Delay Tolerant Bulk Internet Transfers*

Nikolaos Laoutaris and Pablo Rodriguez

Abstract

From its conception the Internet has been a communication network and therefore its development has been largely driven by the overarching assumption that connections and data transfers are sensitive to delay. Therefore, spatial optimization, in the form of routing, has been the main tool for improving the services offered by the network. Temporal optimization, in the form of scheduling, has been limited to millisecond-second scales, to be aligned with the requirements of interactive, delay sensitive traffic. In late years, however, the network has been progressively sifting from communication to content dissemination. Unlike communication, content dissemination can sometimes tolerate much larger delays, *e.g.*, in the order of hours. Such higher tolerance to delay allows scheduling to go beyond congestion avoidance. In this short note we exemplify how to use store-and-forward scheduling to service bulk data transfers that are impossible, or too expensive under current pricing schemes for bandwidth.

Bulk data: Residential and corporate bulk data have fueled an unprecedented increase in overall Internet traffic over the last few years. On the end user side, these include high definition movies from commercial web-sites or P2P networks [2], large scale software updates [4], and remote backups. Adding to the above, datacenters hosting cloud computing applications exchange large amounts of synchronization, accounting, and data-mining traffic, while large corporate and government organizations contribute increasing amounts of economic, engineering, and scientific datasets. At this rate, it becomes increasingly difficult for the ISP and the manufacturers of networking equipment to keep up with the generated traffic without requiring substantial new investments that are hard to find in the competitive ISP market. Thus, the not so distant euphoric belief of infinite network capacity, triggered largely by the fiber glut of the 90's, has been quickly substituted by headlining news on the discriminatory practices of ISPs against P2P traffic [10]. Next we describe two common cases of bulk data bottlenecks and then highlight some basic thoughts on how to solve them through storage-enabled store-and-forward scheduling that takes advantage of the delay tolerant nature of bulk data.

Bulk bottlenecks under flat-rate pricing: Several ISPs throttle [10, 3] P2P traffic from flat-rate residential customers during peak hours in order to

*Telefonica Research, Barcelona. Email: {nikos,pablorr}@tid.es

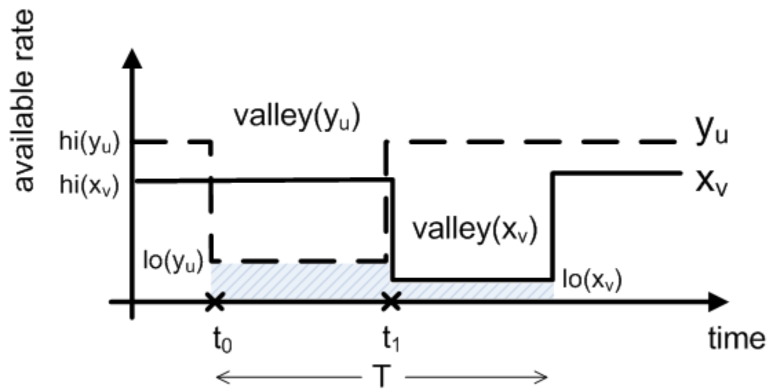

Figure 1: E2E transfers between a sender in ISP v and receiver in a remote ISP u having different peak hours can be constantly bottlenecked in $(t_0, t_0 + T]$ due to alternating long lived bottlenecks in the available uplink rate x_v and downlink rate y_u (valley(x_v) and valley(y_u), respectively).

free capacity for interactive traffic which is valued higher by most users. Currently, it is mostly popular P2P applications like BitTorrent being targeted, but in the future it may be any other bulk application that becomes popular among flat-rate residential customers. Such throttling bottlenecks introduced by cascades of traffic shaping devices can severely impact flows across multiple ISPs with different throttling times. In Fig. 1 we give an example of a difficult case of combined throttling using hypothetical uplink and downlink rates of a sender and a receiver at different access ISPs. In particular, we observe that by chaining such throttling bottlenecks, the combined receiver (valley(u)) and sender throttling (valley(v)) results in a small transferred volume indicated by the shaded area behind the two rates. Notice here that had it not been for the combined effect, each throttling behavior alone would have allowed for much higher volumes (individual areas under the solid and dashed lines). In the extreme case of sender and receiver pairs with long, non-overlapping valleys, the transfer could be throttled across the day. Situations like the above can occur with end points at remote time-zones, or within the same time-zone but on ISPs of different type, *e.g.*, a residential access ISP peaking in the evening and a corporate access ISP peaking at noon.

Bulk bottlenecks under percentile pricing: Similar problems can arise under 95-percentile pricing [5] often applied to corporate customers or hosting services which pay based on (nearly) peak usage. Such pricing is justified by the fact that the cost of networking equipment depends on the maximum load it has to carry under given QoS. Given that customers pay according to peak traffic, and granted that load typically exhibits strong diurnal patterns [9, 6], leaves much already-paid-for offpeak capacity to be used for sending additional bulk data at no extra cost. Nonetheless, as before, time-zone differences can step in the way. For example it might be impossible to use the capacity during the load valley in the early morning hours to send bulk data to a receiver on a distant time-zone that is currently going through its evening peak hours. The end result could be additional transit costs as the bulk flow cannot avoid increasing the peak load.

Store 'n Forward to the rescue: As such bottlenecks (either due to pricing or throttling) become more prevalent, we argue that to restore the performance of bulk transfers and minimize transmission costs will require a new *Store 'n Forward, (SnF)* service based on *temporal redirection* techniques. Existing *spatial redirection* techniques like native and overlay routing perform path selections over short periods to avoid bottlenecks at Internet's core [1] but have no way of escaping complex accumulated constraints that can occur further in the future via the combination of various bottlenecks at different time-zones. To solve the problem, we propose breaking E2E flows into smaller segments and perform SnF scheduling through intermediate storage nodes to achieve the best utilization of the capacity available between two end points for long-lived bulk transfers. Storage nodes decouple the end point constraints: when a sender is not throttled (charged) then it uploads fast (cheap) to a storage node where data accumulates until the receiver can also download it fast (cheap). In [8] we used data from a large transit provider to show that SnF can provide Tbyte-sized daily bulk transfers at low cost, or for free, whereas end-to-end connection transfers, and even parcel delivery services would incur much higher cost. In [7] we have also argued for the benefits of SnF scheduling for residential broadband users. These two works provide some basic initial hints about what we believe to be a promising new field of research on delay tolerant networks.

References

[1] Aditya Akella, Srinivasan Seshan, and Anees Shaikh. An empirical evaluation of wide-area internet bottlenecks. In *Proc. of ACM SIGMETRICS '03*, pages 316–317, 2003.

[2] Kenjiro Cho, Kensuke Fukuda, Hiroshi Esaki, and Akira Kato. The impact and implications of the growth in residential user-to-user traffic. In *Proc. of SIGCOMM'06.*

[3] Marcel Dischinger, Alan Mislove, Andreas Haeberlen, and Krishna P. Gummadi. Detecting bittorrent blocking. In *Proc. ACM IMC'08.*

[4] Christos Gkantsidis, Thomas Karagiannis, Pablo Rodriguez, and Milan Vojnovic. Planet scale software updates. In *Proc. of ACM SIGCOMM '06*, Pisa, Italy.

[5] David K. Goldenberg, Lili Qiuy, Haiyong Xie, Yang Richard Yang, and Yin Zhang. Optimizing cost and performance for multihoming. In *Proc. of ACM SIGCOMM '04.*

[6] Anukool Lakhina, Konstantina Papagiannaki, Mark Crovella, Christophe Diot, Eric D. Kolaczyk, and Nina Taft. Structural analysis of network traffic flows. In *Proc. of ACM SIGMETRICS / Performance 2004*.

[7] Nikolaos Laoutaris and Pablo Rodriguez. Good things come to those who (can) wait or how to handle delay tolerant traffic and make peace on the Internet. In *Proc. of HotNets-VII*, 2008.

[8] Nikolaos Laoutaris, Georgios Smaragdakis, Pablo Rodriguez, and Ravi Sundaram. Delay tolerant bulk data transfers on the Internet. In *Proc. of ACM SIGMETRICS'09*.

[9] Matthew Roughan, A. Greenberg, C. R. Kalmanek, Michael Peter Rumsewicz, J. Yates, and Y. Zhang. Experience in measuring internet backbone traffic variability: Models metrics measurements and meaning. In *Proc. of ITC-18*, 2003.

[10] Slashdot. Comcast Hinders BitTorrent Traffic, Aug 2007.