

The Local and Global Effects of Traffic Shaping in the Internet

Massimiliano Marcon[†], Marcel Dischinger[†], Krishna P. Gummadi[†], and Amin Vahdat[‡]

[†]MPI-SWS and [‡]University of California, San Diego

ABSTRACT

The past few years have seen a dramatic increase in the Internet traffic due to popular bulk data transfer applications, such as BitTorrent. Many access ISPs are deploying traffic shaping to limit the peak network traffic on their transit links, and thereby reduce their wide-area bandwidth costs. In this paper, we show that an ISP can substantially reduce (by a factor of 2 or more) its peak transit link usage by traffic shaping a small fraction of its largest flows, while incurring a minimal penalty on the completion times of these bulk flows. Unfortunately, our analysis also shows that if many ISPs traffic shape flows based on their local interest, the end-to-end performance of the bulk flows will be seriously harmed.

Categories and Subject Descriptors: C.2.3 [Computer-Communication Networks]: Network Operations

General Terms: Design, Economics, Performance

Keywords: Bulk Data Transfers, Traffic Shaping

1. MOTIVATION

The Internet is witnessing explosive growth in demand for bulk content. Examples of bulk data transfers include downloads of music and movie files and distribution of large software amongst others. Recent studies of Internet traffic [1, 8] show that such bulk transfers account for a large and rapidly growing fraction of bytes transferred across the Internet.

The bulk data traffic in the Internet today represents just the tip of the iceberg. Tremendous amounts of digital data are being delivered outside the Internet, for example using hard drives, optical media, or tapes [5, 6, 9], because it is cheaper and faster – though usually not more convenient or secure – than using the Internet. For example, on an average day, Netflix ships 1.6 million movie DVDs [9], or 6 petabytes of data. This is more than the estimated traffic exchanged between ISPs in the USA [10]. It is debatable whether the Internet can ever match the capacity of postal networks. However, the convenience of online transfers is likely to drive the demand to deliver more bulk data over the Internet in the foreseeable future.

Internet bulk data transfers are expensive. A recent study [7] reported that ISPs (or CDNs) charge large content providers, such as YouTube and MSN Live, 0.1 to 1.0 cent per minute for a 200-400 kbps data stream. And higher bandwidth streams will cost even more. The high cost of wide-area network traffic means that increasingly *economic* rather than *physical* constraints limit the performance of




Figure 1: Traffic traversing a university's access link: The traffic shows clear diurnal patterns, with large flows contributing to most of the network bandwidth consumption.

many Internet paths. That is, even when there is plenty of physical capacity available on a given link, ISP policies of charging customers based on peak bandwidth utilization (often measured by the 95th percentile over some time period) result in strong disincentives to approach the full physical capacity of inter-ISP links.

While decades of research in congestion control shows how to manage transfers across physical bottlenecks, there is little understanding of how to manage transfers across economic bottlenecks. Instead, ISPs have developed a variety of ad hoc traffic shaping techniques to control bandwidth costs. This traffic shaping typically targets bulk transfers as they consume the majority of network bytes. However, the techniques are often blunt and arbitrary, and often shut down entire applications (e.g., P2P file sharing [3]) without a sophisticated understanding of the resulting economic benefits. Against this backdrop, we present a systematic analysis of the local benefits of deploying traffic shaping within an ISP, and the global impact wide-spread deployment of traffic shaping has on the performance of bulk flows.

2. ANALYSIS AND FINDINGS

We begin by estimating the potential benefits of traffic shaping. We studied real network traces collected at the access links of universities in the US and Canada. We make two key observations: first, as shown in Figure 1, a small fraction (< 0.5%) of the largest flows are responsible for most of the bytes (> 65%) transferred, and these bulk flows contribute significantly to the peak bandwidth. Second, bandwidth consumption shows a strong diurnal pattern with high peak-to-trough ratios. Our analysis reveals a significant opportunity for intelligent traffic shaping that observes economic incentives and minimizes the peak levels of bandwidth consumption.




Figure 2: Shaped traffic over the same access link: Rate limiting bulk transfers results in substantial peak bandwidth reduction.

We evaluated a composition of traffic shaping and priority queuing techniques that limit the bandwidth consumed by bulk transfers during times of peak utilization. The basic idea is to shift bulk data traffic to periods of low network usage and smooth out bandwidth consumption over the course of the day. Our techniques distinguish between interactive and bulk flow categories based on the flow size. We leave the interactive flows untouched but rate limit the bulk flows once the traffic exceeds a certain threshold. Simulation-based analysis of our proposed techniques shows that they achieve a significant reduction ($> 60\%$) in peak bandwidth (Figure 2), while minimally impacting the completion times of individual bulk transfers. In contrast, we find that naive traffic shaping schemes can dramatically slow or even terminate many targeted flows.

Our results indicate that it will be in the best interest of many ISPs to perform variants of the proposed traffic shaping techniques. Unfortunately, we find that once a significant portion of ISPs perform such *local* traffic shaping, the *global* system behavior degrades significantly. With increased adoption of traffic shaping, bulk transfer performance degrades as the end-to-end distance the transfer travels grows longer. Essentially, differences in the peak transfer times of ISPs in different time zones means that the farther a flow travels (literally), the higher the probability that at least one ISP will throttle the flow at any given time.

We analyzed the performance of a bulk flow traversing the transit links of two ISPs that do traffic shaping to maximize their own benefits. In the case when both the ISPs' transit links are in close proximity, the throughput of the bulk flow decreased by a factor of two compared to traffic shaping at only one ISP. Even worse, when both ISPs are in distant time zones, the bulk flow suffered a 20 fold drop in its throughput. Even with moderate distances, we find that bulk transfers become constantly throttled to the point of delivering largely no utility.

3. IMPLICATIONS

Moving forward, if the trends predicted in this paper hold, many bulk transfers will essentially receive no bandwidth. This limitation would come at a time when the demand for bulk transfers is exploding, consider high-definition video downloads or large scientific data sets. In this context, we will require alternative bulk-transfer architectures that at least consider the economic incentives that led to the traffic shaping in the first place.

One approach would be to keep the traffic local, i.e., within the same ISP. Akamai already offers such a service for client-server workloads and recent efforts like P4P [11] and Ono [2] try to reduce cross-ISP traffic for P2P workloads by biasing neighbor selection to peers within the same ISP.

Another possibility is for ISPs to stop charging for peak levels of utilization but to instead adopt a different pricing model, e.g., per-byte accounting. Unfortunately, this model is likely to result in additional imbalances because it does not recognize that “all bytes are not created equal”. Not encouraging data to be sent during times of otherwise slack usage means that network resources that must still be provisioned for peak demand sit idle. More importantly, per-byte charging would introduce even larger incentives for ISPs to more aggressively traffic shape bulk traffic.

Another approach would be to develop an alternative, incentive-compatible protocol for bulk transfers. While such a protocol is beyond the scope of this work, we outline some high-level possibilities. First, we observe that bulk transfers may still perform well as long as they are subject to traffic shaping by only a single ISP. Next, bulk transfers do not require much of the semantics of TCP, e.g., in order delivery or synchronous end-to-end data acknowledgment. The unique semantics of bulk data flows allow for novel protocols that aim at providing both low bandwidth cost and good performance. We think these protocols could take inspiration from delay-tolerant networks [4] and postal networks that stage the delivery of transfers from point to point in the network. For instance, postal networks often take advantage of capacity as it becomes available to move data towards the destination. Similarly, delay tolerant networks leverage in-network storage to buffer data until connectivity becomes available. One could imagine analogs where data is buffered in the network until traffic throttling abates.

4. REFERENCES

- [1] K. Cho, K. Fukuda, H. Esaki, and A. Kato. The Impact and Implications of the Growth in Residential User-to-user Traffic. In *Proc. of SIGCOMM*, Sep 2006.
- [2] D. R. Choffnes and F. E. Bustamante. Taming the Torrent: A practical approach to reducing cross-ISP traffic in P2P systems. In *Proc. of SIGCOMM*, Seattle, WA, Aug 2008.
- [3] Packet Forgery By ISPs: A Report on the Comcast Affair. http://www.eff.org/files/eff_comcast_report2.pdf.
- [4] K. Fall. A Delay Tolerant Networking Architecture for Challenged Internets. In *Proc. of SIGCOMM*, 2003.
- [5] Google helps terabyte data swaps, Mar 2007. <http://news.bbc.co.uk/2/hi/technology/6425975.stm>.
- [6] J. Gray, W. Chong, T. Barclay, A. Szalay, and J. vandenBerg. TeraScale SneakerNet: Using Inexpensive Disks for Backup, Archiving, and Data Exchange. Technical Report MSR-TR-2002-54, Microsoft Research, May 2002.
- [7] C. Huang, J. Li, and K. W. Ross. Can Internet Video-on-Demand be Profitable? In *Proc. of SIGCOMM*, Aug 2007.
- [8] T. Karagiannis, A. Broido, N. Brownlee, K. Claffy, and M. Faloutsos. Is P2P Dying or just Hiding? In *Proc. of IEEE Globecom*, 2004.
- [9] Netflix Online Movie Rental. <http://www.netflix.com>.
- [10] A. M. Odlyzko. Internet Traffic Growth: Sources and Implications. In *Proc. Optical Transmission Systems and Equipment for WDM Networking II*, 2003.
- [11] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz. P4P: Portal for (P2P) Applications. In *Proc. of SIGCOMM*, Seattle, WA, Aug 2008.