THE BROADBAND INCENTIVE PROBLEM, Part II

a follow-on white paper prepared by the

Broadband Working Group

MIT Communications Futures Program (CFP)

Deleted: Cambridge University Communications Research Network¶

Participating CFP Companies

BT
Cisco
Comcast
Deutsche Telekom / T-Mobile
Intel
Deleted: France Telecom ¶
Motorola
Nokia
Nortel
Orange
Telecom Italia

*** DRAFT OF JULY 19, 2006 ***

The publication of the Broadband Working Group's first white paper on "The Broadband Incentive Problem" stimulated healthy debate among members of the press, industry, and advocacy communities. The most frequently asked questions were "On what data are you basing your arguments?" and "How does this paper relate to the network neutrality debate?" In this brief follow-on paper, we address these and other questions about the original paper, and update the discussion based on newer data that has since become available.

In September 2005, the Broadband Working Group issued a white paper entitled "The Broadband Incentive Problem."[1] Since then, particularly after Congressional hearings on "network neutrality" focused the public's attention on related issues, we have received numerous queries regarding our paper. In this short follow-on piece, we answer frequently asked questions about the original paper, and update the discussion based on newer data that has since become available.

What was the paper really trying to say?

Put bluntly: flat revenues + rising costs would not add up to a healthy, sustainable broadband industry.

Put more grandly: if broadband access providers (both wired and wireless) experience costs that rise faster than revenues, they will be unlikely to invest in the ongoing stream of network capacity upgrades that would be required for broadband to follow a Moore's Law-style curve. If this happens, industries in other parts of the broadband value chain, such as content and consumer devices, will also suffer.

Is this scenario likely to happen? Why might costs rise faster than revenues?

The future can never be known with certainty, and we were not trying to claim that this scenario would definitely come to pass. However, we did want to make the point that **the scenario in which broadband access providers' costs rise faster than their revenues should not be categorically dismissed.**

Formatted: Font: Bold

In particular, this problematic scenario would result from the confluence of the following five trends. The first three trends would limit revenue growth, while the fourth and fifth would raise costs:

Deleted: would result in this problematic scenario:

1. Penetration growth slows as broadband access markets saturate;

2. Network operators, both wired and wireless, continue to offer access pricing plans that provide unlimited data network usage (even if tiered by maximum burst rate) for a fixed monthly fee, otherwise known as "flat-rate" pricing;

3. Network operators are unable to capture as much "value-added services" revenue as their current business plans would appear to expect;

4. The average bandwidth used by each customer rises, perhaps significantly, because of changes to user capability (such as through fatter access pipes) and user habits (such as through adoption of peer-to-peer applications for video distribution); and finally,

[1] The original white paper is available at http://cfp.mit.edu/groups/broadband/docs/2005/Incentive_Whitepaper_09-28-05.pdf

Deleted: May 25

5. Bandwidth costs do not decline fast enough to keep usage-based costs from rising.

What data can shed light on the likelihood of each of these trends?

Deleted: The first three trends would limit revenue growth, while the fourth and fifth would raise costs.¶
¶

We view the instigation of public and private discussions around this question as one of the successful outcomes of the original white paper. These discussions particularly highlight the dearth of credible data in discussions (and indeed in formulation) of broadband policy. By definition, no one has more than speculative data about the future, and extrapolating from current trends has well-known dangers. That said, data about current trends is better than no data at all. Here is what we know about each of the five potential trends outlined above.

Trend 1: Penetration growth slows

Market saturation is by nature a matter of when, not if. The timing of this development will differ by region and technology (for example, wireless broadband markets in Finland will probably saturate in a different year from fiber-to-the-home markets in Japan; and, as the discussion in the original white paper showed, DSL markets have already saturated in S. Korea). As a result, there is no universal answer to the question of "how soon is the problematic scenario likely to appear," although interested readers can certainly do their own extrapolations from broadband penetration growth curves available from market researchers.

Trend 2: Flat-rate access persists

In developing the original white paper, the group conducted a non-exhaustive review of broadband pricing policies around the world, as of winter 2005. This review found significant variety and experimentation, including ISPs in the UK and Australia that charged subscribers based on usage volume; ISPs in Iceland and Portugal that placed explicit limits on usage, with more stringent limits for international traffic; and an ISP in Austria that limited usage only during peak hours. However, our review also found that in the U.S. and Korea – the largest broadband markets at that time by number of subscribers and penetration, respectively – the dominant (though not universal) pricing model was a fixed monthly fee, often tiered by peak access rate, with no explicitly stated limits on any individual user's overall traffic volume. Limits were being implicitly enforced, however, by some providers; not surprisingly, these limitations were most often evident in networks with the most limited (and therefore the most contended-for) resources, such as satellite broadband. This finding is even more true today, as cellular broadband networks have grown in popularity. As [4] documents, U.S. cellular broadband providers with all-you-can-eat pricing are enforcing usage volume limitations in practice, while others are experimenting with usage-based pricing.

Deleted: (

Deleted:)

We understand that these limited data points do not prove that flat rate access will persist. In fact, a key implication of our original paper was that **recognition of the problematic scenario outlined above would actually drive providers away from flat-rate pricing for access.** However, for that to happen, providers must find alternative pricing

Formatted: Font: Bold

Deleted: May 25

mechanisms that can succeed in the marketplace. One of the intentions of the white paper was to stimulate thinking and experimentation in this regard.

Trend 3: Provider revenues from value-added services are limited

The original white paper discussed several reasons why provider revenues from value-added services (i.e. services beyond basic access) might have a more limited upside than many operators would appear to be expecting. Most importantly, not all of the activities that generate value for users will have a "service" component that generates revenues for operators; for example, consider peer-to-peer file sharing applications, which are increasingly being employed for legally obtained content [7]. We also argued that value-added services offered by providers face competition from third parties, and that aggressive attempts by providers to favor their own services would be kept in check by the threat (if not the actuality) of regulatory intervention. This last observation led many readers to ask whether our paper intended to take a position in the network neutrality debates currently raging in Washington, DC.

The short answer is no. However, the discussion in the original paper does provide a framework for analyzing some of the arguments swirling through this debate. Third-party applications or services that consume significant network resources, such as the bandwidth consumed by always-on file sharing, can impose noticeable costs on networks; these costs will need to be recovered from somewhere. As a group, we did not take a position as to where such costs should be recovered from; however, **one of the intentions of the original white paper was to stimulate research and industry experimentation on pricing schemes that would enable such cost recovery directly from users, in market-acceptable ways**. We viewed this alternative as preferable to the common practice of automatically throttling such traffic, without giving users a chance to pay more to get more.

Trend 4: Per-subscriber usage volumes increase, possibly significantly

We focus on per-subscriber usage volumes because of their connection to revenues. While multiple sources have documented the aggregate growth rates of Internet traffic in different places and periods (for example see [8], [10], and [11]), some portion of this growth is driven by the addition of new users and therefore involves the generation of new revenues. The problematic scenario of a cost-revenue growth mismatch arises when existing subscribers raise their usage, without adding new revenues into the system (i.e. assuming flat-rate pricing).

For the original white paper, our understanding of this possibility was informed primarily by public and private data about Korea Telecom's experience, in particular data from KT's 2004 annual report on growth in the number of users, and data from [6] on aggregate traffic volume growth. S. Korea provides a particularly interesting harbinger of broadband trends because its fixed-line broadband market is already saturated. A graph in the original paper illustrated rapid growth in KT's traffic per subscriber. Table 1 complements that presentation by showing the annual percentage growth rates for KT's traffic per subscriber.

Deleted: can

Deleted: increasingly more forms of

Formatted: Font: Bold

Deleted: currently

Deleted: simply

Deleted: In a similar vein, third-party applications and services that do *not* consume significant bandwidth, such as VoIP, are far less likely to impose noticeable usage-related costs on networks. Proposals to charge additional fees to support low-bandwidth activities should be evaluated with care, as they are more likely motivated by the desire to protect a provider's own value-added services revenue, than by cost considerations.¶
In sum, we expect that additional degrees of differentiation in Internet pricing are likely to evolve. This is not necessarily a bad thing, as long as differentiated pricing is clearly based on cost considerations.¶

Deleted: May 25

Compounded, the per-subscriber traffic volume more than quadrupled over the 3-year period, 2001-4:

Table 1: Growth in KT's traffic per subscriber

Period	Growth rate
2001-2	42%
2002-3	75%
2003-4	64%
Total (2001-4)	307%

Deleted: ¶

We were also aware of preliminary results of a novel and rigorous study being carried out in Japan, involving seven Japanese ISPs, several Japanese Internet researchers, and the Ministry of Internal Affairs and Communications, and focusing on the characteristics of residential broadband traffic [1]. Subsequent to the publication of our original white paper, an expanded and updated version of this group's findings became available [2]. This updated version is particularly interesting because it breaks out differences in usage among DSL and FTTH users, possible because by September 2005, Japan had 4 million FTTH users, up by more than a factor of two from the year prior. The results of this study are not only quite interesting in their own right, but consistent with ideas we discussed in the original white paper regarding changes in user behavior and traffic variance that seemed likely as access pipes get fatter.

Specifically, the Japanese researchers found that aggregate traffic increased 45% from November 2004 to November 2005 for the 7 ISPs they observed.[2] Normalizing this increase by the growth rate in total Japanese broadband users over roughly the same period (from about 12m in September 2004 to 14m in September 2005), we estimate that traffic per subscriber grew 24%.[3] This figure is substantially lower than the Korean increases in the 2001-4 period (Table 1), illustrating an expected variability in per-subscriber traffic growth rates across different markets and time periods.

The study's deeper analysis of particular segments within the data provide numerous interesting insights. First, about 70% of the traffic appears constant within the pattern of diurnal traffic fluctuations, suggesting that machine-generated traffic is a significant component of overall usage as well as a significant driver of overall traffic growth. Second, the data clearly demonstrate that subscribers use bandwidth more symmetrically when they have the option. That is, subscribers to symmetrical FTTH services produced more symmetrical traffic patterns, while subscribers to asymmetric DSL services maxed

[2] Participants in the Broadband Working Group confirm a similar growth rate for traffic measured at an Internet exchange during 2004-5. Similarly, [5] estimates 50% annual increases in U.S. Internet traffic.

[3] This assumes that the growth rate of subscribers to the 7 ISPs is not substantially different from the country as a whole. Unfortunately, the growth rate in subscribers for these 7 ISPs is not reported in the paper.

Deleted: May 25

out their upstream traffic capability.[4] Third, rural subscribers do not behave differently from urban ones; lower traffic volumes in rural regions simply reflect their smaller overall number of users, not the nature of usage within each region, which is markedly similar.

Of particular interest is the study's examination of so-called "heavy hitters," defined as users who upload (i.e. send to their ISP) more than 2.5GB per day.[5] The study found that heavy hitters' traffic was more symmetric, and that the top 4% of heavy hitters used 75% of the total upstream traffic, and 60% of the downstream. However, the study also demonstrated that despite the visible change in slope of the cumulative distribution curves, there were no visible gaps in usage patterns. The distribution of usage volumes among users is heavy-tailed, with no clean delineation point between "light" and "heavy" users.[6] As we observed in our original white paper, such a distribution makes the design of tiered or capped volume-pricing schemes particularly difficult.

The study's authors interpret these and other aspects of their data to mean that there is no clear application profile for particular types of users; rather, their results reflect a diverse portfolio of application types (e.g. streaming or downloading, vs. file-sharing) within each user's activities. In their words ([2], p. 12): "[H]igh-volume traffic is generated not only by peer-to-peer file sharing but also by other applications such as content-downloading from a single server. A plausible explanation for the large variance is that the majority of users use both file-sharing and downloading with different ratios." Again, this result reinforces the supposition in our earlier white paper that the variance among users might well increase as access pipes get fatter, making one-size-fits-all pricing that much more difficult to sustain.

Finally, the study found that there were proportionally more heavy hitters among FTTH users than among DSL users. Unfortunately, the study methodology does not enable the authors to distinguish whether the heaviest DSL users are simply the first to migrate to fiber, or whether fiber users become heavy hitters because they can. As the authors write:

> We can no longer view heavy-hitters as exceptional extremes since there are too many of them, and they are statistically distributed over a wide range. It is more natural to think they are casual users who start playing

---

[4] We use "upstream" to refer to data flowing from the user to the ISP; [2] however adopts the ISP's perspective and refers to such traffic as "inbound." As the authors write (p. 13): "The inbound and outbound rates are roughly equal throughout our data sets. Many access technologies employ asymmetric line speed for inbound and outbound based on the assumption that content-downloading is dominant for normal users. However, this assumption does not hold in our measurements."

[5] 2.5 GB per *day* is selected because it is the (rough) knee of the log-log cumulative distribution curves. See especially Figures 15 and 16 in [2]. We note that some provider experimentation with usage caps has involved cutoff values as low as 10 GB per *month*, which may partially explain their unpopularity with users.

[6] In their words: "Note that the difference is only in the slope of the distribution, and the boundary between the two groups is not clear. In other words, users are distributed statistically over a wide traffic volume range, even up to the most extreme heavy-hitters. There is no typical daily traffic volume per user that can be identified by a concave in the slope." See [2] p. 8.

Formatted: Font: Italic

Deleted: May 25

> with new applications such as video-downloading and peer-to-peer file sharing, become heavy-hitters, and eventually shift from DSL to fiber. Or, sometimes users subscribe to fiber first, and then, look for applications to use the abundant bandwidth. The implication is that, if a new attractive application emerges, a drastic change could occur in traffic usage. ([2], p. 13)

This last point is especially worth noting in light of the fact that the huge leaps in usage occasioned by the original Napster's popularity were largely unforeseen by ISPs. Of course, no one can predict with certainty whether some new application or service involving video over the Internet will take ISPs by storm in the same way, but given the ferment in the video-over-IP space at the moment, and recent announcements such as Warner Bros.' use of Bittorrent's peer-to-peer video distribution technology, it is not improbable to imagine scenarios in which per-user traffic volumes grow rapidly.

Trend 5: Aggregate cost of bandwidth rises

This speculation generated a good deal of controversy, particularly from readers who argued that providers will simply be able to "overprovision" themselves out of rising usage scenarios. Overprovisioning was one of the commonly proposed solutions described in the original paper. Its advocates contend that (1) bandwidth costs are declining so fast that usage would have to grow much more than anyone expects in order to cause a problem, and (2) even if that did happen, bandwidth costs are starting from a negligible proportion of major ISP's expenses in any case.

We have no public data with which to confirm or deny contention (2), so we do not discuss it further here, other than to note that the scale and location of an ISP clearly affect the validity of this assumption.[7] Contention (1), however, we were able to examine with some data. The results are mixed.

We began by using public sources to estimate a 35% annual price decline for transit bandwidth prices during the 1999-2005 period. Table 2 lists the estimated transit prices available at major Internet exchange points for a 1 Gbps commitment, and the sources of these estimates.[8] The rate of price decline is a better way to characterize bandwidth price trends than the absolute price quotes shown in Table 2, however. Absolute prices are known to vary considerably based on the size of the circuit purchased, as well as by region (see [10] and [12]). In addition, the prices shown in Table 2 represent marginal prices available at bandwidth exchanges; a provider seeking to balance costs and revenues, however, would concern itself with average costs, which are typically higher.

Deleted: , and t

Deleted: as noted to construct

Deleted: , showing

Deleted:

Deleted: :

[7] [*** consider inserting a reference to latest DSL Prime Newsletter once Burstein puts it out…he makes strong claim that bandwidth is small fraction of operator's costs. ***]

[8] The 2002-3 data points are inferred based on the price decline computed from the 1999 and 2005 endpoints. Put another way, we solve 20.73 = 281.01*(1-r)^6, and find r = 35.2%.

Deleted: May 25

Table 2: Price Declines for Transit Bandwidth

Year	Price ($/Mbps/mo based on purchase of 1 Gbps)
1999	$281.01
2000	$182.09
2001	$118.00
2002*	$76.46
2003*	$49.55
2004	$32.11
2005	$20.73

*inferred, see note 8

Sources:

1999-2001: [3], Figure 7 (p. 49);

2004: http://www.merit.edu/mail.archives/nanog/2004-08/msg00178.html;

2005: http://www.nanog.org/mtg-0505/norton.html

If bandwidth prices are declining 35% annually, how fast can bandwidth usage grow each year such that a broadband providers' overall expenditures on bandwidth remain constant? The answer, 54%, can be seen in Figure 1, which illustrates the relevant computation for bandwidth growth rates from 5 to 100%, and price decline rates from 5 to 50%.[9] The frontier represents constant expenditures on bandwidth, while points above or below it represent increasing or decreasing overall expenditures on bandwidth, respectively.

Deleted: ¶

Deleted: ¶

Deleted: We estimate the 2002-3 prices by computing a 35.2% annual rate of price decline based on the known 1999 and 2005 endpoints. Put another way, we solve 20.73 = 281.01*(1-r)^6, and find r = 35.2%. We assume that this rate of price decline is a better way to characterize bandwidth price trends than the absolute price quotes shown in the table, because prices are known to vary considerably based on the size of the circuit purchased, as well as by region. However, it seems reasonable to assume that wherever transit is competitively provided, the rate of price decline would be similar, regardless of the total circuit capacity.¶
If bandwidth prices are declining 35% a year, then bandwidth per user can grow by 54% annually without squeezing broadband operators' margins.

Deleted: 54% is computed by assuming

Deleted: while the price in year 2 declines to 65% of the price in year 1

Deleted: Intuitively

Deleted: ,

Deleted: More precisely,

Deleted: if $P_1*Q_1 = P_2*Q_2$

Formatted: Subscript

Formatted: Subscript

Formatted: Subscript

Formatted: Subscript

Deleted: $= (1-.35)*P_1 * Q_2$, then $Q_2 = (1/.65) *Q_1 = 1.54 * Q_1$.

Deleted: May 25

[9] The computation proceeds from the assumption that the product of price and quantity of bandwidth used (the total bandwidth cost) remains constant from year-to-year. For a 35% annual price decline, the intuition for 54% can be gotten by observing that if the price in year 2 is 2/3 of the price in year 1, then the quantity can increase by 3/2. Algebraically, $Q_2 = (1+r) * Q_1$, where r is the rate of bandwidth growth shown on the x-axis in Figure 1, and $P_1*Q_1 = P_2*Q_2$. Therefore, $P_2/P_1 = Q_1/Q_2 = 1/(1+r)$. For the graph's y-axis, this ratio is expressed in percentage terms by subtracting 1 from the result. The graphical presentation of this computation was inspired by [11], Figure 4, p. 16.

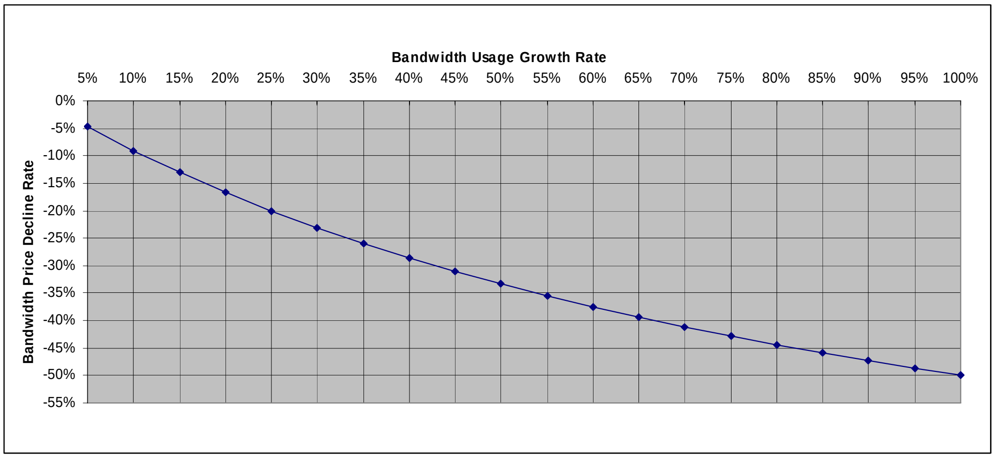

Formatted: Keep with next, Keep lines together

Figure 1: Constant Bandwidth Expenditure Frontier

How does this frontier compare against the per-user traffic growth rates discussed above under Trend 4? Looking back, we see that per-user traffic growth in Japan (24% in 2004-5) and in Korea 2001-2 (42%) fell below the 54% threshold, while later Korean traffic growth per user (75% in 2002-3 and 64% in 2003-4) outstripped it. The year-to-year variability evident in the Korean data is consistent with the interpretations discussed above for the Japanese data; the introduction and adoption of new applications and services is an unpredictable process that leads to correspondingly unpredictable fluctuations in usage levels.

Deleted: at the traffic growth rates described in the previous section

Deleted: this

It would be tempting to conclude that 35% annual price declines provide enough headroom such that broadband operators ought to cross the cost of bandwidth off their list of worries keeping them awake at night. However, such a conclusion would be premature for several reasons. First, even at 35%, some operators have experienced periods of explosive per-subscriber usage growth, putting them above the frontier and thus squeezing their margins. During such periods, the motivation for providers to constrain customer usage is readily apparent. Second, the 35% figure is only one estimate; the rate of price decline can vary significantly, across the full range shown in Figure 1, depending on where and when the bandwidth is used (see [10] and [11] for examples).

Formatted: Left

Most important, evidence is accumulating that at least in the more developed parts of the world, such significant annual price declines are unlikely to continue in the foreseeable future. For example, [10] presents charts suggesting a 2006-7 end to the capacity overhang from the telecom bubble, and argues that this factor, combined with ongoing growth in aggregate demand and industry consolidation, will lead to a closer balance between supply and demand in the future, leading to less steeply declining prices (already evident by 2004; see [10], Figure 2, p. 14). Figure 3 of [11], reproduced below, provides further evidence of the leveling off in prices during 2004-5. This chart also provides a graphic illustration of the geographic differences in bandwidth prices, helping to explain some of the variability

Deleted: May 25

we observed in our global review of ISP pricing structures, and underscoring the point that pricing policies should be expected to vary by locale.

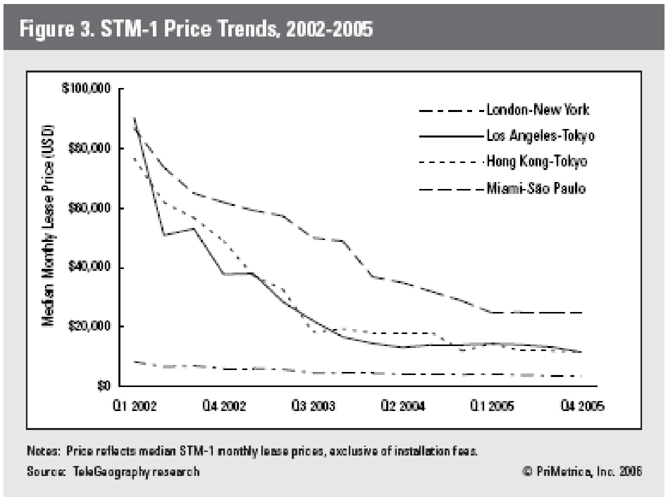

Conclusions

It can be seen from Figure 1 above that if the rate of bandwidth price decline were to slow down significantly, for example by dropping from 35% to 20% or even 10% per year, that the rate of bandwidth growth would have to be cut roughly in half (to 25% and 10%, respectively) for providers to remain on the constant expenditure frontier. With continued growth likely in both subscribership and per-subscriber usage, this outcome seems unlikely. While we are not clairvoyant and cannot predict the future with certainty, if TeleGeography's projections are correct, it would appear highly likely that broadband operators will be spending more on aggregate transit bandwidth in the future, just as they have been in the past (see [11], Figure 4). Whether that increased spending is a problem or not obviously depends on revenue-side factors as well, underscoring the importance of per-subscriber analyses, particularly as a broadband market approaches saturation.

With so much uncertainty surrounding their costs, it should hardly be surprising that broadband providers would seek ways to increase their revenues. As [9] discusses, the motivations for network operators to find ways to successfully price discriminate are very strong, but the track record of success is correspondingly weak. As with our original white paper, our motivation in exploring this topic further has been to stimulate experimentation with a broader range of approaches and experiments, in search of mechanisms that customers find acceptable. The time has come for the broadband industry to think seriously about alternative pricing mechanisms that can remove the perverse incentive for network operators to throttle bandwidth-hungry innovations that they do not "own," but

Formatted: Centered

Deleted: would have to be predicated on the assumption that such significant annual price declines will continue into the foreseeable future. We are not clairvoyant, and cannot predict with certainty whether that will happen. However, the following chart, reproduced from [5], strongly suggests it may not, as long-haul bandwidth providers finally reach the limits of the oversupply engendered by the telecom bubble, leading to more stable bandwidth prices:¶

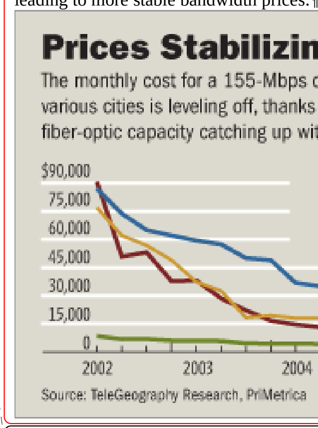

Formatted: Justified

Deleted: The chart also provides a graphic illustration of the geographic differences in bandwidth prices, helping to explain some of the variability we observed in our global review of ISP pricing structures, and underscoring the point that pricing policies should be expected to vary by locale.¶

Deleted: It is clear that flat-rate pricing of consumer broadband access has already led to limitations on the use of bandwidth-intensive applications. Such limits threaten innovation all along the broadband value chain.

Deleted: May 25

that are nevertheless an essential component of the broadband Internet's appeal and value to users.

SOURCES

[1] Fukuda, Kensuke, Kenjiro Cho & Hiroshi Esaki, "The Impact of Residential Broadband Traffic on Japanese ISP Backbones," ACM SIGCOMM Computer Communications 35:1, January 2005. Available via "Related Papers" section of Broadband Working Group page, http://cfp.mit.edu/groups/broadband/broadband.html

[2] Cho, Kenjiro, Kensuke Fukuda, Hiroshi Esaki, and Akira Kato, "The Impact and Implications of the Growth in Residential User-to-User Traffic." Available via "Related Papers" section of Broadband Working Group page, http://cfp.mit.edu/groups/broadband/broadband.html

[3] Ferreira, Pedro and Marvin Sirbu, "The Price of Anarchy for Interconnected Communication Networks Provisioned Selfishly," available at http://bear.cba.ufl.edu/centers/purc/documents/Pedro_Ferreira_14May04.pdf

[4] Sharma, Amol and Dionne Searcey, "Cell Carriers to Web Customers: Use Us, but Not Too Much," *Wall Street Journal*, May 11, 2006, B1.

[5] Heinzl, Mark and Shawn Young, "With Rising Internet Traffic, Spare Fiber-Optic Lines Fill Up," *Wall Street Journal*, April 27, 2006, B1.

[6] Kim, Tae-gyu, "KT Seeks Usage-Based Internet Pricing," *Korean Times*, March 28, 2005.

[7] Sullivan, Mark, "AOL Goes P2P for Video," *Light Reading*, February 3, 2006, available at http://www.lightreading.com/document.asp?doc_id=88118&WT.svl=news1_4

[8] Odlyzko, Andrew, "Internet traffic growth: Sources and implications," in *Optical Transmission Systems and Equipment for WDM Networking II*, B. B. Dingel, W. Weiershausen, A. K. Dutta, and K.-I. Sato, eds., Proc. SPIE, vol. 5247, 2003, pp. 1-15. Available at http://www.dtc.umn.edu/~odlyzko/doc/itcom.internet.growth.pdf

[9] Odlyzko, Andrew, "Pricing and Architecture of the Internet: Historical Perspectives from Telecommunications and Transportation," available at http://www.dtc.umn.edu/~odlyzko/doc/pricing.architecture.pdf

[10] TeleGeography, *International Bandwidth 2005*, Primetrica Inc. (Executive Summary available via free registration at http://www.telegeography.com/ee/supplemental_files/pdf.php?pdf_file=ib2005_exec_sum.pdf&pub_code=free_resources)

[11] TeleGeography, *Global Bandwidth 2006*, Primetrica, Inc. (Executive Summary available via free registration at http://www.telegeography.com/ee/supplemental_files/pdf.php?pdf_file=gb2006_exec_sum.pdf&pub_code=free_resources)

Formatted: Justified, Space Before: 6 pt, After: 6 pt

Deleted: ¶

Formatted: Space Before: 6 pt, After: 6 pt

Deleted: ¶

Formatted: Font: Not Bold

Formatted: Font: (Default) Times New Roman, 12 pt

Formatted: Font: (Default) Times New Roman, 12 pt

Formatted: Font: (Default) Times New Roman, 12 pt

Formatted: Space Before: 6 pt, After: 6 pt

Deleted: May 25

[12] Norton, B., "Transit Cost Survey: Data Collected at NANOG 36," Available at http://www.nanog.org/mtg-0606/pdf/bill.norton.2.pdf

Formatted: Font: Not Italic

ACKNOWLEDGMENTS

The ideas in this white paper reflect discussions among participants in the Broadband Working Group (listed on the CFP web site), as well as feedback on the original white paper received from Dave Burstein, Harold Feld, Jessica Rosenworcel, Simon Wilkie, Cathy Yang, Jeffrey Sterling, and Bob Frankston. We thank them for their stimulating questions and insightful comments.

Opinions expressed in this paper are drawn from consensus views among the working group's participants, and do not represent official views or policies of CFP's sponsoring companies or universities.

Deleted: <insert>, plus¶

ABOUT THE MIT COMMUNICATIONS FUTURES PROGRAM

The Communications Futures Program (CFP) is a partnership between university and industry at the forefront of defining the roadmap for communications and its impact on adjacent industries. CFP's mission is to help our industry partners recognize the opportunities and threats from these changes by understanding the drivers and pace of change, building technologies that create discontinuous innovation and building the enablers for such innovation to be meaningful to our partners. Further information about CFP and the Broadband Working Group is available at http://cfp.mit.edu.

Deleted: <insert>¶

Deleted: May 25