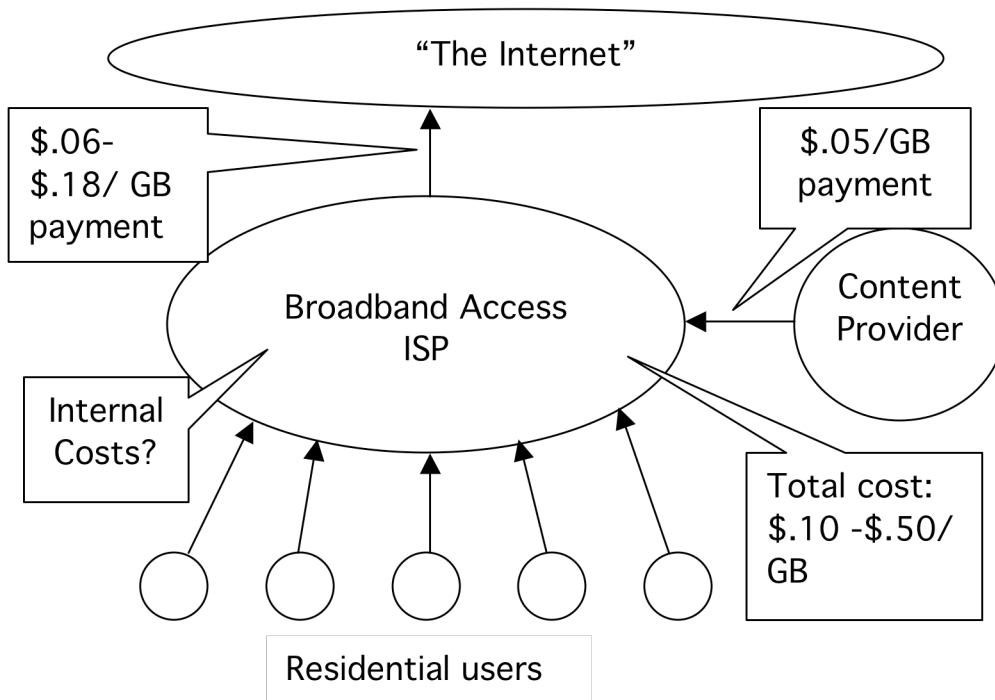


Figure 1: Payment structure for Internet broadband usage

Internal costs

As a cross-check on these numbers, we can build some estimate of the internal cost to the ISP for usage. Again, the ISPs tend to view these facts as proprietary, so we must make some guesses. As I said earlier, I am ignoring the cost of the outside plant. While residential providers of broadband have some very expensive outside plant (the copper pairs, the fiber-coax system or the FTTH system), I am not allocating these costs against usage, since they are fixed costs that are largely sunk. Residential providers presumably have a rich metro infrastructure, which can sometimes be expensive (in terms of cost/mile) in areas where metro capacity is scarce. They also have to reach into rural areas to serve smaller communities. But setting aside the outside plant, an access provider and a provider of high-capacity transit should not look that different. They both have points of presence in metro areas, and metro, regional and (perhaps) national backbone trunks.

To model the cost of such a provider, consider two extremes. A very small access provider might have essentially no internal circuits (so essentially no internal costs), and connect their access network directly to a transit link. In this case, since essentially all their traffic will enter or exit by the transit link, their total cost is equal to the cost of transit service. On the other hand, consider a large provider with its own national backbone. For such a provider, much of their traffic might remain within the network or exit by revenue peering links, but the internal cost of their network would not be that much different from a typical transit provider, since their network would be similar in

structure to a transit provider. We know that the internal costs of a transit provider are capped by the price they charge their customers, which we have pegged at \$.06-\$.18/GB. Assuming that an access ISP has costs at the high end of this range (but allowing for some profit), we could accept that the usage-based internal costs might be \$.15/GB. For those ISPs that serve rural areas, the price might be higher, as I discuss below.

Of all of our cost estimates, this internal cost estimate is the most debatable, because of the question of how many internal costs should be allocated against usage. One ISP may own their own facilities (e.g. fibers) and view them as a sunk cost. Another ISP may rent all of its capacity and take the view that all of its facilities-based costs should be allocated to usage. At the margin, and on the short term, incremental usage adds no cost, but over the longer term, capacity is sized to match usage. If there are massive increases in usage, an ISP may have to re-engineer its internal infrastructure. For this reason, I include internal costs as part of the estimate of usage-related costs. So adding this fact to the picture, we have a final set of estimates as follows:

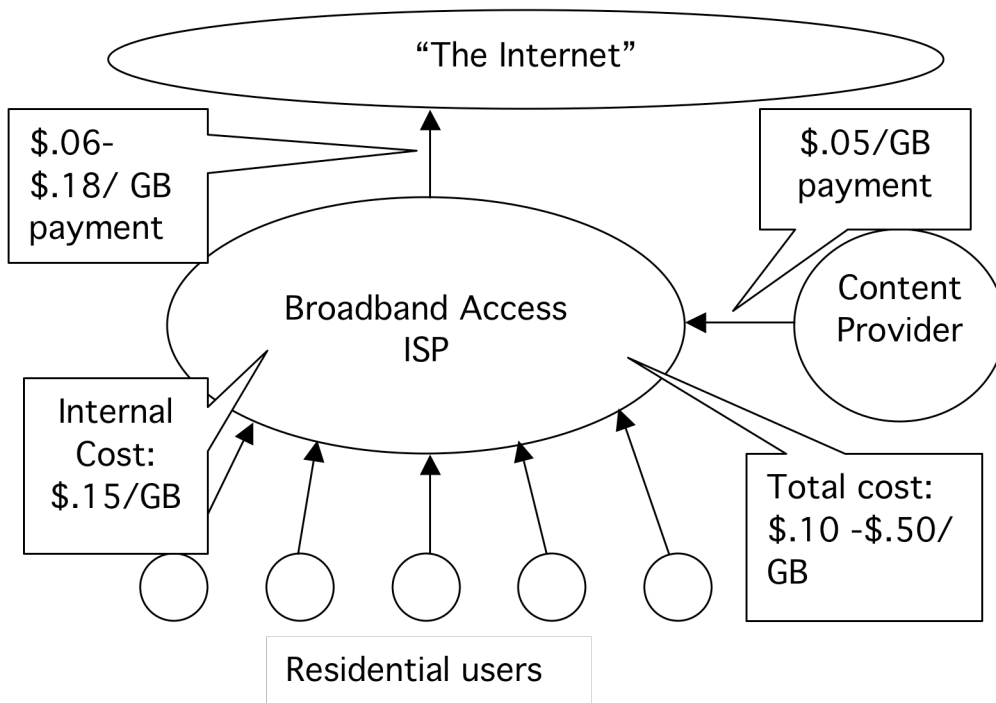


Figure 2: Payment structure for Internet broadband usage

Taking the internal cost as \$.15/GB, to which might be added the transit cost (in some cases) or to which might be subtracted the payment from the content source, we are at least in the same ballpark as the estimate of total cost of \$.10/GB from Burstein.

Implications

The first (and obvious) conclusion is that video downloads will shift the landscape. We have estimates of typical (pre-video) consumer usage as between 2 and 5 GB/month. A single DVD holds 4.7 GB, so the download of a full DVD (most movies do not fill the DVD, of course) could double (or more) the typical monthly usage.

If the consumer watches a total of 8 hours of TV a day (across all his televisions, this may be low, and the data rate of the video is 2 mb/s (again, a low estimate), this would mean that on average, the consumer is downloading almost 220 GB, up from 5GB.

If we take the low end of our estimate of the total cost to deliver a GB, which is \$.10, then this usage would imply a real cost to the ISP of \$22/month. So the ISP would want to add at least \$22 to the monthly \$40, for a (rough) total of \$60/month. With our high estimate of \$.50/GB, this would imply a cost of \$110. Adding that to the current price of \$40 would yield a monthly bill of \$150, which is likely not sustainable by the market. So Internet-based video is viable or not (at today's cost points) depending on whether the actual costs are closer to the high end or the low end of our price bracket.

The next conclusion from the picture is that the total cost of delivering traffic to the customer depends on where it comes from. If it comes into the access ISP over a transit link, there is a transit cost of \$.06-\$.18/GB in addition to the internal costs. If it comes from a content delivery server that pays for access, the internal costs are somewhat offset by a payment we estimate at \$.05/GB. If it comes over a settlement-free peering link, there is no usage-based payment in either direction. More importantly, that estimate of \$.15/GB internal cost (whether it is high or low as an estimate) is an average across all downloads. Some transfers might cost more than this, if they cross expensive facilities. Other downloads might generate essentially no load on internal facilities, if they originate at a point close the user—perhaps a cache at the local head-end. With all respect to Judith Cairncross⁸, distance did not die; it is not even wounded. It is just sleeping, waiting for video to come along. Video that comes from “the other side of the world” might cost that high-end \$.50 per GB, while video that is hosted (e.g. cached) right in the head-end would cost essentially nothing to deliver.

It is estimates such as that \$.50/GB that give rise to dire predictions about the cost of watching a movie. Clearly, the above is an absolute worst-case analysis, since it probably assumes that the content came into the access provider from outside over a transit link. The point of local caching, replicated content sources and CDNs, P2P delivery, broadcast to home-based video recorders and so on, is to find ways to deliver these bits without placing as much demand on the expensive assets of the access ISP, in particular the long-haul and interconnection circuits. Since we have assumed that the incremental cost of increased usage across the access links is very low or essentially nothing, if the content

⁸ Cairncross, F. (2001). The Death of Distance: How the Communications Revolution is Changing Our Lives. Cambridge, MA, Harvard Business School Press.

can be staged “close enough” the total cost might a few pennies per GB, or even a few pennies for a movie.

The rural penalty

As I noted above, the costs for providers in rural areas may be higher. Since I have ignored the cost of outside plant to this point, I will continue to do so, although one should note that the cost of residential access networks in rural areas will often be higher due to less dense population. Some rural networks use wireless technology, which has a different cost structure and also different capacity issues—video of the sort discussed here may not be practical with today’s wireless technology.

The cost of transit service—the connection to the Internet—is the other major cost differentiator for many rural networks. Many rural networks are small, so their internal costs for circuits (aside from the residential access network itself) may be low. But essentially all of their customer traffic will cross the circuit connecting them to the Internet, since few content providers will ask to make a direct connection to a small rural network, and for small ISPs, only a small fraction of the traffic will remain local within the customer pool of the ISP. Transit connections to rural areas are often more costly than in urban areas, because of the length of the circuit to reach the rural head end. Again, precise numbers are hard to pin down and greatly depend on the specifics of the situation, but conversations with rural providers suggest that the cost per mb/s may be twice that of our average \$30/month, or even higher. If we take \$60/m for a mb/s, then we double all the estimates we computed before, so that a GB would now cost \$.36. Small networks (whether rural or urban) will also benefit less from local caching, because the odds of a specific piece of content being retrieved many times is reduced if the pool of customers is smaller. So there may be few options to reduce that \$.36/GB estimate, compared to the options of a large-scale urban ISP.

What this discussion suggests is that both capital and operating costs related to usage will be higher for rural providers, and these costs will end up being passed on to their customers.

An aside: is Internet delivery more costly?

As I discussed above, one of the criticisms that is sometimes launched at the cable industry (in particular) is that they allocate very little of their cable capacity to Internet, and essentially all of it to entertainment TV. This decision results in congestion, which they then have to manage. The framework we have here can shed some light on that issue. Our model has not assigned any part of the cost of outside plant (e.g. the HFC network for cable) to usage. But on the other hand, adding more Internet capacity to the customer’s access link would not reduce any of the costs identified here. How much access capacity to allocate to Internet (as opposed to entertainment TV) and how to reduce the cost of delivery of Internet content are two, almost separate questions. To the extent they are entangled in practice, the analysis is more complex, but the basic issue of cost does not go away—more capacity in the access links does not reduce the other costs.

More and more of entertainment TV is being delivered today using digital encoding—often the same coding that would be used to deliver video over the Internet. So the demand on the access circuit to deliver video should be more or less the same independent of whether it comes over the Internet or over some other share of the access network. The reason that video is cheaper to deliver using “traditional” cable techniques is that the cable architecture has been engineered to reduce the other costs we have identified. For traditional broadcast TV, the cost of bringing the content to the head-end is shared among many customers, and becomes a negligible component of the cost. For the current video-on-demand offerings, the content is hosted on servers in the head-end, so the data need not be fetched on demand across the infrastructure of the ISP, or across any links from another ISP. So both the internal cost and the inter-ISP cost are well-minimized or zero. There is only the cost of the server. As we noted above, disks today cost \$.20/GB retail. If we take that as the rough price of a server and only ten people total watch the video over its lifetime, then this implies that the storage cost per view is only \$.02/GB, which is low enough to tolerate. For a popular video with more viewings, the storage cost per view is no longer material.

So we could surmise that one reason that the ISP is not increasing the capacity of their access network to carry more Internet traffic is that they do not have a model for how to cover the costs we have identified here, as increased usage drives them up, while with current entertainment video, which is a system engineered to reduce cost, these issues do not arise.

Cost reduction through cooperation

The future that emerges will very much depend on whether the ISPs and content providers see each other as partners or adversaries in this space. If they view themselves as adversaries, if ISPs do not facilitate the delivery of video but complain when it is downloaded over expensive transit connections, the resulting costs may not be sustainable. If the ISPs and the content providers cooperate, the cost might be reduced to negligible levels. The cooperation might take the form of caching, support for peer to peer delivery (which, if properly managed, delivers the content from a peer that is close by, so as to make minimal demands on the circuits), broadcast to the home and pre-staging (the DVR of the future), and so on.

This cooperation will include, if only implicitly, a rebalancing of how much consumers and content providers contribute to cover this cost. If the current consumer price point of \$40/month is sustained, it will only be because the content providers are contributing to the cost of delivery. This payment may be in the form of fees to connect their servers to the network (as they are already paying), for storage in caches, or fees to upload the content into the network, or they may take other forms. The costs (if the system is designed properly) might only be a few pennies per movie. But the costs will not be zero, and will have to be covered somehow, either by some increase (perhaps modest) in the monthly charge paid by the residential user or by the content producer.

Two-sided markets

There has been recent interest, both theoretical⁹ and more pragmatic¹⁰, in whether the framework of two-sided markets might be applicable to the situation of the Internet broadband access provider, who attaches on one side to consumers, and on the other side (directly or indirectly) to sources of content. The situation as I have characterized it is a two-sided market structurally, in that the consumer and the content are co-dependent. So there is a valid question as to how these two sides should contribute to covering the costs of delivering the content to the consumer.

Our cost model here gives some basis to determine if there is a practical opportunity to benefit all the parties by shifting costs between the content provider and the consumer.

What would happen if we shifted today's balance of payment between the content provider and the consumer? Our model above had content providers contributing perhaps \$.05/GB to delivery of content, where the total cost was between \$.10 and \$.50/GB. Imagine that the content provider covers another \$.10/GB. If the consumer is using 5 GB/month, his bill would go down by \$.50, or (if he had initially been paying \$40/month), a little more than 1%. The bill to the content provider would triple. It is not necessary to ask detailed questions about user utility curves to speculate that in this version of the story, both sides are *not* better off. The consumers would hardly notice, the content providers would be screaming. The reason (obviously) is that almost none of the consumer's payment goes to cover usage, so there is a tremendous starting asymmetry in the amount that the two sides are paying.

But the situation changes dramatically if the amount of usage goes up. Imagine our earlier example where the user is watching 220GB/month of video, and further imagine that this video is being delivered in a cost-reduced way, so that the cost is \$.05/GB. This implies a monthly charge for usage of \$11, which has to come from either the consumer or the content producer. It might well be that the content producer would cover some of this cost (I estimated that they pay this much today to attach a content delivery server to the network) as part of facilitating delivery of advertising content to the consumer. Morgan Stanley reports that annual Internet spending on advertising is \$21B¹¹, or \$288/household. (I believe this to be on-line households.) This compares to \$40/month

⁹ Economides, N. (2008). "Net Neutrality," Non-Discrimination and Digital Distribution of Content Through the Internet." [I/S: A Journal of Law and Policy for the Information Society \(to appear\)](#).

¹⁰ Houle, J. D., K. K. Ramakrishnan, et al. (2007). The Evolving Internet - Traffic, Engineering, and Roles. [TPRC \(Telecommunications Policy Research Conference\)](#). Arlington, VA.

¹¹ See Meeker, M., D. Joseph, et al. (2007). Technology / Internet Trends. [Future of Media Conference](#). New York.
<http://www.morganstanley.com/institutional/techresearch/pdfs/Future-of-Media-110807.pdf>

(or \$480/year) for residential cable Internet access. Given the relative amounts of the revenues from consumers and from advertisers, it seems reasonable to at least ask what the best balance of payments is to deliver advertising-subsidized content.

Summary

Overall, what this analysis implies is that when the Internet moves to the massive carriage of video, the assumptions about usage, and the payment for usage, will have to be rethought and renegotiated. We will have to “adjust” the user to accept that he cannot watch as much video as he wants for \$40/month, unless the providers of the content contribute to cover the delivery costs, and the ISPs and the providers cooperate to engineer a system that reduces the cost of that delivery to an acceptable level. Otherwise, users will see a pricing model that (for example) has usage tiers, or perhaps even “peak period” usage tiers. However, the cost of usage, if the system is well-designed, should not be so expensive as to destroy the market. If total cost were in the range of \$.05/GB, one could watch TV all month for perhaps \$10 additional to the ISP. ISPs and content providers will have to cooperate to make joint decisions to drive down the cost of video delivery.

Some further conclusions and observations:

- Internet usage is not likely to be a high-margin product. If actual usage costs \$.05/GB, it is hard to imagine that the ISPs will be able to mark this up to \$.10/GB.
- All these prices are from 2007-2008. The video flood is not yet upon us—we now see just the first hints. We can expect that the full shift to Internet video will take several years, and we can expect the cost of usage to decline naturally over that time.

Acknowledgement

I appreciate the comments and suggestions from Steve Bauer and William Lehr on this version of the paper.