Class-of-Service in IP Backbones: Informing the Network Neutrality Debate *

Murat Yuksel
University of Nevada - Reno
yuksem@cse.unr.edu

K. K. Ramakrishnan
AT&T Labs Research
kkrama@research.att.com

Shiv Kalyanaraman
Rensselaer Polytechnic Inst.
shivkuma@ecse.rpi.edu

Joseph D. Houle
AT&T
jdhoule@att.com

Rita Sadhvani
AT&T
rsadhvani@att.com

ABSTRACT

The benefit of Class-of-Service (CoS) is an important topic in the "Network Neutrality" debate. Proponents of network neutrality suggest that over-provisioning is a viable alternative to CoS. We quantify the extra capacity requirement for an over-provisioned classless (i.e., best-effort) network compared to a CoS network providing the same delay or loss performance for premium traffic. We first develop a link model that quantifies this *Required Extra Capacity (REC)*. For bursty and realistic traffic distributions, we find the REC using ns-2 simulation comparisons of the CoS and classless link cases. We use these link models to quantify the REC for realistic network topologies. We show that REC can be significant even when the proportion of premium traffic is small, a situation often considered benign for the over-provisioning alternative.

Categories and Subject Descriptors:
C.4 [Computer Systems Organization]: Performance of Systems
C.4 [Computer-Communication Networks]: Network Architecture and Design

General Terms: Performance

Keywords: Network neutrality, Class-of-service, Economics, Performance

1. INTRODUCTION

Currently there is a wide ranging debate on the issue of "network neutrality" [1]. One key technical aspect of the network neutrality debate is whether best-effort application traffic should be carried along with other (so-called "premium") traffic for which SLA commitments have been made (or are expected, either explicitly or implicitly) without differentiation. Some network neutrality proponents, at one end of the opinion spectrum, suggest that there should be no differentiation of traffic and all performance requirements should be met by overprovisioning 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 overprovision the network?

*This work is supported in part by AT&T and NSF contracts 0721600 and 0721609.

SIGMETRICS'08, June 2–6, 2008, Annapolis, Maryland, USA.
ACM 978-1-60558-005-0/08/06.

Our study focuses on this specific question. We compare a classless network which is over-provisioned against an engineered network using per-class queuing to offer Class-of-Service (CoS) (i.e., differentiated service) and meet user expectations and SLAs.

The hypothesis of this paper is that an over-provisioned single-queue network service for meeting the SLAs of performance-sensitive traffic and regular best-effort traffic is inefficient (from a capacity viewpoint) compared to an engineered network offering *simple* 2-queue CoS differentiation. Though this basic fact is well-known, our paper refines it to identify parametric regions where this inefficiency exists and is pronounced. We estimate the *required extra capacity (REC)* for a classless link to match the performance (in delay or loss) provided by its CoS-based correspondent. We show that *the inefficiency, as measured by the REC, is significant even for moderate utilizations and particularly for low fractions of premium traffic*. We generalize the single link model to an ISP network taking into account the network topology, traffic matrices (based on a gravity model), and shortest path routing. Based on the quantitative modeling of Sprintlink ISP topology obtained from Rocketfuel [4], we provide estimates of the REC. In [5], we briefly introduced the basic insights for quantifying REC at the *link* level. In this paper, we focus on extensively studying REC within a *network* setting. Major contributions and results of this paper are as follows:

- *Quantification of REC in a network setting.*
- *REC is higher when the proportion of traffic using the premium class is smaller*. Our network model analysis lends insight into the *significant REC needed even when the proportion of premium traffic is small*.
- *Effect of long-range dependent (LRD) traffic*. We show that traffic patterns with more realistic burstiness, as in LRD traffic, cause REC to increase by orders of magnitude in comparison to short-range dependent traffic like a Markov-Modulated Poisson Process (MMPP) [2]. LRD models are known to approximate the real Internet traffic at various time-scales, although they are more applicable on longer time-scales [3].

We consider two traffic classes on a CoS link: a *premium class* and a *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 (neutral) link to achieve the same performance target for both traffic classes. Let the aggregate traffic rate be λ_D to be served by a CoS link with a capacity of μ_D. Also let a fraction g of this aggregate traffic be premium class traffic - a rate of $\lambda_{Prem} = g\lambda_D$, while the remaining is best-effort (BE) class traffic with a rate of $\lambda_{BE} = (1-g)\lambda_D$. For the premium class traffic, we define a performance target ζ, in terms of delay or loss. We first formulate the necessary classless link capacity μ_N to achieve the same performance target ζ for the aggregate traffic λ_D.

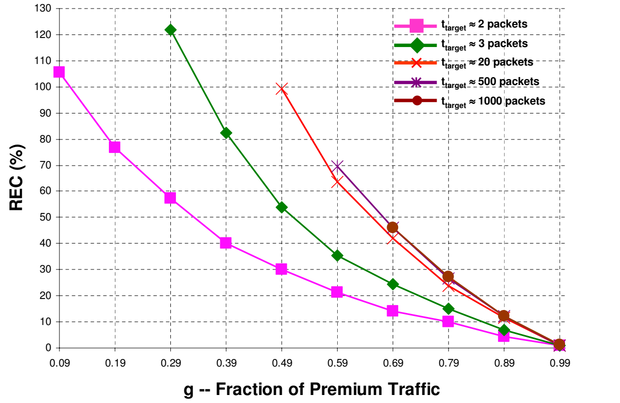

Figure 1: MMPP/M/1 link model with t_{target} in normalized units of packets, e.g., "1000 packets of delay" equals to 8.192ms delay over a 1Gbps link carrying 1KB packets.

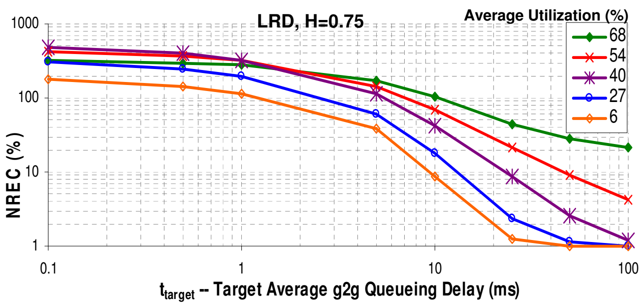

Figure 2: Sprintlink g2g delay – NREC against the target average g2g queueing delay under LRD or MMPP traffic.

From this, we can calculate REC for the CoS link in terms of *rate* as $\mu_N - \mu_D$ (or as a *percentage* $100(\mu_N/\mu_D - 1)$). In this model, we use average delay t_{target} or average loss probability p_{target} as the performance target. Figure 1 shows the how link REC behaves as the proportion of premium traffic g increases for fixed delay targets. With loss, an additional parameter is the buffer size, which we consider to be K packets for each of the traffic classes in the *CoS link* and $2K$ for the aggregate traffic in the *classless link*.

2. NETWORK MODEL AND RESULTS

We generalize our single link model to a network model, reflecting a typical ISP's backbone. Crucial components of a network model include a realistic *topology* and a realistic *traffic matrix*. Given the topology and the traffic information, we then calculate REC for the complete network. We call this required extra capacity for the complete network as "*network REC*"(*NREC*). We first calculate a routing matrix R for the ISP network from the link weight information. With a realistic traffic matrix T, we can then calculate the traffic load pertaining to individual links by taking the product of T and R. For each of these link traffic loads, we apply the link model described earlier and find the RECs for each individual link. Finally, we calculate NREC by averaging the individual link RECs across all links of the network. Briefly, $NREC$ expresses the *average extra capacity needed on each link* of the network.

We calculate NREC both for MMPP and LRD traffic by performing a lookup of the link model simulation result [5] for a given utilization and performance target. The MMPP traffic allows us to observe NREC values under well-behaved and short-range-dependent traffic with very conservative burstiness behavior, and LRD to see how much larger REC would become with more bursty and realistic traffic. In both cases, we use *conservative* parameters for the traffic burstiness, i.e., $r = 4$ for MMPP (where r is the traffic rate ratio of the two states), and Hurst parameter=0.75 for LRD. Real

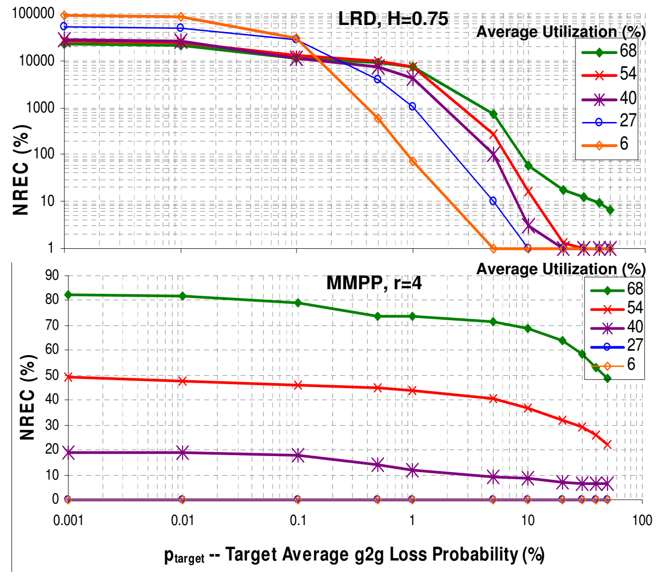

Figure 3: Sprintlink g2g loss – NREC against the average g2g loss probability under two kinds of traffic: LRD vs. MMPP. The buffer size at each link of the network is K=100ms.

IP traffic is often more bursty than this [3]. Also, when loss probability is the performance goal, we use a buffer size of $K = 100ms$, which is conservative in comparison to conventional buffer sizes.

Figures 2 and 3 show the NREC values for the Sprintlink topology for five levels of traffic load. Fig. 3 includes results for both LRD and MMPP traffic. It is clear that NREC increases as the average link utilization increases, especially when the target average g2g queueing delay is smaller (Fig. 2). Also, as expected, LRD traffic results in an order of magnitude larger NRECs in comparison to the case with MMPP traffic. For example, for a g2g queueing delay target of 5ms and a 40% utilized Sprintlink network, NREC under MMPP traffic is about 20% while it is about 100% with LRD traffic. This difference becomes more evident when the target g2g queueing delay is smaller. Figure 3 shows NREC when the performance target is the g2g loss probability. Again, it is clear that NREC increases as the average link utilization increases.

3. SUMMARY AND DISCUSSION

We have quantified 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 diff-serv network. We used IP backbone topologies from Rocketfuel along with a careful and rigorous procedure for synthetically generating traffic matrices based on relative user populations while ensuring link capacities are sufficient to support the traffic. We observed that NREC increases with the average utilization of the links in the network and when the relative proportion g of the premium traffic *is smaller*. Moreover, NREC grows rapidly as the acceptable delay and packet loss targets become tighter (smaller) and as the traffic becomes burstier.

4. REFERENCES

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

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

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

[4] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP Topologies with Rocketfuel. In *Proc. of SIGCOMM*, 2002.

[5] M. Yuksel, K. K. Ramakrishnan, S. Kalyanaraman, J. D. Houle, and R. Sadhvani. Value of Class-of-Service Support in IP Backbones. In *Proc. of IEEE IWQoS*, 2007.