THE BROADBAND INCENTIVE PROBLEM

a white paper prepared by the Broadband Working Group

MIT Communications Futures Program (CFP)

Cambridge University Communications Research Network

September, 2005

THE BROADBAND INCENTIVE PROBLEM

a white paper prepared by the Broadband Working Group
MIT Communications Futures Program (CFP)
Cambridge University Communications Research Network

September, 2005

Participating CFP Companies

- British Telecom
- Cisco
- Comcast
- Deutsche Telekom / T-Mobile
- France Telecom
- Intel
- Motorola
- Nokia
- Nortel

The broadband value chain is headed for a train wreck. Any business that expects to reach its customers or employees through ever-better mass-market broadband Internet access, whether wired or wireless, is in for a rude awakening. Unless the broadband incentive problem is recognized and dealt with now, it is unrealistic to assume that Moore's Law-style improvements in broadband access will be available as a growth engine for companies that produce networked content, devices, applications and services, and as a productivity engine for mobile and home-based workers throughout the global economy.

The broadband locomotive left the station with a critical missing piece: the incentive for network operators to support many of the bandwidth-intensive innovations planned by upstream industries and users. The intent of this white paper is not to prescribe any particular solution to this problem, but to ensure that stakeholders across the value chain recognize now that the broadband train is headed for a crash, and are motivated to redirect it — each in ways appropriate to their situation — before it is too late. This motivation is especially important in markets where the locomotive is just beginning to puff, including wireless broadband in most regions, and any form of broadband in less economically developed regions. The earlier the broadband incentive problem is recognized, the easier it will be to avoid in the future.

Executive Summary

The future of broadband faces a crisis: an incentive problem derailing the ability of mass-market Internet users to take advantage of Moore's Law. Today's prevailing business models give wired and wireless broadband operators the perverse incentive to throttle innovative, high-bandwidth uses of the Internet. If this problem is not addressed now, many commonly foreseen broadband developments are unlikely to happen as planned. These include the next generations of videoconferencing, interactive video and television (broadly defined), collaborative gaming, peer-to-peer applications, grid-oriented computing, network-based backups, data-capable wireless networks (3G and beyond) and the sophisticated portable networked gadgets that will use them, and fiber-to-the-home networks. Delays in these innovations will hurt the makers and users of networks and all of their upstream complements, including content, applications, services, and devices (Figure 1).

The incentive problem is already evident in leading broadband markets. Popular flat-fee pricing models have encouraged penetration but also led innovative users to adopt bandwidth-intensive behaviors that impose additional costs on network operators, an especially noticeable problem once penetration saturates and revenue growth flattens. Leading network operators have considered or imposed restrictions on user behavior, employing a range of schemes that vary in sophistication. Most simply focus on limiting user traffic, while a few also seek to monetize additional usage, typically in coarse ways that may bear little relation to actual usage costs imposed.

Effective solutions to the broadband incentive problem are not obvious. Today's responses will prove inadequate as broadband markets expand. In our observation, operators have not yet found access pricing mechanisms that both make sense to users and effectively align user behaviors with the costs they impose. Overly broad limitations on user behavior will be unpopular with users and, by unduly curtailing the activities that motivate users to pay for broadband in the first place, will ultimately prove unsatisfactory to providers as well (Figure 2). Many operators have also proposed to respond to rising usage-based costs by extracting additional revenue from value-added services beyond basic access, such as voice-over-IP and IP-based television. We do not believe this response is adequate to solve the problem, for two reasons. First, we expect that operators' revenues from value-added services will be insufficient to cover rising usage costs, because service revenues will be limited by competition from a growing set of third parties, and by legal or regulatory attention to any perceived constraints on such competition. Second, some bandwidth-intensive broadband innovations will not have an associated revenue-generating service. Reliance on value-added service revenues does not give operators an incentive to support this subset of applications, even though the innovations they represent will be valuable to users and upstream industries.

The intent of this white paper is to ensure that stakeholders across the broadband value chain recognize the reality of the incentive problem, and are motivated to deal with it now, before it becomes more difficult to solve. This paper does not propose particular solutions, but rather intends to motivate stakeholders to work now to address the problem, each in ways appropriate to their particular situation and perspective. It is in everyone's interest — network operators, users, upstream value chain participants, and government stewards of our economic well-being — to find solutions that will support the ongoing network investments needed for broadband to follow Moore's Law, and for application innovation to flourish. Early recognition of the incentive problem is especially important in less mature broadband markets, such as wireless broadband in most regions, and any form of broadband in less economically developed regions. The mistakes of the past do not have to be repeated, and the earlier the broadband incentive problem is recognized, the easier it will be to avoid in the future.

Why is the Broadband Incentive Problem Important?

Moore's Law has proven to be an engine for growth in all aspects of mass-market (or, in other words, consumer-oriented) digital technology. As the price-performance of digital devices — from personal computers to point-and-shoot cameras — has improved over time, consumers have been able to take advantage of an advancing array of capabilities, increasing their willingness to pay. The resulting revenue growth has fueled ongoing research and development, and today we observe a virtuous cycle of industry innovation, user adoption, and market growth.

The broadband incentive problem needs to be solved because it threatens to turn today's virtuous cycle into a vicious one. Mass-market digital technologies increasingly expect to communicate over the Internet — a trend that is already evident in the next generation of innovative digital devices, from iPods to Skype-based mobile phones to set-top boxes for IPTV. As the developers of Internet-enabled devices, applications, content, and services take advantage of price-performance improvements in computing power, they will depend on complementary improvements in communications — and in particular access networks, where the bottleneck often lies — to make their products more valuable to customers.

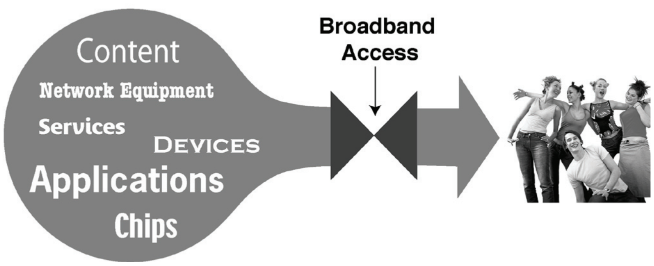

Figure 1: Broadband Value Chain

Ongoing growth in all of the upstream industries shown in Figure 1 will therefore depend on mass-market broadband access — both wired and wireless — being on a price-performance improvement curve compatible with the rest of the broadband ecosystem. While the full range of creative network-dependent innovations is impossible to predict, commonly expected drivers of user and upstream-industry growth include the next generations of videoconferencing, interactive video and television (broadly defined), collaborative gaming, peer-to-peer applications, grid-oriented computing, network-based backups, data-capable wireless networks (3G and beyond) and the sophisticated portable gadgets that use them, and fiber-to-the-home networks. All of these bandwidth-dependent developments are threatened by the broadband incentive problem.

We do not foresee growth in broadband access on the order of Moore's Law unless the incentive problem is recognized and dealt with effectively now. Today's prevailing revenue models for mass-market access give broadband operators a perverse incentive: to throttle many innovative, high-bandwidth uses of the Internet, rather than to invest in the additional network capacity needed to support the next generation of bandwidth-intensive applications. The first part of this white paper explains the origins of this rational but perverse incentive: how user behavior is changing, and why these changes are likely to increase network operators' costs without necessarily increasing their revenues.

Stakeholders who recognize this problem but believe it is easily remedied are not likely to grasp the importance of considering alternative solutions. Consequently, the second part of this white paper discusses why commonly proposed remedies are either ineffective or incomplete solutions to the broadband incentive problem. The discussion in part II is not at attempt at exhaustive examination of proposed solutions. Rather, the intent is to convince stakeholders that solutions must go beyond the obvious candidates discussed in this white paper, if they are to properly align the investment incentives of network operators with the value derived by participants upstream in the value chain.

Timely recognition of the broadband incentive problem is particularly important for stakeholders in less mature broadband markets, including emerging forms of wireless broadband (such as cellular 3G and WiMax), and broadband in less economically developed regions. Stakeholders in these markets have the opportunity to establish sustainable user expectations regarding pricing and usage, at an earlier stage of market development. The incentive problem may be more easily dealt with in newer markets, if stakeholders do not follow the initial "all you can eat" pricing policies employed in leading fixed broadband markets.

Solving the broadband incentive problem is important to the health of the entire broadband value chain, including network operators, users and other participants upstream, and government stewards of our economic well-being. As Figure 2 illustrates, stagnation in what users can do with networks will ultimately lead to stagnation in user demand for the networks themselves. While it may be rational for access providers to under-invest in network capacity in the current environment, more capable networks would enable more

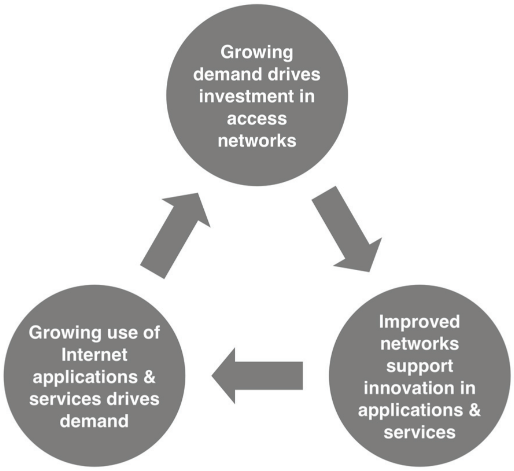

Figure 2: Ideal Broadband Virtuous Cycle

capable applications and services, increasing total industry revenues by raising customers' willingness to pay. Profits for all could be larger if operators' revenue models aligned their incentives for investment with benefits upstream.

PART I: ORIGINS OF THE BROADBAND INCENTIVE PROBLEM

How Is User Behavior Changing?

The roots of the incentive problem lie in changes that are already observable in how customers use always-on, high-speed broadband connectivity. Narrowband (dialup) access constrained user behaviors such that: (a) average usage levels were similar across users, and (b) for any individual user, the difference between average and peak usage rates was not large. As always-on, high-peak-rate broadband access lifts constraints on application and user behavior, the variability in broadband usage is increasing. Customers continue to use applications and services, such as text-based email, that were developed to work well under narrowband constraints; but they also have a widening array of broadband-enabled applications and services to choose from.

As broadband diffuses in the marketplace, a growing number of users can be expected to exploit the capabilities offered by higher peak rates. Some will do so only occasionally, for example to perform software upgrades requiring large downloads. Others will do so more routinely, as they "podcast" their own audio or video productions, watch IP-based TV or other forms of on-demand or streaming media, videoconference via Skype, or participate in other peer-to-peer file sharing or distribution applications, like BitTorrent, that are capable of transferring data at the peak rate of the access link on a more or less continual basis.

How Will Changes in User Behavior Affect Broadband Traffic?

These structural and behavioral changes suggest that broadband traffic will differ from narrowband along several dimensions. First, higher peak rates will increase the variation within any single user's traffic. Second, the dispersion among users will be greater as penetration increases and the customer base comes to mirror the diversity of the general population. Finally, the convenience of "always-on" and the availability of usage-intensive applications will cause the average traffic per user to increase.

Evidence from Korea, where broadband penetration is the highest in the world, suggests that these changes are more than hypothetical (see box, "Lessons From the Korean Broadband Marketplace"). The mean traffic per user is clearly rising rapidly: aggregate network traffic nearly doubled every year since 2001, while the number of subscribers grew at a much slower rate during that period. Additionally, a small fraction of users generates a much larger proportion of the traffic, indicating large dispersion among users.

How Does Increasing Broadband Traffic Raise Operators' Costs?

In essence, broadband networks consist of switching equipment, communication links (whether wired or wireless), and management servers whose dimensions are selected to handle an expected level of traffic. Once traffic grows beyond those expectations, additional investments in these components are required to increase network capacity. KT's large investments in upgrading their backbone capacity (see box, "Lessons From the Korean Broadband Marketplace") illustrate the financial impacts of growth in aggregate broadband usage.

The costs of increased usage can take the form of operational or capital expenses. Additional capital investments are required when design limits are reached in the network components that the access provider owns, such as in the "internal" access and aggregation network shown in Figure 4. Capacity limitations in this portion of the network are particularly relevant in cellular-style wireless broadband infrastructures, where spectrum is not readily available to support additional usage.

Operational expenses arise when access providers lease capacity from others, such as for the communications links for backhaul transit shown in Figure 4. Increased aggregate usage may not necessitate capital investments for long-haul carriers (the backhaul providers) if they operate with plenty of excess capacity. However, the access provider typically

LESSONS FROM THE KOREAN BROADBAND MARKET

Korea Telecom recently announced plans to switch from a policy of charging Internet services on a flat-rate pricing schedule to a usage-based system. This is a watershed event in the development of broadband access over the last decade. Korea has led the world in broadband adoption, and the problems they face may well foretell similar issues in other markets.

KT's experience shows that once a broadband market becomes saturated, revenue growth stagnates, and network infrastructure and operational costs do not decline as rapidly as usage-sensitive costs rise.

- Korea tops the world in residential broadband penetration with 12 million of its 15 million households subscribing.[1] Penetration rates are highest among the young: 2004 penetration rates were 95.3% and 88.1% for people in their 20s and 30s, respectively.[2]
- With the market approaching saturation, broadband subscriber growth rates have flattened out, dropping from 75.3% in 2000, to 11.2% in 2003, to 4.7% in 2004.[2]

While revenue growth has slowed considerably, Internet usage continues to rise quickly (Figure 3). The continual growth in network usage has forced KT to invest repeatedly in expanded network capacity. The aggregate traffic on KT's network has nearly doubled every year since 2001, and KT has had to invest over $150M since 2000 in upgrading their backbone capacity.[1]

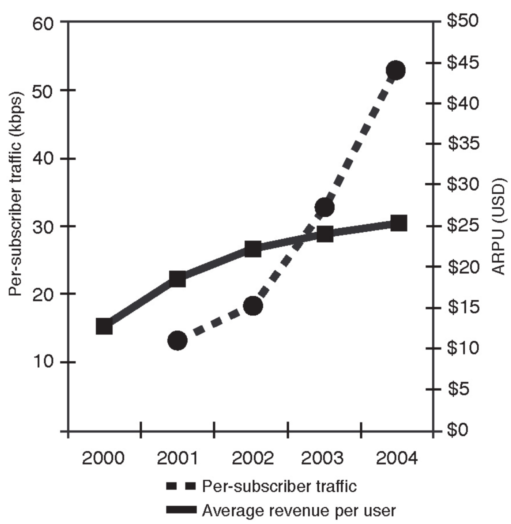

Figure 3: KT Traffic vs. Revenue Growth, Per Subscriber[1,4]

KT's announcement of usage-based pricing has met with stiff public resistance. Even though KT claims that average customers will not see a rate increase (since 5% of users account for nearly half of the total traffic on KT's network[3]), the Korean public has not been receptive to the change. KT has maintained that they will institute the usage-based pricing sometime in the next two years but that it will only affect a small percentage of their customer base.

In short, there are no easy answers for broadband providers, but the problem is clear — increasing usage costs are not economically sustainable when coupled with flat revenues.

[1] Korean Times, "KT Seeks Usage-Based Internet Pricing", March 29, 2005.
[2] Korean Times, "Internet Penetration Rate Tops 70%", Feb. 1, 2005.
[3] Economist Intelligence Unit, "South Korea: Broadband Blues", April 9, 2004.
[4] KT Corp. (2004). 2004 Annual Report. Retrieved August 10, 2004, from KT Corp. web site http://www.kt.co.kr/kthome/eng/ir/fi_reports/ar.jsp

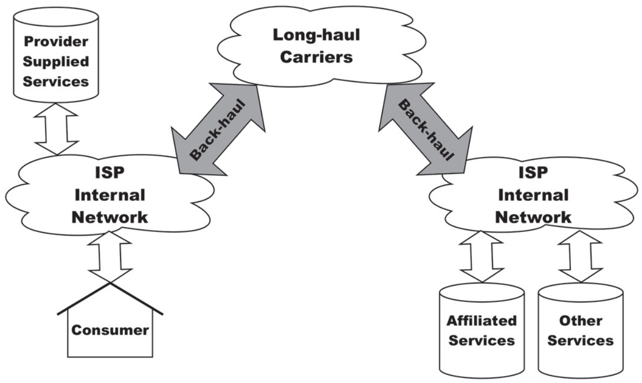

Figure 4: Traffic-Sensitive Operational Costs

only pays to lease a fraction of the long-haul provider's capacity — and as that fraction increases, so does the access provider's operational (lease) expense.

Increased usage can also raise operational expenses under other forms of operator interconnection, beyond the backhaul transit illustrated in Figure 4. Nominally expense-free peering agreements among operators can also include clauses that require inter-provider payments in the presence of significant deviations from balanced traffic flows. Rapid adoption of new applications that change the balance of traffic flow, such as peer-to-peer applications, can thus create operational expenses where none were expected.

What Are Typical Broadband Revenue Models Today?

All broadband operators derive revenues directly from the provision of network access to customers. Some also derive revenues in other ways, including directly from customers for services beyond access, or indirectly from other participants in the broadband value chain.

For access-based revenues, the flat monthly fee is currently the dominant industry model. Customers are typically given a choice of a few (2-3) different recurring fee levels; peak access bit rate, overall monthly traffic volume, or both are the primary factors distinguishing the different pricing tiers. Further details can be found in the Broadband Working Group's informal review of broadband pricing around the world. One key finding of this review is that price tiers based on overall traffic volume are not common in S. Korea (the leading broadband market in terms of penetration) or the U.S. (the largest broadband market measured by absolute number of subscribers). The review also found that among providers who do offer volume-based price tiers, both nominal and actual enforcement can be quite variable. Some providers will simply generate a letter after several months to subscribers who "consistently" exceed their volume tier, while others claim they will charge additional usage-based fees in real time once customers exceed volume thresholds, analogous to fees for extra minutes of cell-phone usage above one's monthly allotment.

Non-access-based broadband revenues come in many forms, and it is not our intent here to catalog them all, but rather to highlight two that we will discuss further below in the context of potential "solutions" to the incentive problem: vertical integration into add-on services, and payments from third-party affiliates. Value-added services could include, for example, a broadband operator's offering of voice-over-IP, home networking setup, or IP-based television. Third-party affiliate payments can also take various forms. For example, consider an online music service that chooses to give a broadband operator a percentage of customer revenues in return for promotion of their service to the operator's customer base. Alternatively, consider an online movie delivery service that chooses to pay selected broadband operators to ensure that the service's customers on those particular networks will not experience operator-imposed limitations on their movie downloading activity, or will be afforded priority service during periods of congestion.

PART II: WHY TODAY'S APPROACHES ARE INADEQUATE

Flat-Fee Pricing

The popularity of flat monthly access fees is not surprising. Their simplicity and predictability make them appealing to customers — an especially important consideration at this early stage of the broadband market, when penetration rates are still growing rapidly. Once penetration levels off, however, the "all-you-can-eat" nature of flat rates will prove problematic. The Korean example demonstrates that even a small group of users with insatiable appetites for bandwidth can easily cause aggregate usage to keep rising even as penetration rates and corresponding access revenues remain flat.

As bandwidth-intensive behaviors diffuse further into the customer base, this problem will only get worse. If broadband operators do not shift their pricing away from today's flat rates, they will be increasingly motivated to curtail rather than encourage many innovative uses of their networks. The perverse long-term implications of this incentive structure for operators and upstream value-chain participants alike lead us to consider it a problem for the entire broadband industry.

Why Not Just Raise the Flat Rate?

Given the appeal of flat rates to users, and the simplicity and lower cost of flat-fee billing, it is worth considering whether the broadband incentive problem could be solved if providers simply raised their flat rates to reflect increased mean usage. Perhaps, but the likely increased dispersion among users must also be considered in evaluating the impact of this solution. Increased variance in the user base may make "one size fits all" pricing less desirable and more difficult to sustain.

Not all networked applications are usage-intensive, and the Korean experience is consistent with others in suggesting that increases in the mean can be driven by intensive usage from even a small fraction of the user base. The mean, in other words, does not represent any "typical" user. While many users are clearly willing to pay some premium for the predictability of a fixed bill, increasing variance among users suggests this premium would grow for most users if the flat rate were raised.

If flat rates are the only option available, some low-intensity users will be priced out of the broadband market. That outcome is undesirable for upstream value-chain participants, since it limits their addressable customer base. Alternatively, if competitive options exist that better meet light users' needs (for example, with usage-based prices), low-intensity users will be inclined to defect from the flat-rate network. That outcome is clearly undesirable for the flat-rate network operator. In other words, raising flat rates fails to achieve the alignment of incentives that is needed if a solution is to benefit both network operators and upstream value-chain stakeholders.

Pricing Tiered by Peak Rate

As described in the section above on revenue models, a common access pricing scheme today consists of tiered flat fees differentiated by the peak rate at which a user can send and/or receive traffic over the access link. Peak-rate tiering bears some intuitive similarity to tiered pricing schemes in other networked industries — buckets of cell-phone minutes, for example — and may be an effective market segmentation mechanism. However, peak-rate tiers are actually not a very good proxy for the costs imposed by user traffic. In particular, peak-rate tiers do not protect broadband operators from high-volume users. Operators dimension their networks based on an expected busy-period duty cycle that is far less than 100% for every tier, regardless of its peak rate. Thus, even a subscriber to a low peak-rate tier whose duty cycle is much higher than expected (for example, a subscriber who sends or receives traffic continuously) can raise an operator's traffic-sensitive costs.

This problem will get worse as access networks become capable of much higher peak rates, because the volume of traffic from high-duty-cycle users will be able to vary much further from the operator's expectation. For example, usage tiers of 1, 2, or 3 Mbps are not likely to induce as much variance in traffic volume as the usage tiers of 1, 10, 100, and even 1000 Mbps (i.e., 1 Gbps) possible with fiber-to-the-home networks. Put another way, peak-rate tiering exposes operators to increased risk from heavy users, since as capacity expands, such users are capable of sending ever larger volumes of traffic. This increased risk creates a disincentive for providers to make their access networks capable of much higher peak rates.

Not only do peak-rate price tiers not address the incentive problem, they can also have harmful side effects. When networks are not congested, the most efficient way to handle traffic is to clear it through the network as quickly as possible. Peak-rate price tiers, however, impose throughput limits on the traffic coming from users who subscribe to lower-priced tiers. Slower clearing of that traffic can at times lead to congestion later, or in other words, to higher costs. Peak-rate price tiers can also harm upstream innovation and overall industry revenues by unnecessarily limiting the adoption of applications requiring high peak (i.e., burst) rates. Some applications and services, for example some networked games, will depend on low latencies but not require sustained high traffic volumes. It is not particularly in anyone's interest (other than, perhaps, the parents of game-addicted teenagers) to artificially suppress the adoption and use of such applications.

Pricing Tiered by Traffic Volume

Less common today, but becoming more so, are access pricing schemes that differentiate price tiers based on the monthly volume of traffic sent by a user. Since the volume of user traffic is more closely related than the peak bit rate to providers' recurring (i.e. non-fixed) costs, volume-tiered schemes have the potential to better align user and provider incentives. Achieving that potential, however, requires at least three types of refinements to today's somewhat crude volume-tiered schemes.

First, providers need to continue developing tools and techniques, beyond the first generation of fixed "usage estimators" that have begun to appear, to help make volume-based tiers more palatable to users. Today's pricing schemes generally express volume tiers in units (e.g., Gbytes) that are not terribly meaningful to users. Unlike the number of hours spent dialed into a narrowband ISP, or the number of minutes spent talking on a mobile phone, the number of bytes sent or received by a digital communications application is largely invisible to and beyond the control of users. Broadband operators may find inspiration in approaches adopted by other consumer utilities, such as electricity and water, to metering, pricing, and related techniques (e.g., energy efficiency rating and labeling schemes). Whatever techniques providers adopt, they must allow for the fact that networked applications can communicate without the knowledge or participation of users, and not always for legitimate or expected purposes (e.g., various forms of malware).

Second, volume-based tiering schemes could be aided by research into effective analytic techniques useful for deriving appropriate volume cutoff points based on the characteristics of network traffic, especially since those characteristics are likely to change over time as new applications and user behaviors emerge. In the dialup and cellular context, where usage limits have been based on the time that a user spends communicating, it has been relatively simple to design cutoffs that distinguish "normal" from "extraordinary" usage. This kind of differentiation is much less straightforward in the always-on broadband context, and remains an important topic for ongoing research.

Third, and most important, research and experimentation needs to continue to refine volume-based pricing schemes so that they can more closely align user willingness to pay (or not) with the need for actual traffic management on provider networks. The aggregate monthly traffic volumes used to set price tiers today are only gross approximations for the actual costs that any user's traffic imposes on a network. Those costs also depend on factors such as when the user's traffic arrives (at the same time as everyone else's, such that the traffic congests network resources, or at a time when those same resources would otherwise go unused?) and where the traffic is headed (for example, does it need to exit the access provider's network?).

Ideally, access pricing would perfectly align user and provider incentives such that only traffic that imposes additional costs is subject to either additional fees (if the user is willing to pay the extra costs), or to limitations (if the user not willing to pay). The ideal scheme would strike the perfect balance between encouraging upstream innovation and encouraging investment in next generation networks. In reality, the preferences of many users for predictable pricing may mean that approximations to such a perfect balance may be a more realistic goal. Regardless, the state of network pricing research suggests that a better balance could actually be achieved in practice than is being accomplished today using the blunt instrument of monthly volume caps.

Why Not Just Throttle User Traffic?

We consider this question because throttling of user traffic is clearly already being employed by broadband providers as a solution to the problem of rising usage costs. Throttling is not without merit as an approach, in cases where users are not willing to pay to cover the costs that their traffic induces. The problem with most throttling schemes in place today, however, is that they do not generally observe this key distinction. Rather, they are typically blanket schemes that effectively place an infinite price on traffic once either congestion occurs, or (possibly arbitrary) usage limitations are exceeded. The negative effects of this blunt approach on innovators upstream, and eventually on providers too, have already been discussed above.

Throttling should be thought of more generally as a way of using prices to shape user behavior. Done right, it should allow willing users to pay more for their traffic that costs more, thus generating the revenues needed for ongoing operator investments in network capacity. In other words, more fine-grained throttling reduces to the same kind of pricing mechanisms discussed above in the context of better aligning the interests of users, providers, and innovators upstream.

Recovering Usage Costs from Non-access-based Revenues

As discussed above, non-access-based broadband revenues come in many forms; it would not be feasible to list them all here, let alone judge the magnitude of their potential contribution to a broadband operator's revenue. Yet two sources of non-access-based revenues — vertical integration, and payments from third-party affiliates — are so commonly mentioned as potential solutions to the incentive problem that it is important to highlight the problems we foresee with exclusive reliance on these mechanisms.

As expansion in broadband access enables a mass market in services for which consumers are willing to pay — including "old" services such as web and email hosting, and new ones such as voice-over-IP and IP-based video-on-demand— operators expect to capture additional revenues by vertically integrating into the provision of such services. Using such revenues to help offset any additional costs imposed by growing broadband traffic would be a familiar model to most network operators, since historically, consumers have paid directly for services (telephony, whether fixed or mobile, and television) and not for their transmission over a network.

Operators rightly foresee benefits from vertical integration. It can realize economic efficiencies, such as economies of scope when retailing costs are shared across multiple (possibly bundled) products and services. Integration may also help broadband operators expand their market reach and differentiate themselves. Relying on vertical integration as an exclusive solution to increasing per-user traffic costs, however, is problematic: service revenues may be insufficient to offset growing bandwidth costs, for several reasons.

One reason is market-based competition for value-added services. For example, suppose a broadband access provider plans to offset bandwidth costs with revenues from offering voice-over-IP (VoIP) service. Their VoIP service faces direct competition from third-party, PSTN-interconnected VoIP services, such as Vonage and SkypeOut. Such competition constrains the prices the operator can charge for VoIP to the marginal costs experienced by the third-party competitors. Because of the efficiencies of vertical integration, the provider may still be able to charge prices well above their own marginal cost. However, the provider also faces competition at the margin from application-only forms of VoIP, such as basic Skype, which generate no service revenue at all (i.e., the price to consumers is zero, clearly below the provider's marginal cost to provide VoIP service). Both types of competition will combine to limit the broadband operator's ability to price their VoIP service above their marginal cost and thus earn revenues that can be used to offset bandwidth costs.

Service revenues may also be limited by actual or threatened legal and regulatory interventions that constrain the extent to which network owners can differentiate their value-add services from those of third parties. From the VoIP example discussed above, it is evident that with vertical integration, operators have rational incentives to block or degrade third-party applications and services that compete with their own services. Such actions are clear targets for legal challenges, and there is at least one precedent for a regulatory response: in March 2005, the U.S. Federal Communications Commission ordered Madison River Communications, a small phone company serving over 200,000 customers in North Carolina, to stop blocking Vonage's service, and imposed a $15,000 fine. Although the revenues threatened in this case were from traditional voice (POTS) rather than broadband VoIP, the regulatory rationale was that "the Internet should remain open to all types of traffic." Regulatory responses will of course vary across countries and over time. For example, two months after the FCC levied its fine, the U.S. Supreme Court's Brand X decision affirmed vertical integration for cable-based broadband, and shortly thereafter, the FCC announced its intention to do the same for DSL. Whatever the communications-specific regulatory regime, however, broadband operators will still find themselves constrained to some extent by competition (antitrust) laws in most countries.

An additional problem is posed by bandwidth-intensive applications that do not have a revenue-generating service associated with them. Examples include peer-to-peer applications, BitTorrent, podcasting, and basic Skype. Some of these, like Skype, use relatively little bandwidth, but others, like video podcasting, can use a lot. Neither vertical integration nor payments from affiliates work as a model for recovering the traffic-sensitive costs imposed by such innovations. If only these revenue models are in place, providers will have no way to capture value from such applications. With no way to recover the costs imposed by bandwidth-intensive applications of this sort, providers will have every incentive to block the development and use of such applications, despite their obvious value to users and innovators upstream.

For example, exactly this situation is evident in the cellular industry today as it transitions to 3G. Operators are trying to keep content and services as closely integrated with their networks as possible, and have not generally encouraged the kind of third-party activities that would lead to free-wheeling mobile application innovation, analogous to BitTorrent and podcasting in wired broadband. While such innovations might enhance 3G's appeal with consumers, the incentive problem looms large, given the scarcity of spectrum and bandwidth available for broadband applications, and pricing models that generally follow the usage-insensitive precedent set by wired broadband access.

Finally, reliance on payments from third-party affiliates is also less than ideal, because it presumes a relatively static landscape of services and relationships among operators and service providers. New applications and services can emerge quickly on the Internet, and this is generally considered to be a valuable feature. But this would not be encouraged if, for example, traffic associated with new services — that have not yet established relationships with broadband operators — is throttled by default.

Engineering a Network to Reduce Traffic Costs

A final approach to the incentive problem is to assert that in fact it will not be a problem after all: traffic-sensitive costs will not rise significantly, because increases in traffic volume will be matched by corresponding declines in the cost per bit transported. This solution is often referred to as "overprovisioning," a commonly used shorthand that rather imprecisely refers to both static and dynamic aspects of bandwidth costs.

Statically, bandwidth exhibits economies of scale: at any given point in time, the price per bit per second (bps) declines as the overall capacity of a link increases. Dynamically, bandwidth prices are declining over time: at any given total capacity level, ongoing technical innovation will make the price per bps lower tomorrow than today. Examples of such innovation include Moore's Law effects that make switching equipment more capable, and improvements in optical and radio transmission that make communications links more capable. The proposed "overprovisioning" solution is to support higher levels of traffic at a lower cost per bit by exploiting bandwidth economies of scale and price declines over time.

The key issue with overprovisioning is how the rate of price decline (whether static or dynamic) compares with the rate of traffic growth. While both rates are inherently uncertain, a study conducted by one member company of the Broadband Working Group shows that in several plausible scenarios for the future, the rate of traffic growth will outstrip the rate of price decline. The analysis is based on historical trends in the rate of decline in bandwidth costs, and a range of scenarios for the addition of video-based traffic to the broadband applications mix. Anecdotal evidence from cable operators suggests that the study's most aggressive scenarios for the growth of video traffic are at least plausible.

In addition, trends toward localized ad insertion, and rapid growth in both the scale and variety of on-demand video, are increasingly personalizing even traditional (i.e. non IP-based) video traffic. Customized and personalized content undermine the potential of a number of engineering techniques commonly employed or proposed to conserve bandwidth by using existing network capacity more efficiently. Personalized TV limits the potential benefits of multicasting. Applications in which each user's content is unique (such as file-sharing programs or videoconferencing) do not benefit much from caching and mirroring approaches that move commonly accessed content "closer" to users. Hosting popular content on servers internal to an operators' network may indeed reduce opex associated with backhaul transit (Figure 4), as less traffic is exchanged with other networks. It is unknown, however, what proportion of user activity will relate to popular content.

In sum, plausible scenarios can be constructed in which traffic grows faster than the costs of capacity decline. Cutting costs is likely to be part of the solution to the incentive problem, but it is unlikely to be enough by itself.

Conclusions

This paper has articulated the consensus of Broadband Working Group participants that a critical problem exists which, unless solved, will ultimately stunt the growth of the industries that constitute the broadband value chain. The "all you can eat" pricing models that are common today create incentives for providers to limit usage growth rather than invest to support it. These incentives, while rational for providers today, are damaging for users and other upstream value-chain participants and therefore, ultimately, for providers as well (Figure 2). Today's most commonly proposed solutions — in particular higher flat fees, revenue models based on value-added services, and cost-reducing network engineering techniques — may be part of the answer, but are unlikely to prove sufficient by themselves. Good solutions to this problem need to align the incentives of network operators and upstream stakeholders, for example by enabling monetization of usage that imposes costs on providers. Solutions that achieve this alignment will produce the revenues necessary to support ongoing operator investments in more capable networks, enabling innovation and growth to continue in all parts of the broadband value chain.

The broadband incentive problem will manifest in different time frames depending on market maturity. It is already evident for wired broadband in Korea, where penetration has saturated and revenue growth flattened out, leading Korea Telecom to propose alternatives to flat-rate pricing. In mid-stage markets, where revenues are expanding rapidly along with penetration, the problem is often less obvious. Nonetheless, it has been recognized, as evidenced by the acceptable use policies imposed by some U.S. providers of cutting-edge, higher-bandwidth broadband alternatives, both wired and wireless, that give operators the option to terminate a customer's service if they use too many high-volume applications.

In all markets, the time to act is now. The earlier the problem is recognized, the easier it will be to address. The conclusions of this paper are therefore especially relevant for stakeholders in the least mature broadband markets, including emerging forms of wireless broadband (such as cellular 3G and WiMax), and any form of broadband in less economically developed regions, including developing countries. Stakeholders in these markets have the opportunity to establish sustainable user expectations regarding pricing and usage, at an earlier stage of market development. The incentive problem may be more easily dealt with in newer markets, if stakeholders do not follow the initial "all you can eat" pricing policies employed in leading fixed broadband markets.

Achieving and articulating the consensus reflected in this paper has been a necessary first step toward solving the broadband incentive problem. Further progress will require internal discussions within the stakeholder companies, as well as government policy makers, involved in the broadband value chain. We expect that such discussions will result in a range of possible actions, including:

• Initiation of stakeholder-internal efforts targeted at this problem. For example, corporate stakeholders may model the strategic impact of the incentive problem on their business plans, under a range of scenarios for the time frame within which the problem manifests; similarly, government stakeholders may choose to model the problem's overall impact on the economy, and examine its implications for public policy. Stakeholders may also consider how they might solve the problem based on their own perspective within the value chain.

• Mapping of this problem into ongoing techno-economic research, for example empirical and theoretical work on pricing of congestible resources that is being carried out in academic and industrial labs.

• Identification of external efforts needed to progress toward industry-wide solutions. For example, solutions may involve defining or specifying additional flows of information among value-chain participants. When appropriate, such efforts may be conducted within the neutral forum provided by MIT's Communications Futures Program.

The expected outcome of these actions will be to avoid the train wreck that would otherwise lie ahead for the broadband value chain, by steering the industry towards an incentive structure that can support sustainable innovation.

Acknowledgments

The ideas in this white paper reflect numerous discussions among participants in the Broadband Working Group. John Watlington, France Telecom, led this effort. He can be reached via *wad* "at" *ftrd.us*. Text was also contributed by Sharon Gillett, William Lehr, and Steven Bauer, all of the Massachusetts Institute of Technology and reachable via *sharoneg*, *wlehr*, and *bauer* (respectively) "at" *mit.edu*. Other CFP members who participated regularly in Working Group discussions of this topic include Bob Briscoe, Dave Brown, Lesley Gavin, Dave Payne, and Steve Whittaker (BT), Jeff Dean (Cisco), Jeffrey Burgan (Comcast), Hossein Moiin (T-Mobile), Dan Dahle (Intel), Phil Fleming, Dan Grossman, Gerald Joyce, Whay Lee, and Marie-Jose Montpetit (Motorola), Dirk Trossen (Nokia), Kathy Bharrathsingh and Deborah Stokes (Nortel), and Marvin Sirbu (Carnegie Mellon University). Special thanks to Natalie Klym and Mary Turner of CFP for external review and to David Clark, CFP Principal Investigator, for inspiration. 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.