The Evolving Internet - Traffic, Engineering, and Roles

Joseph D. Houle[1], K. K. Ramakrishnan[1], Rita Sadhvani[1], Murat Yuksel[2] and Shiv Kalyanaraman[3]
jdhoule@att.com, kkrama@research.att.com, rsadhvani@att.com, yuksem@cse.unr.edu, shivkuma@ecse.rpi.edu

[1]AT&T, [2]University of Nevada-Reno, [3]Rensselaer Polytechnic Institute

Abstract

Entertainment and real-time applications like voice-over-IP, medical telemetry, network gaming and streaming video are quickly becoming prevalent applications over packet-based communication networks. Not only do these applications demand a very diverse set of network performance requirements, networks themselves are also experiencing rapid growth in the number of users and traffic per user. There have traditionally been two canonical approaches to handling such needs: (i) provide substantially increased bandwidth to create sufficient overall capacity in a best-effort network, or (ii) provide class-based differentiated service to meet each application's performance requirements. The current network neutrality debate reflects the tension between providing substantial additional capacity to meet the most demanding application needs while exploring viable business models to fund this capacity growth.

We examine three dimensions to lend insight to this ongoing debate. First, we discuss traffic growth projections followed by a review of application performance requirements. Next, while assuming a mix of emerging as well as traditional applications, we develop queueing models for an IP backbone to quantify the economies achievable with a differentiated network compared to a best-effort network with enough capacity to achieve the required performance. Finally, we examine a variety of content delivery models to understand the flexibility needed to achieve the objectives of sustainability and customer expectations.

Importantly, our analysis quantifies the amount of over-provisioning required for an IP backbone that provides best-effort service while still meeting the needs of emerging real-time traffic with delay and loss targets. We calculate the Required Extra Capacity (REC) for a best-effort network to meet the same delay and loss performance for premium class traffic provided by a relatively simple, two-class differentiated-service network. Our results demonstrate that the REC increases as the network utilization increases or as the traffic becomes burstier, or the proportion of premium class traffic decreases. With conservative assumptions about the burstiness of the traffic (2-state MMPP parameters), REC approaches 60% even at average link utilizations of 60%, for a relatively small proportion (e.g., 20%) of premium class traffic.

Various business models have evolved to cover the distribution cost of delivering news and entertainment from content provider to consumers. These fall along a wide spectrum based on their differing proportions of distribution costs that are borne by advertising as compared to that borne by consumer subscription revenue. At one extreme, there are some newspapers whose delivery costs (and in fact all their costs) are entirely funded by advertising revenues and at the other extreme, there are media distribution companies (e.g., satellite radio) that are almost completely funded by the consumer. As the Internet plays a bigger role in the distribution of information and entertainment services, it too as a delivery vehicle will require flexibility so that both network operators and content providers alike can design sustainable pricing models that meet the needs of their customer base. Maintaining this flexibility is critical as the Internet's role in media and entertainment distribution continues to evolve.

1. Introduction

With Internet connectivity and web-usage becoming nearly ubiquitous, many day-to-day activities have migrated to the Internet. User expectation of Internet use has also evolved from a "best-effort" connectivity to a high performance medium meeting the bandwidth demands for all types of applications. Media and entertainment are also likely to generate a variety of applications which ride over IP networks. Service providers and consumers alike would prefer to use a single "pipe" for all their communication and entertainment needs, if possible. Quality-of-Service (QoS)-sensitive applications like VoIP, IPTV [39], gaming, and telemedicine will be offered over such a converged IP network. To satisfy requirements posed by these applications, ISPs need to engineer their networks to meet the service level agreements (SLAs) of their business customers, while simultaneously supporting highly unpredictable consumer traffic. Such provisioning must meet SLAs while being resilient to changes in customer demand, changes in application mix, and network failures.

Currently there is a wide ranging debate on the issue of "network neutrality", with both economic and technical aspects [12][13]. One key technical aspect of the net neutrality debate is whether best-effort traffic should be carried along with other (so-called "premium") application traffic for which SLA commitments have been made. At one end of the opinion spectrum are network neutrality proponents who suggest that there should be no differentiation of traffic and all application performance requirements should be met by over-provisioning the network. The question then is whether this can be done with a small amount of additional capacity or is there a need to significantly over-provision the network? Though QoS has been extensively studied, we believe it is important to quantify the extent of over-provisioning needed in a network while meeting the SLAs achieved by a corresponding differentiated-service network. The second aspect of our study focuses on this specific question by comparing a classless over-provisioned network with an engineered network using per-class queuing to offer Class-of-Service (CoS) (i.e., differentiated-service) network with the same offered load, user requirements, and SLAs.

The hypothesis of this paper is that a single-queue based over-provisioned service for meeting the SLAs of QoS-sensitive traffic and regular best-effort traffic is less efficient (from a capacity viewpoint) compared to an engineered network offering simple 2-queue Class-of-Service (CoS) differentiation. With this best effort approach, the network needs to allocate enough excess capacity over-and-above the capacity provisioned by CoS network in order to meet these multiple objectives. We show that this excess capacity required in a classless network is significant. Analytically, we estimate the required extra capacity (REC) for a classless link to match the QoS (in delay and/or loss) provided by its CoS-based correspondent while modeling the basic SLA requirements. We generalize this single link model to a network model taking into account the network topology, traffic matrices (based on a gravity model) and shortest path routing.

Further, we examine business models from other industries while recognizing that traffic growth and application diversity will continue to require significant network capacity growth, even with differentiated-service networks. The intent is to consider alternative models as the Internet evolves from a casual communication medium to a major distribution vehicle for news and entertainment services.

Structurally, the paper starts with motivation for our work by illustrating traffic growth of the Internet in Section 2. Section 3 outlines performance requirements of legacy and new applications for the

Internet. In Section 4, we position our work relative to the existing literature. Next, in Section 5, we cover our modeling framework for quantifying the amount of over-provisioning needed to match the application performance requirements that can be achieved by premium class of a class-of-service network. Section 6 describes role and management trends in similar markets and how they relate to the Internet as an information delivery channel. Conclusions are presented in Section 7.

2. Internet Traffic and Capacity Growth

As a broad platform for various network-centric applications, the Internet has grown faster than any other comparable technology platform. Internet growth is driven by the growth of applications which translate into the growth of traffic demands and more importantly, the growth of *user expectations* of Internet performance. Hundreds of millions of users are no longer casually using the Internet: they are *depending* upon it for their work and home lives. In this section we first look at the growth in Internet traffic *demand* and the growth of *capacity* (or bandwidth) to match that demand. In particular, our focus will be on the growth in traffic demand per-user and understanding the drivers of this per-user traffic growth rate. Since the user's willingness and ability to pay often does not grow in proportion with the costs associated with meeting the increasing per-user traffic demand, there is a need to explore alternative economic models to fund these incremental costs.

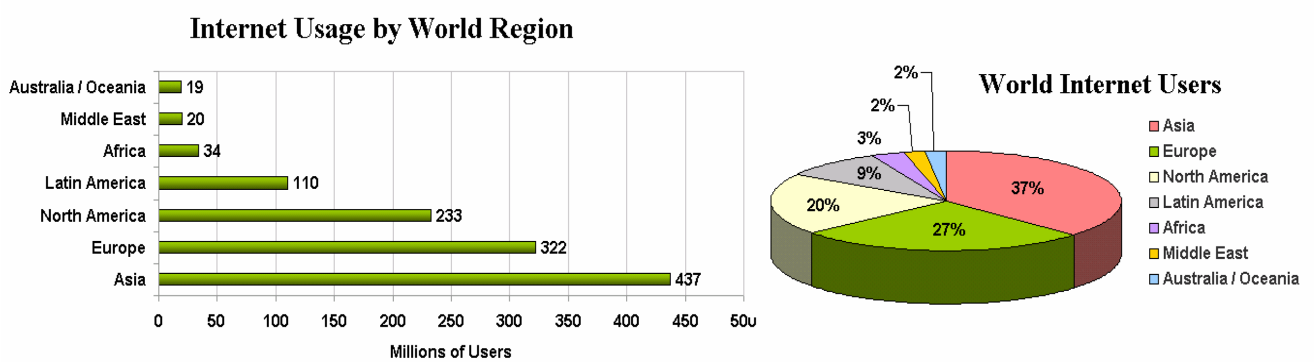

Figure 1: Internet Users Worldwide: Total (1.173 B) and Distribution (37% Asia vs. 20% N. America) [Source: InternetWorldStats [35]]

First, we take a look at Internet traffic or demand, which continues to grow briskly. Internet World Stats [35] estimates that as of June 2007 there are about 1.2B users, representing 18% of the world population. The mix of users also indicates globalizing shifts: ~35% users in 2000 were in North America compared to ~20% of the users in North America and 37% in Asia in 2007 (see Figure 1). Importantly, the recent annual growth rate of number of users is about 12.5% [35], a rate likely to continue. In comparison to this, the aggregate Internet traffic is expected to grow annually at 37% till 2011 according to a recent projection from Cisco Systems [43]. This higher rate of Internet traffic growth compared to the user growth implies that the *per-user traffic intensity* will grow at a rate of roughly 25% per-year (or doubling every 3 years). Cisco [43] also projects that consumer IP traffic (both Internet and non-Internet IP) will grow annually at 57% (driven by video traffic and broadband penetration). Business IP traffic is expected to grow at a slower 21%. Further, consumer IP traffic is projected to surpass business IP traffic in 2008 [44]. The shift of the majority of traffic from business to consumer combined with the fast growth of consumer IP video traffic changes ISP economics because consumers tend to pay much less per MB for video than for any other IP-delivered content or service [43]. We now explore the dynamics and impact of broadband access and video traffic on per-user traffic in more detail.

Broadband access is past the inflection point globally (over 17% of global users and 38% of North American users having broadband access) with approximately 300M broadband subscribers at the end of Q1 2007[1]. Assuming an average download speed of 250 Kbps per subscriber and a 5% duty cycle, this implies a global broadband-driven traffic estimate of 1.2 exabytes/month[2]. This compares well Cisco's projection for consumer Internet IP traffic of 1.4 exabytes/month for 2007 [44]. These numbers suggest that the *average utilization of broadband access links is increasing to the 5% mark* beyond Odlyzko's [36] 2002 observations of broadband access usage being < 1%. The number of broadband subscribers will almost double to 567M by 2011[3] (faster than the 12.5% annual growth in the overall Internet user base). In summary, both the usage intensity and number of broadband users are growing faster than the corresponding aggregate numbers. The application driver for this is video traffic which we examine next.

Total Internet IP traffic in 2006 is 2.4 exabytes/month [50], forecast to grow to 10.7 exabytes/month in 2011 [44]. Out of this, the total consumer internet IP traffic is expected to grow from 1.4 exabytes/month (2006) to 7.8 exabytes/month in 2011 [44]. Video is the dominant traffic component of the Internet today; peer-to-peer (P2P) video, YouTube video [45], and IPTV video content (multicast and on-demand) are important sub-categories. YouTube involves short HTTP-based videos streams mostly 1-3 min long, corresponding to about 10 MB each [45]. YouTube serves over 100M videos/day [37], leading to over 18 petabytes/month, or 4% of consumer IP traffic in 2006. Bittorrent and other P2P systems involve downloads of larger files, accounting for 649 petabytes/month or 53% of consumer Internet IP traffic in 2006. Between 2007 and 2011, Internet video-to-TV (or IPTV) is forecast to increase by a factor of 10 (to 1.2 exabyte/month) due to the rapid growth of HDTV streams and on-demand downloads.

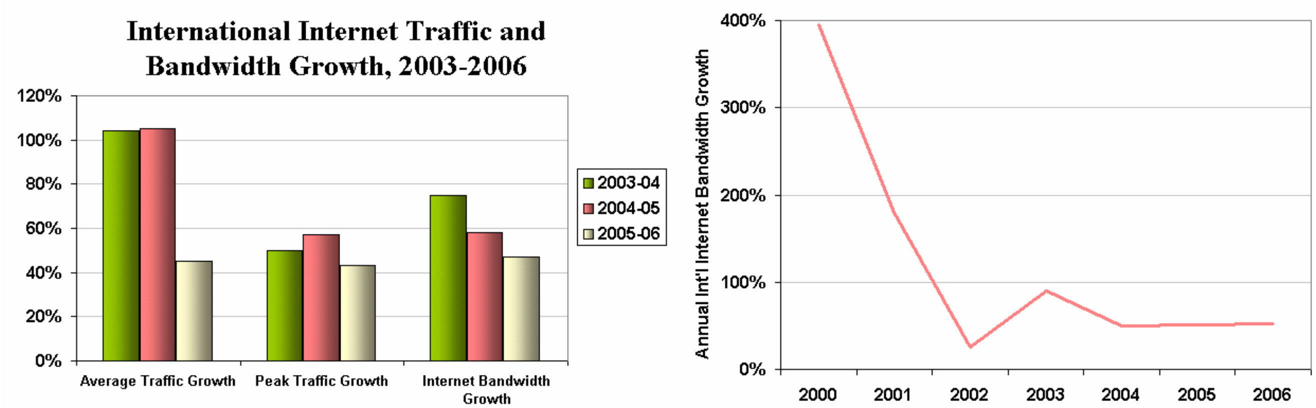

(a) (b)

Figure 2: (a) International Internet Traffic Demand vs. Bandwidth Growth. (b) International bandwidth growth (2000-2006) [Source: TeleGeography Research [38]]

We have so far discussed the demand (or traffic) side of the equation. Capacity (or bandwidth) deployment thus far has kept pace with traffic demand growth. During the late 1990s, spurred by optimistic demand growth projections, there was significant investment in backbone capacity relative to demand. But that has not been maintained in recent years (see Figure 2(b), which plots the growth of

[1] http://point-topic.com/home/press/dslanalysis.asp
[2] 1 exabyte = 1 million terabytes, 1 petabyte = 1000 terabytes and 1 terabyte = 1000 gigabytes.
[3] http://www.itfacts.biz/index.php?id=P8573

International backbone capacity growth). TeleGeography Research [38] reports a 75% growth in the average traffic on the world's *international backbones* in 2006 that outpaced the 47% growth of international backbone capacity (see Figure 2(a)). Interestingly, this trend continued for the third consecutive year. The capacity growth has also come down from the highs of 300+% in 2000-2001 and has been relatively stable since 2002 at well below 100% (between 40-50% recently). As the aggregate capacity growth remains below demand growth for a number of years, we would reasonably expect more balance between the levels of demand and capacity over time[4]. In terms of absolute capacity in international backbones, TeleGeography reports a total of roughly 1.6 Tb/s (or 0.5 petabytes/month) as of 2006 (see Figure 3 for inter-regional breakdowns). Though we do not have access to aggregate intra-regional capacity estimates, the aggregate Internet IP capacity needs to be significantly larger than total demand (estimated at 2.4 exabytes/month [44] in 2006) to maintain low average utilizations. For example, to achieve average utilizations below 20%, the aggregate Internet IP capacity today must be at least 12 exabytes/month (> 36 Tb/s). The open question is whether this aggregate capacity will continue to grow at rates that match demand growth, or there will be an increase in utilization (i.e., better balance). The initial evidence supports the latter hypothesis.

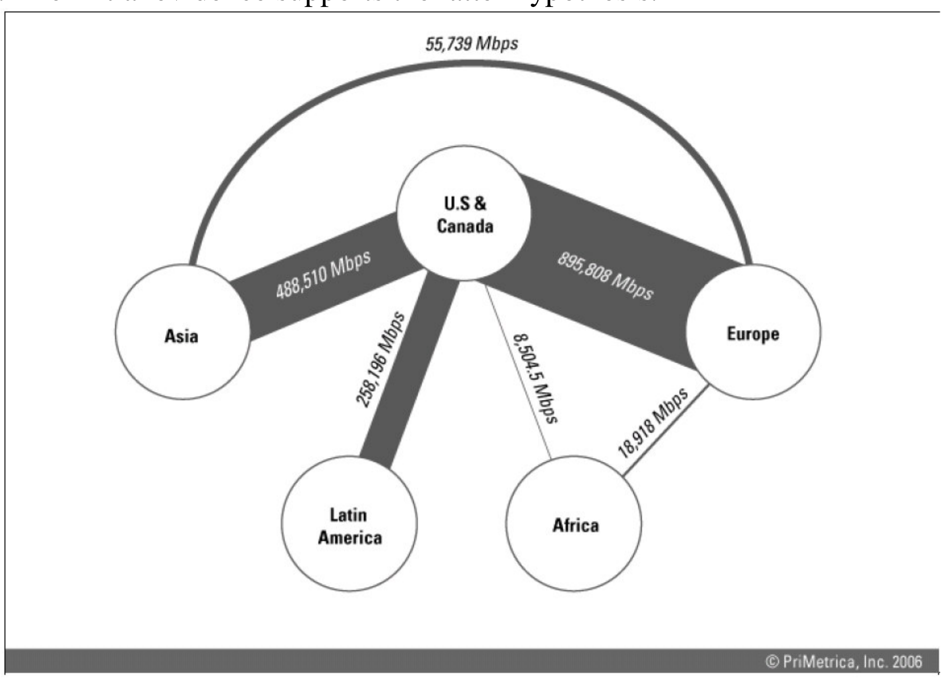

Notes: Data as of mid-2006. Interregional bandwidth below 1,000 Mbps not depicted.

Source: TeleGeography research © PriMetrica, Inc. 2006

Figure 3: International Backbone Capacity (2006). This does not include terrestrial bandwidth within each region. [Source: TeleGeography Research [38]]

In summary, the Internet is experiencing faster aggregate traffic growth (37%) than the growth in number of users (12.5%), driven primarily by various video applications. This implies a per-user traffic growth of 25% per year. Capacity data from international Internet backbones show increased balance between capacity and demand, suggesting that the days of dramatic over-capacity may be over. The shift in aggregate traffic from business to the consumer, combined with consumer IP video traffic over broadband access changes ISP economics. This is because consumers tend to pay much less per MB for video than for any other IP-delivered content or service [43].

[4] Note that a significant amount of capacity is set aside for restoration purposes as well.

3. Application Performance Requirements

Entertainment and real-time applications like voice-over-IP, medical telemetry, network gaming, and streaming video are quickly becoming prevalent applications over packet-based communication networks. These applications impose an extremely diverse set of network performance requirements. First, the demand on network bandwidth for an individual video stream to a user is quite significant – a standard definition television channel needs up to 2 Mbit/sec., while high definition television channel requires between 6 to 10 Mbits/sec, depending on the encoder. While multicast ameliorates the demand on the backbone, the growth in the number of channels and the use of video-on-demand are likely to continue maintaining pressure on backbone capacities. These requirements are in the form of desired delay target above which there is a potential for the application to be impaired. In addition, the application and the protocols used for communication (e.g., TCP) often impose a packet loss target above which performance (in terms of application throughput and end-end delay) may be adversely impacted. To satisfy requirements posed by such applications, ISPs need to provision their networks to meet the service level agreements (SLAs) for business customers, despite high and variable background traffic from customers without SLAs.

APPLICATION	MAXIMUM ONE-WAY DELAY	PACKET LOSS IN THE NETWORK
IPTV	<100 msec	<0.01%
Video-on-Demand	<50 msec	<0.001%
VoIP	<150 msec	<0.1%
Video Conferencing	<150 msec	<0.05%
Gaming	<50 msec	<0.1%

Figure 4 - Performance Requirements by Application

Figure 4 illustrates some typical performance requirements for a selected class of entertainment and real-time applications. For example, a toll-quality IP Telephony service typically imposes performance requirements. These include: (i) Low delay: end-to-end packet delay must be small enough that it does not interfere with normal voice conversations, and (ii) Low packet loss: packet loss must be small enough to not perceptibly impact either voice quality or the performance of other equipment that use it as the underlying communication medium (e.g., legacy fax). The maximum one-way delay, acceptable for most interactive voice usage is about 150ms. Based on the analysis in [33], the delay budget for queueing in the backbone network is approximately 10ms after taking into account propagation, coder, silence suppression, de-jitter buffer, and access network delays. Typically, the requirements for packet loss for encoded speech are 1% or less. While loss concealment algorithms can be used to reproduce intelligible speech even with higher loss rates, the resulting performance may often be considered to be inadequate. However, in addition to such QoS under typical conditions, premium application traffic expects to have the service protected under transient failure conditions as well. While the above may be a worst-case situation, it is important to note that interactive real-time applications impose non-trivial loss and delay constraints on the network.

Video also imposes stringent loss requirements because of the variety of players that may be used by consumers. These have different amounts of storage available as well as varying loss-concealment capabilities. Further, the use of User Datagram Protocol (UDP) as the transport protocol also imposes limitations on the ability of the end-point to recover from packet losses. Impairments due to packet losses can be easily visible, and it is highly desirable for the service provider to keep these at a minimum. Thus, the packet loss requirements for video are quite stringent for high-quality video distribution.

Multiplayer network-based gaming is another application that has stringent performance requirements. For example, [34] suggests that the combination of packet loss and delay (and in some cases, jitter) play a significant role in user-perceived quality of the game. In [34], the authors suggest that user-perceived quality and player sensitivity to achieved network performance can be measured by the departure rate of game players. A user that is satisfied with the quality tends to stay longer on the game; and hence, when both loss and delay increase users leave the game with increasing frequency. Target loss rates of 0.1% are suggested and delays on the order of 50 to 100 ms are required even for role-playing games with limited interactivity. Given that users are distributed across a wide geographical area, the obvious limitations of propagation delay will limit the achieved latency, but the perceivable allowance for queuing delays in the backbone network will be increasingly stringent. We anticipate that the guideline we used for VoIP of having 10 ms for the queuing delay budget in the backbone network would apply for the gaming case as well.

In the future, with the use of the network for applications such as interactive medicine and telemetry, it may be such that the performance requirements for these applications will also be stringent. Since the delay requirements will ultimately be limited by propagation delays, the significant requirement will likely be on packet loss requirements, because the budget for recovery of lost packets from retransmissions will be smaller and smaller.

4. Related Work

Costs and opportunities of class-of-service (CoS) in Internet have been researched heavily recently, especially within the context of "network neutrality" debate [12]. The technical side of the debate involves mainly *value or cost quantification* of several issues such as provisioning of the diff-serv forwarding behaviors [22] and estimations of capacity versus demand scaling [8][23]. Many of the papers on this topic [24][13] imply that an ISP should be allowed to provide *differentiation* in its services, a result inline with the decades of QoS research. Our work focuses on a simple 2-class vs. 1-class model at the aggregate level without admission control. An analysis similar to ours was done by Sahu et al. [22] in comparing loss performance of forwarding behaviors (i.e., discard eligibility vs. priority) of the diff-serv architecture. Instead of services specific to the diff-serv architecture, our work compares the classless service to the class-of-service in general. We also provide the quantitative comparison at the edge-to-edge (g2g) level with full consideration of network-specific issues such as the topology and the traffic matrix. While others [23] have examined the benefit of over-provisioning to overcome traffic uncertainty and to accommodate scaling up of the network, we examine the relative benefit of CoS support in terms of capacity savings.

Policy and regulation issues related to the network neutrality debate have been a very active topic of research as well. Several studies have argued for *deregulation* as the basis for the Internet's successful growth with historical analysis. For example, [25] suggested that the Internet has evolved to its current level of success due to minimal regulation in policies, while [26] suggests that the market and operational flexibility are the reasons why the Internet has been replacing the Public Switched Telephone Network (PSTN) even though it was initially an overlay on the PSTN. [27] further suggests that regulatory actions should be made in minimal incremental steps to avoid any undesired outcomes, maybe unforeseen for the lifetime of regulators initiating such action. Researchers also attempted to identify factors involved in Internet access demand and inferred that subsidization of Internet access may not be an effective way of enabling a "universal access" environment [28]. Interestingly, income and population density were identified to be small factors in contrast to conventional intuition [32].

There is a vast body of network QoS literature studying different queuing, scheduling and buffer management mechanisms to allocate finite capacity and delay (given an average utilization) amongst flows at a statistically multiplexed resource [9][10]. Recent work by Ciucu et al. [11] proposes a provisioning strategy based upon a statistical service curve characterization in conjunction with admission control or shaping and policing, and argues that scheduling has little value added above such provisioning. In our work, a key difference is that we do not have admission control or shaping/policing of input traffic. But, since the network must still honor premium-traffic SLAs, we show that simple Class-of-Service (CoS) scheduling is valuable. In [8], Kelly suggested that CoS scheduling may become redundant in a future Internet if end-nodes intelligently tune their demand based on congestion indications [40] from the network. However, we envisage the Internet to be an environment where end nodes may be non-cooperative and not necessarily react to congestion in a single prescribed fashion. Furthermore, the use of UDP with varying degrees of cooperation and congestion avoidance will likely persist as usage of multimedia streaming grows. Previous research has also suggested that simple (in terms of being responsive to congestion in the network) access to the network is a given and its existence in the history of telecommunications is prominent [29][30][31].

QoS research has recognized the need to simplify and de-couple QoS building blocks to promote implementation and inter-network deployment. The IETF Int-serv work [16] and Differentiated services [17] models simplify the architecture for supporting CoS in the core IP network. Core-stateless fair queuing (CSFQ) [20] further simplifies the core architecture and moved the data-plane complexity to the "edges," and allowed flexibility in the choice of control-plane options [18][19]. IETF's diff-serv has been used in private network services (e.g., Virtual Private Networks (VPNs) [21]). The objective is to primarily provide good service for those applications and customers that need and request it. In this paper, we look at the quantitative benefit of diff-serv/CoS in a common Internet, where we will have both premium-traffic SLAs to honor, and best-effort traffic. Instead of explicit class differentiation, the flow-aware networking approach [14][15] suggests the use of implicit differentiation by using per-flow queuing and per-flow admission control.

5. Quantifying the Value of Class-of-Service Support

As part of our effort to outline policies, quantification of the value of supporting class-of-service (CoS) or the cost of classless service is crucial. We devise a modeling framework to perform a comparative quantification of the capacity required for a classless service versus a CoS network for various ISP topologies. We start with a simple bottom-up comparative model of the two services, which poses the question: "How much extra capacity needs to be provisioned for the classless service to meet the same *performance* (e.g., in terms of *delay, loss*) as the premium class traffic achieves in a CoS link with a particular *aggregate (including both premium and best-effort) traffic load*?" We use this link model to derive the relationship between the **required extra capacity** (REC) and the following parameters: premium class delay and/or loss and aggregate traffic load. We derive equations for REC with non-preemptive priority scheduling for the CoS link. We then extend this link-level model to a network model where edge-to-edge (g2g) premium class performance goals are built upon link-level performance goals throughout each g2g path. This enables us to use the link-level REC model for calculating the needed extra capacity for each link in the network.

5.1. Preliminaries

We start by considering two traffic classes on a CoS link: *premium class* and *best-effort class*. We set a performance target of delay or loss for the premium traffic on the CoS link, and then seek to find the required extra capacity (REC) for a classless link (which treats both traffic classes identically in a single FIFO queue) to achieve the same performance target for both the traffic classes. Figure 5 illustrates the comparison of the two cases of the link service. Let the aggregate traffic rate be λ_D that is served by a CoS link with a capacity of μ_D. Also let a fraction of this aggregate traffic be premium class traffic with a rate of $\lambda_{Prem} = g\lambda_D$, with the remaining being the best-effort (BE) class with a rate of $\lambda_{BE} = (1-g)\lambda_D$. For the premium class traffic, we define a performance goal ζ in terms of delay or loss.

Given the parameters as illustrated in Figure 5, we formulate the necessary classless link capacity μ_N to achieve the same performance target ζ for the aggregate traffic λ_D. From this, we can calculate REC in terms of rate as $\mu_N - \mu_D$ (or as a percentage by $100(\mu_N/\mu_D - 1)$). With this model, one can use average delay t_{target} or average loss probability p_{target} as the performance goal. In the case of loss, an additional parameter is the buffer size, which we express as K packets for each one of the traffic classes in the *CoS link* and *2K* packets for the aggregate traffic in the *classless link*. Notice that, for a fair comparison, we use the same total buffer size of *2K* in both scenarios: in the CoS link as well as the classless link.

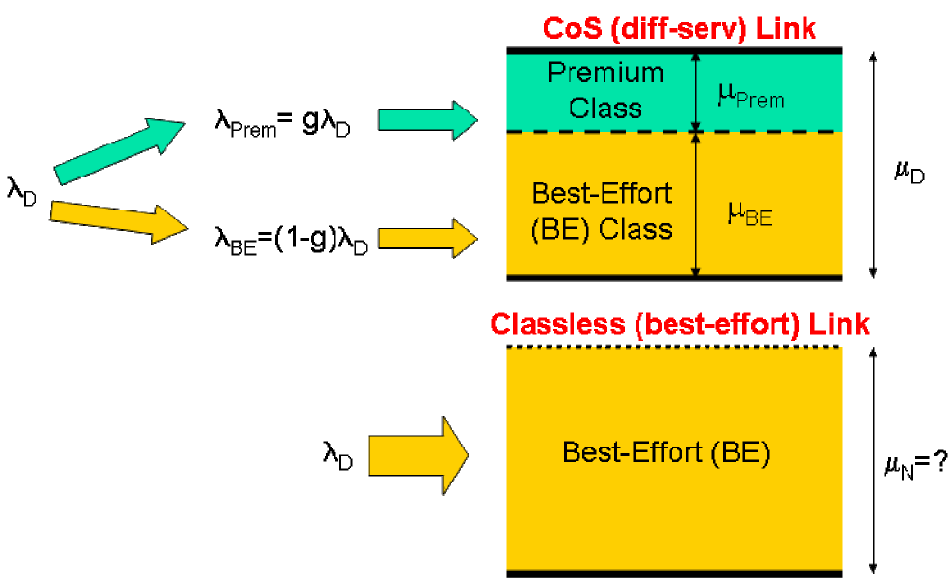

Figure 5: Link-level comparison of two different service types: neutral vs. CoS.

Due to the fact that *non-preemptive priority queuing* is one of the most prominent packet scheduling policies in networks supporting CoS in addition to being analytically tractable, we chose to further our analysis with it. We develop the link models based on two different traffic models: Poisson traffic (to provide us an initial understanding) and a Markov-Modulated Poisson Process (MMPP) [1] (to give us a more realistic traffic model). For simplicity of analysis, we assume that service time (i.e., packet size) distribution is exponential. We note that our estimates of REC will be conservative because we assume that the aggregate traffic (i.e., premium + BE) exhibits the same relationship for the first two moments for the traffic (i.e., the relationship between the mean and the variance) in each class. That is, if we are modeling the premium class traffic with a Poisson process of rate $\lambda_{Prem} = g\lambda_D$ and the BE class traffic with a Poisson process of rate $\lambda_{BE} = (1-g)\lambda_D$, then we assume that the aggregate traffic for the classless service is also a Poisson process with rate λ_D. Indeed, the superposition of two such traffic streams is known to yield a more bursty traffic than the burstiness of each individual traffic stream. If, as is often the case, the premium traffic (e.g., VoIP) is less bursty than the BE traffic, the REC would in fact be higher for the classless service in comparison to the CoS case.

In our earlier work [7], we have provided simple analytical derivations of REC. Our analytical study clearly showed that the REC depends on the rigor of the performance goal and the aggregate traffic rate λ_D of the CoS link. The average delay that the premium class experiences at the CoS link is dependent on three factors: (i) the aggregate traffic rate λ_D, (ii) the fraction *g* of the premium class in that aggregate, and (iii) the CoS link capacity μ_D. We also showed that REC and the performance achieved by the premium class of the CoS link can be expressed in terms of only two parameters: *g* (premium traffic fraction) and ρ (link utilization). In our graphical plots below, we will display the performance targets (t_{Prem} and p_{Prem}) as shades of color/grey on graphs that plot REC vs. *g* and ρ.

5.2. Link Model

To observe how REC behaves as the traffic becomes more bursty, we examine the case when the traffic is characterized by Markov-Modulated Poisson Process (MMPP) models [1]. The traffic model is obtained by composing multiple Poisson processes. The simplest MMPP model can be developed by means of two states (*i = 1, 2*) each corresponding to a particular sending rate λ_i of a Poisson process, with a target average sending rate of λ_t. As the ratio of λ_2/λ_1 gets higher the generated traffic becomes more bursty. Let the sending rate of the first state be a fraction *0<a<1* of the average rate (i.e., $\lambda_1 = a\lambda_t$) and the ratio of the traffic rates of the two MMPP states is a generic number *r* (i.e., $r = \lambda_2/\lambda_1$). Specifically, to generate a traffic with average rate of λ_t, we set the two traffic rates as:

$$\lambda_1 = a\lambda_t, \pi_1 = (ar-1)/(ar-a)$$
$$\lambda_2 = ar\lambda_t, \pi_2 = (1-a)/(ar-a)$$

where *0<a<1*, *r>1/a*, and π_1 and π_2 are the state probabilities. Note that our performance graphs use the parameters *a* and *r*, where the product *ar* is a measure of burstiness. In particular, we will use the values *a = 0.5* and *r = 4* or *r = 8* in our graphs.

Though we have derived the analytical formulations for the link models MMPP/M/1 and MMPP/M/1/K, we only present the results from a simulation-based link model for the MMPP traffic since the analytical formulas only hold when both states of the MMPP model send at a rate less than the link's capacity (i.e., $a\rho < 1$ and $ar\rho < 1$). Simulation-based calculation of REC eliminated this limitation of the analytical link model.

We used ns-2 simulations to calculate the REC. We simulated both the CoS link and the classless link for various ρ and g values, and matched the empirical performance of the premium class in the CoS link to the empirical performance of the aggregate traffic on the classless link. To simulate the CoS link, we used priority queuing of two flows passing through the link. For the classless link simulation, we used a FIFO queue for a single flow with a rate equal to the aggregate of the two flows of the CoS case.

In order to find the REC values by simulation, we match the performance (i.e., delay or loss probability) experienced by the premium class flow in the CoS link with the one experienced by the single aggregate flow over the classless link, within 1% error (6 repetitions of each case). We first simulate the CoS link 6 times for a given capacity (e.g., μ_D = 10Mb/s), utilization ρ, fraction of Premium traffic $g = \lambda_{Prem} / \lambda_D$, and buffer size K (if loss probability is the performance target). This empirically gives us the performance goal i.e., t_{Prem} (delay) or p_{Prem} (loss). We then matched this performance goal in the classless link simulations within 1% error, e.g., $t_{target} = t_{Prem} \pm 1\%$. To find the classless link capacity μ_N yielding a match to the premium class performance, we iteratively updated the link's capacity and observed whether the classless service performance matches that of the premium class traffic in the CoS case. At every iteration, we repeated the classless link simulation 6 times.

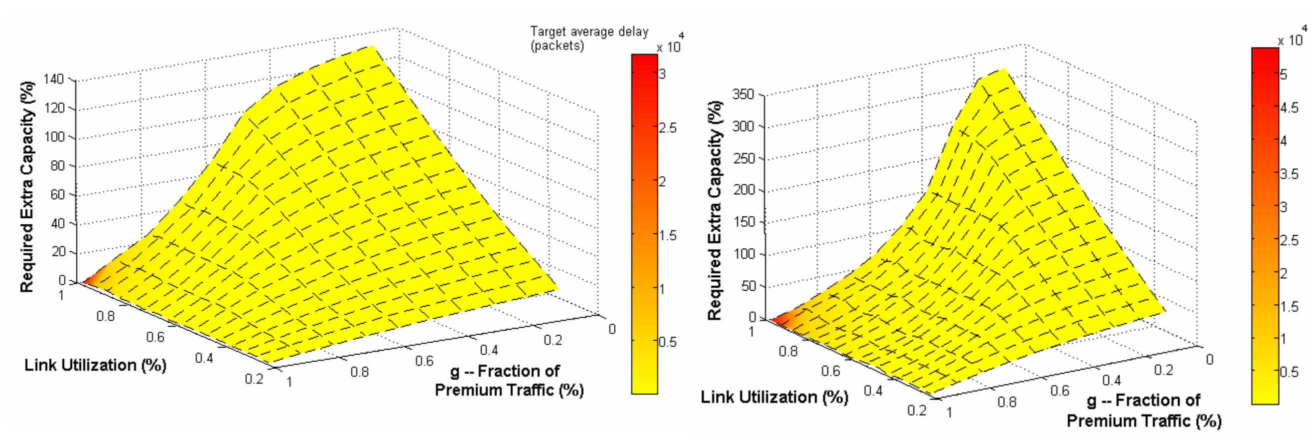

(a) *a*=0.5, *r*=4 (b) *a*=0.5, *r*=8

Figure 6: MMPP/M/1 link model results: The graphs plot REC as a function of link utilization ρ and premium traffic fraction *g*. The darkness (color) of the surface shows the target average delay (also a function of ρ and *g*), normalized in units of packets. For example, "1000 packets of delay" equals to 819.2 ms delay for 10 Mbps links and 8.192 ms for 1 Gbps links. Notice that even the light (yellow) areas on the surface have several thousand packets of delay, irrespective of link speeds.

5.2.1. Achieving a Delay Target: MMPP/M/1 Model

By using simulation-based calculation of REC as described above, we obtained REC values when average delay is the performance goal. Figure 6 shows the results while varying the traffic load and *g* for selected cases of burstiness (i.e., *a* and *r*). The darkness of the REC surface shows the target delay (i.e., t_{target} and p_{target}) in terms of the number of packet service times. As the product of *a* and *r* (i.e., burstiness of the MMPP traffic) increases, the REC also increases accordingly. Also, the regions where the analytical model is valid are apparent as the surface becomes steeper starting from the points satisfying the condition of $g = 1/ar\rho$, which is the minimum g value for each of the MMPP states for the best-effort class traffic to have sending rates less than the link capacity.

As we see in Figure 6, the REC grows as the link utilization becomes higher, but more so when the fraction of premium traffic (*g*) is smaller. On the other hand, when the traffic is predominantly of the premium class, there is less benefit from the differentiation. When the proportion of premium class traffic is small, even an individual arrival of that class at a classless queue would have to be serviced quickly. Thus the classless service would require a higher service rate than the CoS-based service which would treat the premium-class arrival with priority, keeping the delay experienced by that arrival small. As a result, when *g* is small, the REC would be higher. We believe this is important as we anticipate that current networks will likely see only a gradually increasing amount of premium class traffic. But even at $g = 0.5$, the REC can be quite significant (e.g., 50%) at high utilizations (e.g., 0.8). As we go from Figure 6(a) to (b), the increased burstiness also causes the REC to become significantly larger. Note that the target average delay (the bar on the right) is in multiples of 10^4, and the scale of the average delay we are considering here is large, of the order of 500 to 50000 packet service times.

5.2.2. Achieving a Loss Target: MMPP/M/1/K Model

We now look at the MMPP/M/1/K case when the performance goal is in terms of average packet loss probability, i.e., p_{target}. Note that the CoS link provides an equal buffer of *K* packets to each of the traffic classes, and that the classless link uses all the buffers (i.e., total of *2K* packets) for the aggregate traffic. Figure 7 shows the REC in percentage for various *a* and *r* combinations. We used three buffer sizes $K = 6, 15, 60$ packets with the motivation that conventional buffer sizes are similar for a 10Mb/s link carrying 1KB packets. Also, to keep g2g delay sufficiently low for applications, these buffer sizes are reasonable.

The darkness of the REC surface (scale shown in the vertical bar on the right) shows the target average loss probability (i.e., p_{target} and p_{Prem}) in terms of percentage. Similar to the MMPP/M/1 case, as the *ar* product (i.e., burstiness of the MMPP traffic) increases, the REC also increases accordingly. Our simulations treated loss probabilities less than 10^{-5} as zero. Therefore, REC values for small *g* and ρ values will likely to be higher than what is shown in Figure 7.

From Figure 7, we observe once again that the REC grows with utilization, particularly when the fraction of premium traffic (*g*) is small. Also, as one would expect, as we increase the amount of buffering *K* from 6 to 60 packets, the REC reduces, and the range of utilization where there is little or no REC required slowly increases. As the utilization increases, the loss probability goes up (increasing darkness). If the acceptable packet loss target is small, the REC also has to be higher. Going from Figure 7 (c) to (d), we observe that the increased burstiness (larger *ar* product) causes the REC to also increase, for the same p_{target} (reflected in the shade of darkness).

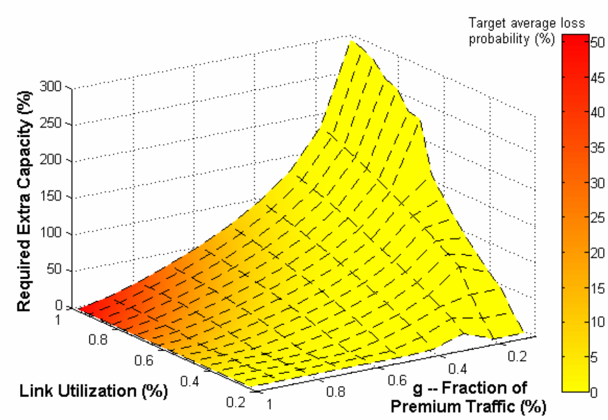

(a) a=0.5, r=4, K=6

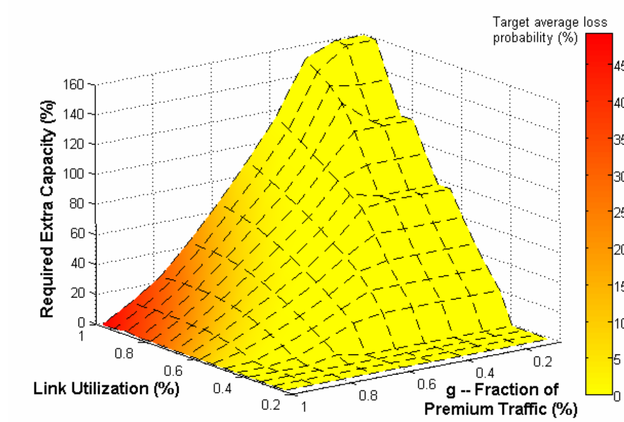

(b) a=0.5, r=4, K=15

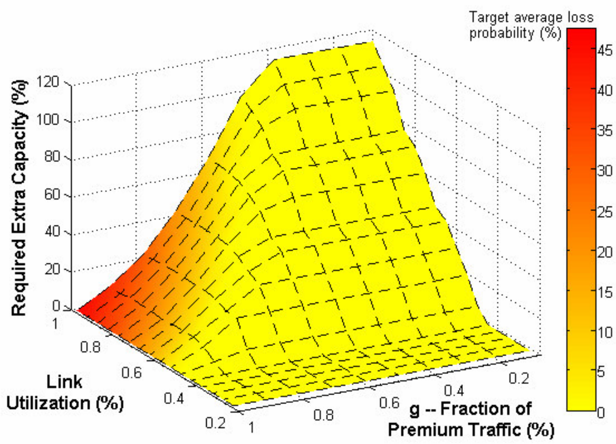

(c) a=0.5, r=4, K=60

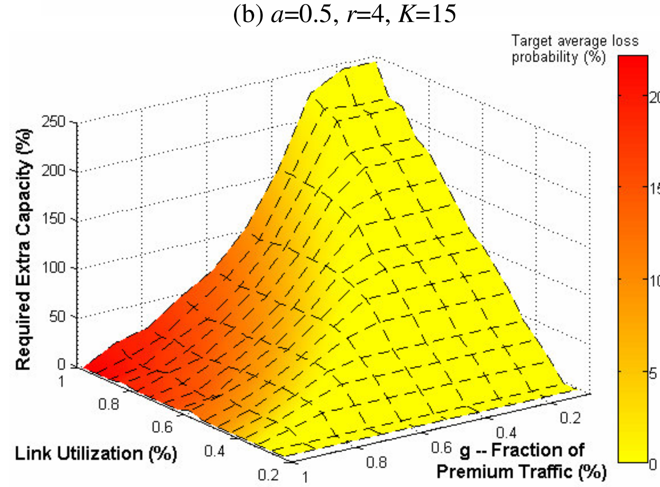

(d) a=0.8, r=4, K=60

Figure 7: MMPP/M/1/K link model: The surface darkness shows the target loss probability. Three buffer sizes K = 6, 15, 60 were simulated, roughly corresponding to 0.1ms, 0.25ms, and 1ms of buffer times (since the total buffer is $2K$ packets by our model) for a 1Gb/s link carrying 1KB packets.

5.3. Network Model

The final step involves generalizing the single link model to a network model. We focus on developing our network model for a typical ISP's backbone network. Crucial components of a network model include (i) a realistic topology (i.e. adjacency matrix, link weights, link propagation delays, link capacities) and (ii) a realistic traffic matrix. Given the topology and the traffic information, we calculate the REC for the complete network instead of a single link as was discussed in the previous section.

Briefly, we first calculate a routing matrix R for the ISP network from the link weight information. With a realistic traffic matrix T (we a gravity model to calculate one), we then calculate the traffic load pertaining to individual links by performing the product of T and R which shows the distribution of traffic loads on individual links. For each of these link traffic loads, the link model described earlier will apply. The link-load distribution will thus lead to a distribution of REC over links for the network.

The goal of the network model is to determine the *additional percentage capacity needed for a classless network over a CoS network* on an edge-to-edge (g2g) basis. The network model takes the following steps to calculate the network REC:

- *Step 1*: Construct the *routing matrix* $R_{F\times L}$ based on shortest path first (Dijskstra's) algorithm.
- *Step 2*: Form the *traffic vector* $A_{F\times 1}$.
- *Step 3*: Calculate the *traffic load on each link* by performing the matrix operation $Q = R^T \lambda$, where $Q_{L\times 1}$ is the link load vector (in Mb/s).
- *Step 4*: Check and fix the *feasibility* of the traffic load and routing.
- *Step 5*: Calculate the *per-link* REC by using Q_i as the traffic rate λ_D for the ith link and the performance goal t_{target} or p_{target} for that link i.
- *Step 6*: Calculate the *network* REC by averaging the per-link RECs from Step 5.

To obtain realistic topology information for ISPs, we used the Rocketfuel [11] data repository which provides router-level topology data. In order to assign estimated capacity values for individual links, we used a technique based on a Breadth-First Search (BFS) algorithm. We first select the maximum-degree router in the topology as the source node for BFS to start. Then we assign higher capacities to the links close to this max-degree router. We use gravity models [2][3] to construct a feasible traffic matrix composed of edge-to-edge (g2g) flows. We select the routers with lower degree and higher distance to the max-degree node as edge routers originating traffic into the network. We also ensure that there exists at least one edge router for each PoP in the topology. In the gravity model, we used the CIESIN [4] dataset to calculate the city populations.

To evolve from the link model to the network model, we split the g2g performance goals on individual links of the g2g path. In order to split the g2g delay target t_{target} on individual links, we simply divide the delay requirement equally on each link of the path assuming that t_{target} is only the queueing and insertion delay into the links. After *equally* splitting the g2g delay on individual links for all g2g flows, we collect the tightest (i.e., minimum) delay requirement on each individual link among the delay requirements imposed by each g2g flow traversing the link. Similarly, given a g2g loss probability target p_{target}, we assign the loss probability requirement on each link of the path as follows. Specifically, for a path with l links, we equally assign the survival probability to each link as the lth root of the overall path's survival probability $1 - p_{target}$.

In order to generate the network REC results, we use link model REC results for a given utilization and performance target. We perform a lookup from the simulation results of the link models (i.e., MMPP/M/1 or MMPP/M/1/K) and a linear interpolation on the link model REC values using the available datapoints to obtain the appropriate link REC. We use $a = 0.5$ and $r = 4$ as parameters for the MMPP traffic model. Real IP traffic is considered to be more bursty than what these values represent [5]. Also, when loss probability is the performance goal, we use a buffer size of $K = 60$ packets. For delay as the primary performance metric and since the g2g delay can be reduced by smaller buffer sizes [6], we chose a relatively small buffer of 120 packets (i.e., $K = 60$ packets, corresponding to a buffering time 100ms and 1ms for 10Mb/s and 1Gb/s links respectively with 1KB packets). Also, we use 0.1-10ms and 0.01-0.5% as the g2g queuing delay (excluding propagation) and loss probability targets respectively, representing required performance ranges for current and potential network applications, as described earlier.

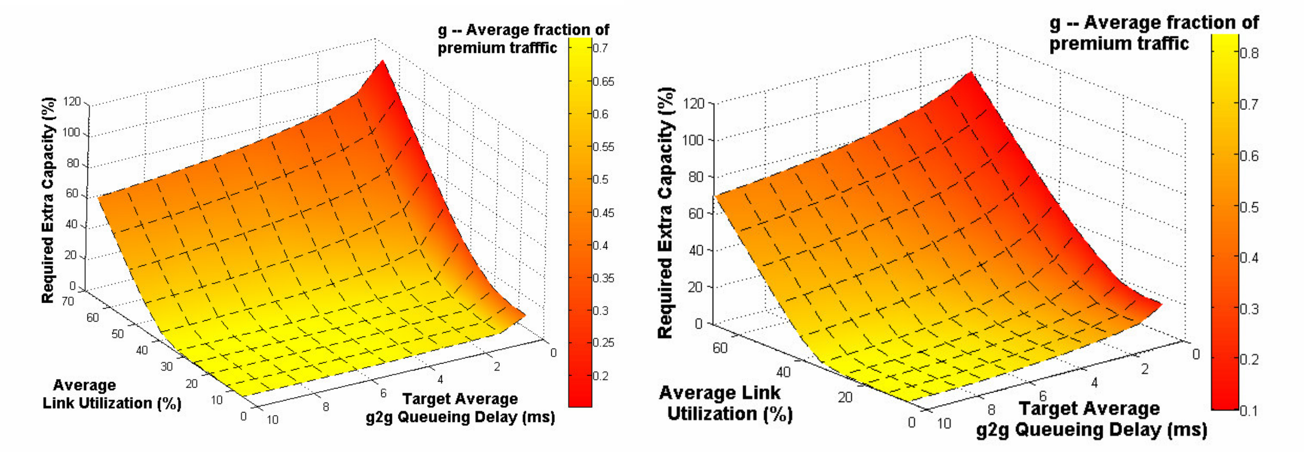

(a) Abovenet (b) Sprintlink

Figure 8: MMPP/M/1 network model: The surface darkness shows the g2g target queueing delay. MMPP's burstiness is defined by $a = 0.5$ and $r = 4$.

Figure 8 shows the behavior of average REC across all links for two different network topologies, when the target edge-to-edge delay is the criterion, for the MMPP traffic case. Note that the fourth dimension, reflected in the shading and the vertical bar on the right, represents g, the proportion of premium traffic. As the average link utilization goes up, the REC goes up, especially when the target average delay is smaller. As we see, when g is small, REC is higher, because the extra capacity needed for the classless service has to be higher to ensure that the arrivals for the premium class are served as quickly as the CoS case would, with non-preemptive priority scheduling. Although the exact amount of REC changes with each topology, the REC needed on each link is similar across the topologies (in the range of 50-100% at average link utilizations of 80%).

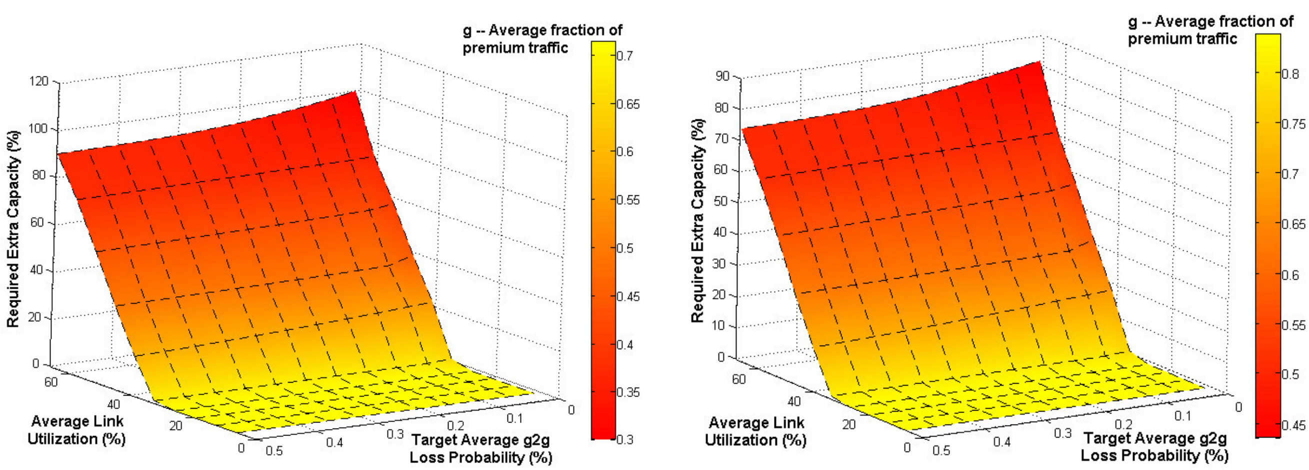

(a) Abovenet (b) Sprintlink

Figure 9: MMPP/M/1/K network model: The surface darkness shows the g2g target loss probability. The buffer size is $K = 60$ packets and the MMPP's burstiness is defined by $a = 0.5$ and $r = 4$.

Figure 9 shows how the average REC changes for the two topologies for the edge-to-edge packet loss criterion. We see across the network topologies, the REC increases as the utilization increases beyond a threshold (below which the buffering enables the classless service to avoid losses), and increases as the proportion g becomes smaller (darker shade). However, the increase in the REC is not as rapid when the target loss probability reduces from 0.5% down to 0.1%, which again reflects the role of buffering at each of the links. It is important to note however, that for average link utilizations of 80%, the average REC can be up to 100% for the case examined in the figure. It is also important to note that under failure situations link utilizations can easily get well over 80% carrying protection traffic, even in a well-engineered and provisioned network.

6. The Internet's Emerging Role

In this section we explain how some of the new applications are transforming the Internet from its traditional role of simple packet communications for applications such as email and file transfer into new delivery vehicle roles for rich media content. The needs of these applications can be supported by a variety of business and economic models. We survey some of the models that may develop in response to these emerging roles of the Internet.

6.1. Revenue Flows

6.1.1. The Traditional Internet

Revenue flows associated with the traditional Internet have been relatively simple as illustrated in Figure 10. Consumers and content providers/businesses buy Internet access services from their respective Internet Service Provider (ISP). In some situations, smaller ISPs buy service from one or more larger ISPs. The larger ISPs are connected to each other in a "peering" arrangement. In a peering arrangement neither ISP pays the other to carry traffic, but rather they agree to interconnect in a revenue-free model.

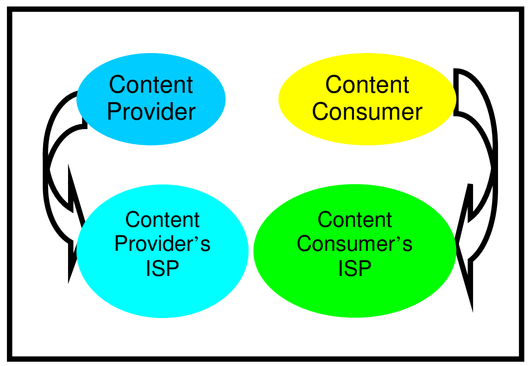

Figure 10: Traditional Internet.

6.1.2. The Internet as a Delivery Vehicle

As the Internet evolves from being merely a communication medium to being the delivery vehicle for the products and the services of content providers, a new revenue dynamic surfaces. As shown in Figure 11, content providers collect revenue from content consumers for their product. These content providers/aggregators are using the Internet to distribute their content globally and depend on the ISPs to build a high performance delivery vehicle for their product. This model leaves it primarily to the consumers to support the costs for network improvements necessitated by delivery of the content provider's product over the consumer's ISP.

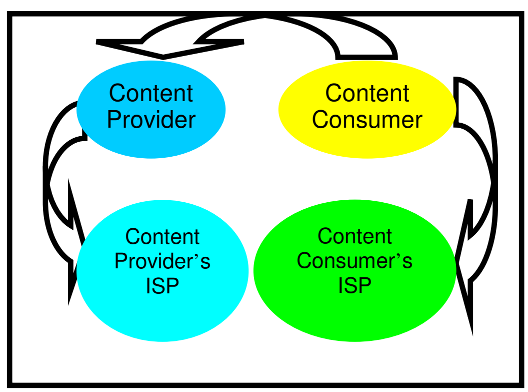

Figure 11: The Internet as a Delivery Vehicle

6.2. Subscriber Revenue–Advertisement Revenue Continuum

Could there be a way to re-allocate the costs of the delivery in the Internet "channel" to keep the total price to the consumer manageable? Examining the news and entertainment industry, we observe that various business models have evolved to cover the distribution cost of delivering the content from content producer to consumer. These fall along a wide spectrum based on their differing proportions of distribution costs that are borne by advertising vs. user-generated subscription revenue. As shown in Figure 12, at one extreme, there are newspapers whose delivery costs (and in fact all their costs) are entirely funded by advertising revenues. At the other extreme, there are media distribution companies (e.g., satellite radio) that are almost completely funded by the subscriber.

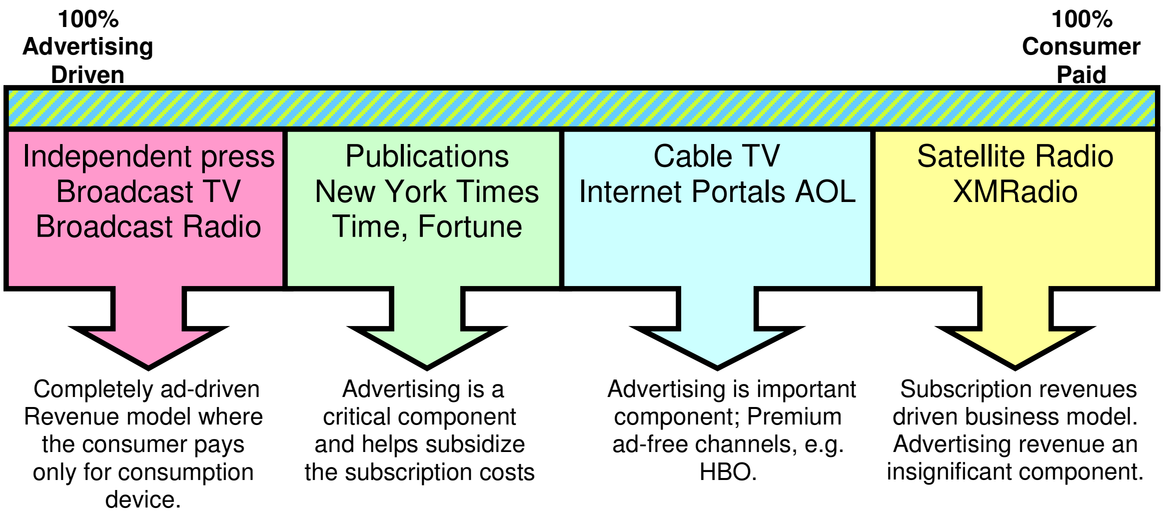

Figure 12: Spectrum of business models

The question then is whether the delivery costs for using in the Internet as a delivery vehicle could also be more aligned with advertising to subscription ratios seen in other media distribution environments? In the quest for a satisfying answer to this question, we examine some recent news events as well as financial analysis of media firms in the following subsections.

6.2.1. Case I: Wall Street Journal

One of the biggest potential changes with the News Corp's acquisition of Dow Jones is the possibility of moving the online version of Wall Street Journal, WSJ.com, to a free, fully ad-supported model. While the current subscription model likely helps WSJ command higher CPMs (*Cost per mille* [47], also called *"cost ‰"*) within its walled gardens (as well as the fact that almost 50% of WSJ.com revenue comes from subscription fees), by moving to a fully ad-supported model the journal can point to a more engaged and demographically attractive reader. It is widely believed that a free, fully ad-supported model will prove to be more sustainable in the long term and the benefits of expanding reach and gaining share of overall online advertising dollars will outweigh the opportunity and absolute dollars associated with the paid subscription model.

Analysis indicates that free WSJ.com would overtake the paid WSJ.com in revenue and profitability in a very short time. More importantly, the view is that the brand will become more valuable and get not only additional traffic but also audience growth. It is estimated that of the advertisement revenues, about $75M is generated by WSJ.com. In addition, WSJ.com will generate roughly $65M in subscription revenue in 2007, putting advertising/subscription revenues at a 54% / 46% split, or $140M in total [46]. Moreover, a free WSJ.com, with global distribution through News Corp would likely see an increase in visitors and advertisement revenue, potentially siphoning ad dollars from incumbent leading financial sites, such as Yahoo! Finance, MSN Money, and AOL Money & Finance.

6.2.2. Case II: New York Times

Financial analysis of the newspaper New York Times (NYT), largest daily circulation of all seven-day newspapers in the United States as of September 2006, indicates that circulation revenues bring in 26% of the revenues as compared to advertising which brings in 68% of the revenue. Subscriber revenues alone do not even cover the production costs of the newspaper (see Figure 13). Without an advertisement based model, a subscription to NYT could be three times higher, costing the subscriber $15/week. Moreover, its Internet revenues grew over 20%, such that the online businesses accounted for more than 10% of NYT's revenue growth every year. Essentially, NYT is diversifying its revenue by growing its online edition which is primarily advertisement funded.

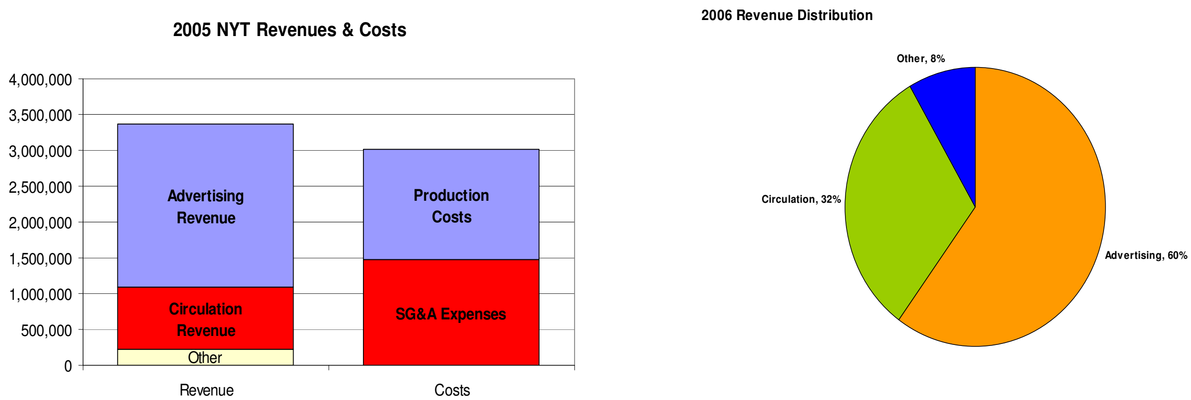

Figure 13: New York Times financial figures.

6.2.3. Case III: Sirius Radio

Financial analysis of Sirius Radio indicates that subscription revenues bring in 92% of the revenues as whereas advertising brings in 3% (see Figure 14). Moreover, with this revenue model, Sirius Radio has continually posted significant net losses year over year. This has led to a proposed $11B merger of struggling satellite-radio operators XM Satellite Radio Holdings Inc. and Sirius Satellite Radio Inc.

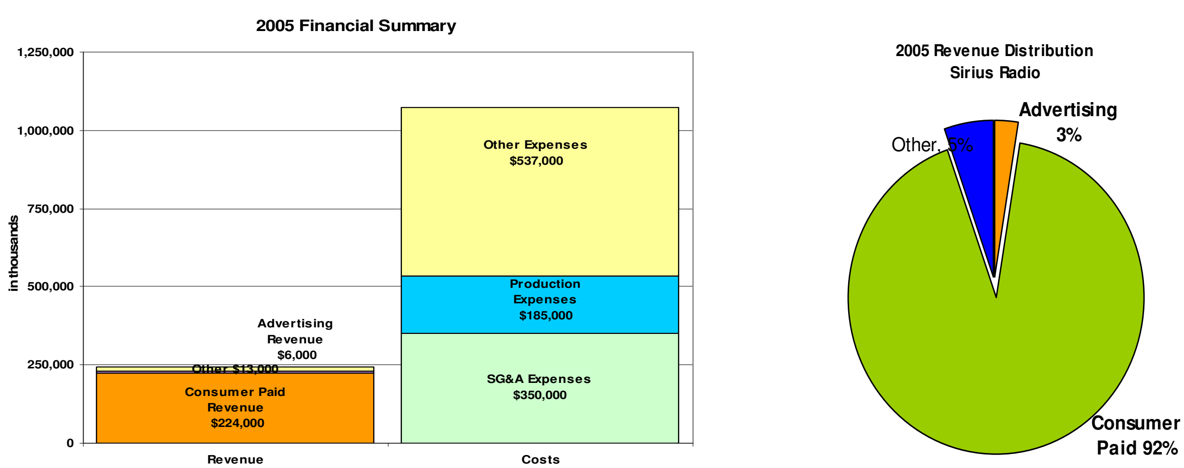

Figure 14: Sirius Radio financial statement - 2005.

6.2.4. Case IV: AOL

Year over year, AOL has seen a drop in subscriber revenues and increase in advertising revenue (see Figure 15). This appears to have pushed AOL to announce its transition from reliance on access and subscription-based services to a dual revenue model based on access and advertising.

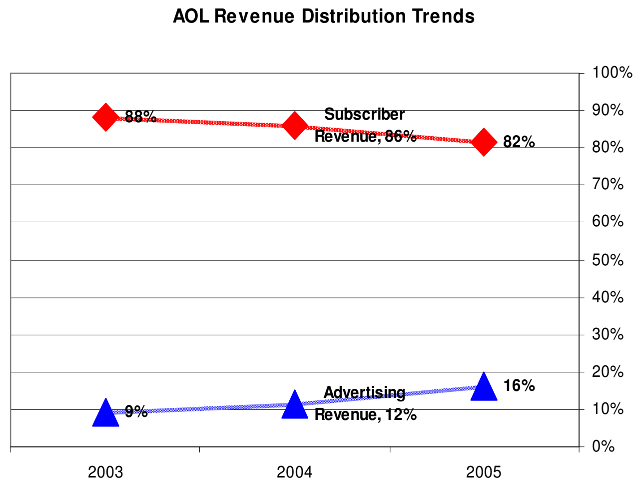

Figure 15: AOL revenue distribution trends in recent years.

6.2.5. Implications and Propositions

These examples illustrate a broad spectrum of successful subscription-to-advertising ratios, as well as changes of ratios in the same industry (even the same company) over time. We believe that a blended mix of cost assignment between end user and content owner subsidized by advertising would help the end user. Pushing the entire costs to the subscriber instead would increase the cost to the consumer, lowering the penetration and value of the service to the population at large. If the total costs of the broadband infrastructure are borne by the subscriber, a larger fraction of the population would be priced out of the services furthering the digital divide. Prophetically, five years ago O'Donnell [42] suggested: "if and when consumers switch their consumption of entertainment media from broadcast and broadcast cable technologies to IP-based delivery, we would expect that paid content - or advertiser supported content - will have to take up a larger share of the burden of maintaining the core of the Internet".

A strict imposition of network neutrality, as proposed at one extreme, would prevent new potential flows of revenue such as the advertising revenue (that is flowing to content providers today) to be used to offset the end-to-end cost of content delivery over the Internet.

6.3. Multisided Markets

As the Internet evolves from being a casual information channel for consumers into being a vehicle for enabling content providers to deliver their product to their customers, the Internet starts looking like a mediating platform in a multisided market. The Internet, through the investments of Internet Service Providers, brings together content providers and content consumers. Both content providers and consumers get more value out of their Internet connectivity as the other group becomes larger and more diverse. Consumers get access to entertainment and content providers get a delivery vehicle to deliver their product to their customers.

Price structure in multisided markets refers to the decomposition or allocation of the total price between one side and the other [41]. It is unlikely that historical Internet access cost ratios between content providers and subscribers is ideal for continued Internet growth. One element of a multisided market is that when one side of the market gets significant benefit out of the arrangement, there is often a willingness to sponsor participation on the other side of the market. Flexibility in price structure will be critical as the Internet continues to be a mediating platform for content delivery and other multisided markets.

6.4. Other Emerging Models

Revenue flow flexibility on the Internet has allowed other models to surface, with varying degrees of success.

6.4.1. Content Delivery Networks

Content delivery networks (CDN) are overlay networks on-top-of the Internet. The goal of a CDN is to distribute content from the content provider to a large number of consumers either more economically and/or with better performance than if the traffic were to flow end-to-end over the Internet natively. The CDN buys Internet access service from multiple ISP, deploys local caches, and then charges the content provider a fee to use the local caches for delivery of their content.

6.4.2. ESPN360

ESPN offers a service called ESPN360. Any broadband provider that wishes to carry it must pay ESPN a fixed fee for each subscriber. Unless your ISP pays ESPN, you won't be able to use ESPN360. In other words, the content provider blocks the ISP, not the other way around. Those in favor of net neutrality might argue that such a structure highlights the need for a net neutrality law that bans content providers from charging ISPs as well as the other way around. Those opposed to mandatory net neutrality might point out that a service like ESPN360 may not exist at all if ESPN were not allowed to experiment with a new pricing model.

As the Internet plays a bigger role in the distribution of information and entertainment services, it too will require flexibility so that both network operators and content providers alike can design pricing models that meet the needs of their customer base and match their business plans. Maintaining this flexibility is critical as the Internet's role in media and entertainment distribution continues to evolve. Internet companies need the ability to experiment with different business models depending on application or user-base to optimize the consumer-business value and continue investment in the packet transport infrastructure. Allowing firms, both those in the content provider market and those providing the enabling platform, to experiment with pricing is likely to be an important component in ensuring continued innovation in content and infrastructure.

7. Summary and Conclusions

The Internet continues to grow. We should expect to see a five fold increase in consumer traffic in the next five years. The number of users continues to grow at an annual rate of 12.5%. The average traffic per user is expected to grow even faster at 25% per year. A growing set of diverse applications will put broad performance requirements on the network. The Internet industry will need a cost effective approach to respond to the performance needs of these applications while simultaneously responding to the growth requirements.

A well-established approach for networks to meet the diverse application performance requirements is to exploit the use of Quality of Service support. There has been a large body of work on supporting Quality of Service in the network, and quantifying the benefits in terms of reducing both the magnitude and variability of delay and loss experienced in the network. On the other hand, there has also been considerable debate and advocacy on the benefits of having a simple classless (i.e., best-effort) service.

The current network-neutrality debate reflects the tension between providing substantial additional capacity with a classless service in comparison to a network providing QoS to meet the most demanding application needs.

This paper provides a quantitative analysis of the amount of extra capacity required of such a classless network to support traffic that requires delay and loss performance, which has not been explored previously. In this paper, we provided quantitative analysis of the required extra capacity (REC) for a classless network to meet the same delay and loss assurances that would be provided by a relatively simple two-class Differentiated Services-based network. First, using a realistic two state MMPP traffic arrival process, we quantified the REC for a single network link. We observed that REC grows with utilization, and is of particular concern when the proportion of premium class traffic requiring delay or loss assurances is small. We also note that the REC increases as the traffic becomes more bursty. We examined REC in the context of both delay and loss SLAs.

We then examined the behavior of REC with more complex/realistic network models for an IP backbone. We observe that the network-REC increases as the average utilization of the links in the network increases and as the relative proportion of the premium traffic reduces. Moreover, network-REC grows rapidly as the acceptable delay and packet loss targets become tighter (smaller). As an example, using conservative assumptions on the burstiness of the traffic (2-state MMPP parameters), network-REC approaches 60% even at reasonable average link utilizations of 60%, for a relatively small proportion (e.g., 20%) of premium class traffic.

In addition to understanding the need and extent of additional capacity required in providing a classless, undifferentiated network, we also examined the issue of the allocation of the costs for distribution of content over the network. Towards this, we examined the business models that have evolved over the years to cover the distribution cost of delivering news and entertainment from the content provider to the consumer in other media. We briefly examined various business models along the revenue source continuum. At one end of the continuum, are newspapers whose delivery costs (and in fact all their costs) are entirely funded by advertising revenues. At the other end, there are media distribution companies (e.g., satellite radio) that are almost completely funded by the consumer through subscription revenue alone. We observe that these business models have evolved over the years, and continue to evolve for each of these media (newspapers, television and radio).

We believe that as the Internet plays an increasing role in the distribution of information and entertainment services, the business models that are sustainable will also evolve. Quintessential for this evolution is flexibility in the business models that propagate information and content on the Internet. This flexibility would enable network operators and content providers to design pricing models that meet the needs of their customer base and match their business plans. This ability to support different business models depending on application/user is critical for enhancing the value that the Internet provides to the consumer and continued investment in the network infrastructure. Network Neutrality regulation at this point in an overly broad simplistic framework may unnecessarily restrict the network service provider to participate at only at one end of the continuum of business models, thus potentially limiting opportunities for investment and innovation. We conclude that maintaining flexibility in the possible network service provider-content provider business relationships that are allowed is very important as the Internet's role in media and entertainment distribution continues to evolve.

References

[1] W. Fischer and K. Meiher-Hellstern, "The Markov-Modulated Poisson Process (MMPP) cookbook," *Performance Evaluation*, vol. 18, pp. 149–171, 1992.

[2] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot, "Traffic Matrix Estimation: Existing Techniques and New Directions," in *Proceedings of ACM SIGCOMM*, 2002.

[3] Y. Zhang, M. Roughan, C. Lund, and D. Donoho, "An Information-Theoretic Approach to Traffic Matrix Estimation," in *Proceedings of ACM SIGCOMM*, 2003.

[4] "The Center for International Earth Science Information Network (CIESIN)," http://www.ciesin.columbia.edu.

[5] H. Jiang and C. Dovrolis, "Why is the Internet Traffic Bursty in Short Time Scales?," in *Proceedings of ACM SIGMETRICS*, 2005.

[6] G. Appenzeller, I. Keslassy, and N. McKeown, "Sizing router buffers," in *Proceedings of ACM SIGCOMM*, 2004.

[7] M. Yuksel, K. K. Ramakrishnan, S. Kalyanaraman, J. D. Houle, and R. Sadhvani, "Value of Class-of-Service Support in IP Backbones," in *Proceedings of IEEE IWQoS*, Chicago, IL, June 2007.

[8] F. Kelly, "Models for a self-managed Internet," *Philosophical Transactions of the Royal Society*, A358, pages 2335-2348, 2000.

[9] G. Cortese, R. Fiutem, P. Cremonese, S. D'Antonio, M. Esposito, S. P. Romano, and A. Diaconescu, "CADENUS: Creation and deployment of end-user services in premium IP-networks," *IEEE Communications Magazine*, 41(1):54–60, Jan 2003.

[10] T. Engel, H. Granzer, B. Koch, M. Winter, P. Sampatakos, I. Venieris, H. Hussmann, F. Ricciato, and S. Salsano, "AQUILA: Adaptive resource control for QoS using an IP-based layered architecture," *IEEE Communications Magazine*, 41(1):46–53, Jan 2003.

[11] F. Ciucu, A. Burchard, and J. Liebeherr, "A network service curve approach for the stochastic analysis of networks," in *Proceedings of ACM SIGMETRICS*, pages 279–290, New York, NY, USA, 2005.

[12] Network Neutrality, http://en.wikipedia.org/wiki/Network_neutrality.

[13] J. Crowcroft, "Net Neutrality: The Technical Side of the Debate: A White Paper," *ACM Computer Communications Review*, 37(1):49–55, January 2007.

[14] S. Oueslati and J. Roberts, "A new direction for quality of service: Flow aware networking," in *Proceedings of NGI*, pages 1389–1399, 2005.

[15] J. Roberts, "Internet Traffic, QoS and Pricing," in *Proceedings of the IEEE*, volume 92, pages 1389–1399, 2004.

[16] R. Braden et al., "Integrated services in the Internet architecture: An overview," *IETF RFC 1633*, Jun 1994.

[17] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An architecture for differentiated services," *IETF RFC 2475*, Dec 1998.

[18] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, and J. Rexford, "Traffic engineering for IP networks," *IEEE Network Magazine*, 14(2):11–19, Mar/Apr 2000.

[19] E. Rosen et al., "Multi-protocol label switching architecture," *IETF RFC 3031*, Jan 2001.

[20] I. Stoica, S. Shenker, and H. Zhang, "Core-stateless fair queueing: A scalable architecture to approximate fair bandwidth allocations in high speed networks," *IEEE/ACM Transactions on Networking*, 11(1):33–46, Feb 2003.

[21] B. Gleeson, A. Lin, J. Heinanen, G. Armitage, and A. Malis, "A framework for ip based virtual private networks," *IETF RFC 2764*, February 2000.

[22] S. Sahu, D. Towsley, and J. Kurose, "A quantitative study of differentiated services for the Internet," in *Proceedings of IEEE GLOBECOM*, 1999.

[23] Y. Huang and R. Guerin, "Does Over-Provisioning Become More or Less Efficient as Networks Grow Larger?" in *Proceedings of IEEE ICNP*, 2005.

[24] X. Yang, G. Tsudik, and X. Liu, "A technical approach to net neutrality," in *Proceedings of HotNets*, 2006.

[25] C. McTaggart, "Was the Internet ever neutral?" in *Proceedings of TPRC*, 2006.

[26] D. Clark, S. Bauer, W. Lehr, P. Faratin, and R. Sami, "The Growth of Internet Overlay Networks: Implications for Architecture, Industry Structure and Policy," in *Proceedings of TPRC* 2006.

[27] N. C. Stolleman, "Effects of Regulation on Investment and Innovation Incentives," in *Proceedings of TPRC* 2006.

[28] A. Chaudhuri, K. Flamm, and John Horrigan, "An Analysis of the Determinants of Internet Access," in *Proceedings of TPRC* 2005.

[29] A. M. Odlyzko, "Pricing and Architecture of the Internet: Historical Perspectives from Telecommunications and Transportation," in *Proceedings of TPRC* 2005.

[30] A. M. Odlyzko, "Internet pricing and history of communications," AT&T Labs, 2000.

[31] R.J. Edell and P. Varaiya, "Providing internet access: What we learn from index," INDEX Project Report No. 99-010W, 1999.

[32] K. Flamm, “The Role of Economics, Demographics, and State Policy in Broadband Competition: An Exploratory Study,” in *Proceedings of TPRC* 2005.
[33] P. Goyal, A. Greenberg, C. Kalmanek, B. Marshall, P. Mishra, D. Nortz, and K. K. Ramakrishnan, “Integration of Call Signaling and Resource Management for IP Telephony,” *IEEE Network Magazine*, 13(3):24–32, May/June 1999.
[34] K.-T. Chen, P. Huang, and C.-L. Lei, “How Sensitive are Online Games to Network Quality,” *Communications of ACM*, Vol. 49, No. 11, pp. 34-38, Nov. 2006.
[35] Internet World Stats, http://www.internetworldstats.com/emarketing.htm
[36] A. M. Odlyzko, “Internet traffic growth: Sources and implications,” *Optical Transmission Systems and Equipment for WDM Networking II, B. B. Dingel, W. Weiershausen, A. K. Dutta, and K.-I. Sato, eds., Proceedings of SPIE*, vol. 5247, 2003, pp. 1-15. http://www.dtc.umn.edu/~odlyzko/doc/complete.html
[37] YouTube serves up 100 million videos a day online. http://mashable.com/2006/07/17/youtube-hits-1-million-videos-per-day
[38] Global Internet Geography, TeleGeography Research, September 2006, http://www.telegeography.com/index.php
[39] IPTV World Forum. http://www.iptv-forum.com
[40] K.K. Ramakrishnan and R. Jain, “A Binary Feedback Scheme for Congestion Avoidance in Computer Networks with a Connectionless Network Layer,” *Proceedings of the ACM SIGCOMM*, August 1988.
[41] J.-C. Rochet and J. Tirole, “Two-sided Markets: A Progress Report,” *The RAND Journal of Economics*, Volume 37-3, pp. 645-667, 2006, http://idei.fr/doc/wp/2005/2sided_markets.pdf
[42] S. O’Donnell, “An economic map of the Internet,” *Proceedings of TPRC*, September 2002, http://tprc.org/papers/2002/102/InternetMap.pdf.
[43] Cisco Systems, “The Exabyte Era,” White paper, 2007, http://www.cisco.com/application/pdf/en/us/guest/netsol/ns537/c654/cdccont_0900aecd806a81a7.pdf.
[44] Cisco Systems, “Global IP Traffic Forecast and Methodology, 2006-2011”, White Paper, 2007, http://www.cisco.com/application/pdf/en/us/guest/netsol/ns537/c654/cdccont_0900aecd806a81aa.pdf.
[45] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, “You tube Traffic Characterization: A View from the Edge,” *Proceedings of Internet Measurement Conference (IMC)*, San Diego, CA, October 2007.
[46] Lehman Brothers analyst Doug Anmuth, Equity Research, Market Commentary/Strategy August 3, 2007. http://www.paidcontent.org/images/uploads/lehmanweekly.pdf
[47] Permille, http://en.wikipedia.org/wiki/Permille