Modeling of Aggregate Available Bandwidth in Many-to-One Data Transfer

S. C. Hui and Jack Y. B. Lee
{schui3, yblee}@ie.cuhk.edu.hk
Department of Information Engineering
The Chinese University of Hong Kong
Hong Kong

Abstract – **This work investigates the modeling of aggregate available bandwidth in multi-sender network applications. Unlike the well-established client-server model, where there is only one server sending the requested data, the available bandwidth of multiple senders when combined together does exhibit consistent properties and thus can be modeled and estimated. Through extensive experiments conducted in the Internet this work proposed to model the aggregate available bandwidth using a normal distribution and then illustrates its application through a hybrid download-streaming algorithm for video delivery. This new multi-source bandwidth model is especially suitable for the emerging peer-to-peer applications, where having multiple sources available are the norm rather than the exception.**

Keywords **– multiple-sender transmission, network measurement**

I. INTRODUCTION

Today's Internet only provides best-effort data delivery and so does not guarantee bandwidth availability. While the best-effort model works well for data applications such as the WWW and email, it presents significant challenges to bandwidth-sensitive applications such as video streaming.

Specifically, to successfully stream a video we need to ensure that the video bit-rate does not exceed the network bandwidth available, or else the client will run into buffer underflow, leading to playback hiccups. Unfortunately the available network bandwidth between a sender and a receiver is not known a priori and worst, often varies from time to time.

Ideally, if the bandwidth availability can be accurately modeled by a random process, then the sender can simply select a video bit-rate such that performance can be guaranteed probabilistically. However, as we will show in Section III, modeling the bandwidth availability for a single sender is very difficult, if not impossible.

Given the limitation of this single-sender approach, researchers have begun to investigate approaches employing multiple senders [1-3] to exploit three potential benefits: (a) increasing the throughput by combining the bandwidth of multiple senders; (b) adapting to network bandwidth variations by shifting the workload among the multiple senders; and (c) reducing bursty packet loss by splitting the data transmission among the multiple senders.

In this work we go one step further to argue that if there are sufficient numbers of independent senders, we not only can achieve higher throughput, but also be able to model the aggregate available bandwidth as a normal distribution according to the Central Limit Theorem. We verify this conjecture experimentally by conducting streaming experiments in the global PlanetLab testbed [4]. Our experimental results strongly suggest that this model is applicable in the current Internet and thus, can be used for designing multi-sender streaming protocols that supports probabilistic performance guarantees [5].

This work has three contributions. First, to the best of our knowledge, this is the study investigating the modeling of aggregate available bandwidth of multiple senders. Second, this is the first study to report experimental results to show that the aggregate available bandwidth is normally distributed, and under what conditions. Finally, this discovery opens a new way to providing probabilistic performance guarantees in bandwidth-sensitive applications such as video streaming even in the current best-effort Internet.

II. BACKGROUND AND RELATED WORK

Modeling of network traffic has been studied extensively in the literature. It is generally accepted that the Internet traffic cannot be adequately modeled by simple models such as a Poisson process [6]. A number of studies showed that network traffic is in fact self-similar [7-10], exhibiting long-range dependency with heavy-tailed distribution. There are many other traffic models proposed in the last decade but due to space limitation they will be not reviewed here.

It is worth noting that the abovementioned studies primarily focused on modeling properties of the network traffic itself. By contrast, our work focuses on the modeling of the bandwidth available for streaming media data in an *end-to-end* manner. In particular, our measurements include the effects of network link capacity, competing traffics, limits and variations of the sender itself (e.g., due to other concurrently running applications), as well as dynamics of the transport protocol (e.g., TCP).

Not surprisingly, with so many system factors in the equation the resultant bandwidth availability between a sender and a receiver can vary significantly across

This research is funded in part by a Direct Grant, an Earmarked Grant (CUHK4211/03E) from the HKSAR Research Grant Council, and the UGC Area of Excellence in Information Technology Scheme (AoE/E-01/99).

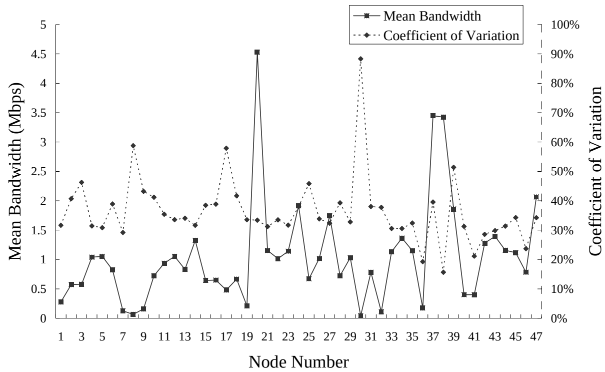

Fig. 1. Average bandwidth and Coefficient-of-variations of the 47 senders.

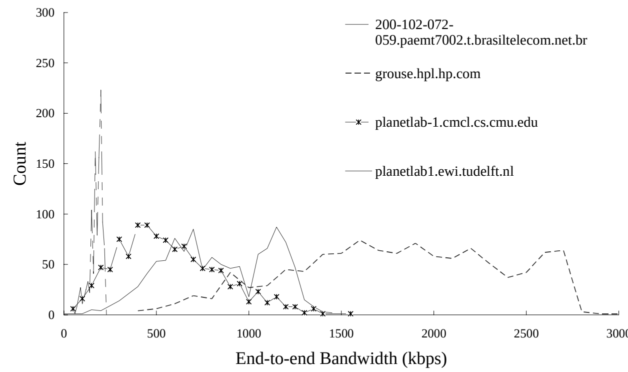

Fig. 2. End-to-end bandwidth distribution, sample from (a) planetlab-1.cmcl.cs.cmu.edu (b) planetlab1.ewi.tudelft.nl (c) grouse.hpl.hp.com (d) 200-102-072-059.paemt7002.t.brasiltel ecom.net.br

different senders and as a result does not conform to any consistent models. On the other hand, if we combine the available bandwidth of multiple senders, then the aggregate available bandwidth will become far more consistent.

Specifically, let X_i $\{i \in 1..N\}$ denotes a set of N independent random variables representing the available bandwidth from sender i to the receiver. Assume each X_i to have an arbitrary probability distribution with finite mean μ_i and finite variance σ_i^2. Then according to the Central Limit Theorem (CLT), the combined available bandwidth has a limiting cumulative function which approaches a normal distribution. Note that for the CLT to be applicable, we need to ascertain that the X_i's are independent, i.e., the senders' available bandwidths are not correlated. We investigate this issue in Section IV by computing the correlation coefficient [11] of different senders' bandwidths.

III. Single-Source Bandwidth Availability

A. Measurement Methodology

To obtain realistic results it is necessary to conduct experiments in the Internet rather than in a simulator or a closed test-bed. Therefore we conducted all experiments in the PlanetLab [4] global test-bed which has hundreds of hosts residing in many different countries around the world connected through the Internet. For the actual bandwidth measurement we used the Iperf [12] tool, which can measure the network throughput averaged over a given period of time. A total of 47 different hosts in PlanetLab are employed in the experiments. We manually removed hosts local to the receiver host to prevent overflowing the receiver and skewing the results. We also tested the receiver's throughput to ensure that the local network and the receiver will not become the bottleneck in the measurements.

We installed the Iperf server in the 47 sender hosts, and let the receiver connect to the sender to initiate data transmission. We measure the bandwidth availability by sending data using TCP from the senders to the receiver. The senders all send data as fast as TCP will allow. The receiver captures the average throughput for each source once every 10 seconds (the default setting in Iperf). The measurement lasts for 3 hours, which generated 1,080 measurement samples.

Although Iperf also supports the use of UDP in measurements, we choose TCP for two reasons. First, sending UDP datagrams at very high data-rate will likely cause serious network congestion and affect other users. Second, even for video streaming it is desirable to keep the video data traffic TCP-friendly to minimize impact to other traffic flows. Thus we employed TCP instead of UDP in the bandwidth measurements and the results should also be applicable to other TCP-friendly protocols (e.g., TFRC [13], etc.).

B. Measurement Results

We first examine the throughput of individual senders. Fig. 1 plots the mean throughput and the coefficient-of-variation (CoV) of the 47 senders. We can observe that the bandwidth availability of the 47 senders varies substantially from a minimum of 0.04 Mbps to a maximum of 4.53 Mbps. Moreover, the senders' temporal bandwidth variations, represented by their CoV, also vary substantially across different senders, ranging from 0.16 to 0.88.

Fig. 2 plots the bandwidth distribution of 4 out of the 47 senders, over the measurement period of 3 hours. We observe that their distributions also vary substantially from one sender to another and do not conform consistently to any known distributions. These results clearly illustrate the difficulty in estimating and modeling the bandwidth availability of individual senders.

IV. Multi-Source Bandwidth Availability

While the properties of individual senders are difficult to model and predict, the properties of the aggregate bandwidth of multiple senders are far more consistent. We examine in this section the characteristics of aggregate available bandwidth from multiple sources using the methodology in Section III.

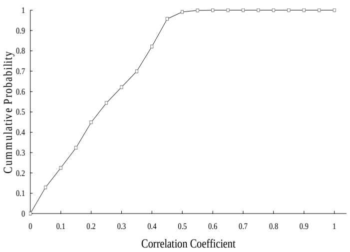

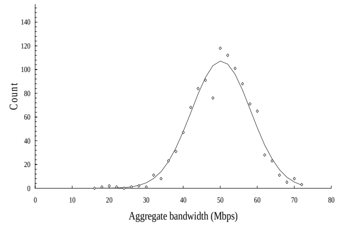

Fig. 3. (Left) Correlation of senders' bandwidth availability.
Fig. 4. (Right) Distribution of aggregate bandwidth of 47 senders.

Fig. 3 plots the cumulative distribution for the *correlation coefficient* [11] of the 47 sending nodes in the PlanetLab. The correlation coefficient measures the degree of correlation between the bandwidth availability of two senders. The result shows that half of the sender-pairs have a correlation coefficient less than 0.2 while the correlation coefficients of all sender-pairs are less than 0.6. This suggests that the bandwidth availability of most nodes is relatively uncorrelated. Therefore if we treat the bandwidth availability of each sender as a random variable, we will expect the sum of these random variables and hence the aggregate available bandwidth to approach the normal distribution.

This is confirmed in Fig. 4 which plots the distribution of the aggregate bandwidth of all 47 senders as well as the normal distribution with the same mean and variance as the measurement samples. By inspection we can see that the empirical distribution closely follows the normal distribution. To further quantify the similarity, we apply the Shapiro-Wilk test [14] which computes from the measurements a p-value to quantify the measurements' conformity to the normal distribution. The range of the p-value is from 0 to 1, with larger values representing better conformity to the normal distribution. For example, a p-value of 0.05 represents a 95% confidence level and is generally considered to be conformance to normal.

To further investigate the effect of the number of senders on normal-conformance, we vary the number of senders from 2 to 45 and plot the resultant p-value in Fig. 5. Note that each data point is computed from the average of up to 1,000 different combinations of senders.

There are three observations. First, as expected the p-value increases with the number of senders (up to 18 senders). Second, we also note that beyond 18 senders the p-value actually decreases slightly. This is an artifact of the Shapiro-Wilk test as the test is more sensitive to non-conformity when there are more samples. Third, using the p-value threshold of 0.05 as a threshold for normal-conformity [14], the results show that the measured distribution becomes normally distributed even when there are only 4 senders. This suggests that even with only a few senders, we can still approximate the aggregate available bandwidth using the normal distribution.

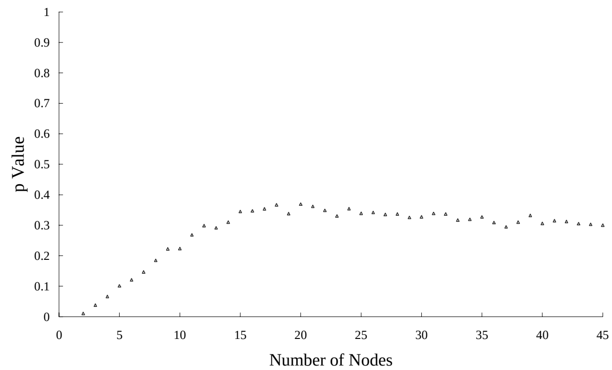

Fig. 5. The effect of the number of senders on normal conformance.

V. Applications

The significance of multi-source streaming is that the aggregate available bandwidth can be described by the normal distribution despite the fact the underlying Internet is a best-effort network. This discovery enables one to build content delivery systems with probabilistic performance guarantees which can improve the quality of service to end users.

A. Hybrid Download-Streaming

As an illustration we consider the typical downloading of media content (e.g., MPEG video) from a web server for local playback at the client. While streaming can reduce the latency significantly, streaming may not be possible if the available bandwidth is lower than the media data rate. In that case the only option is to first download the media file completely before playback to ensure that playback will be continuous, or else risks frequent playback interruptions which can be very annoying.

However, if the media file is to be downloaded from multiple web servers (with the data properly divided across the servers), then the aggregate data transfer rate will exhibit the normal distribution as demonstrated earlier. This enables us to estimate the data transfer time and thus start the playback process even before the download is completed, and still be able to guarantee (probabilistically) continuous playback.

Specifically, let C_i be the aggregate data transfer rate at time interval i after the beginning of the downloading process and assume playback begins w intervals after download begins. For simplicity, the length of a time interval can be chosen to be the same as the measurement interval as explained in Section III.

To ensure that playback is continuous, we then need to ensure that the total amount of media data received at any time interval i, denoted by A_i, must not be lower than the total amount consumed by playback, denoted by B_i, i.e.,

$$A_i \ge B_i, \forall i \ge 0 \qquad (1)$$

where $A_i = \sum_{j=1}^{i} C_j$ and $B_i = \begin{cases} R(i-w), & \text{if } i > w \\ 0, & \text{otherwise} \end{cases}$

Substituting the definition of A_i and B_i into (1) and rearranging we can obtain

$$\sum_{j=1}^{i} C_j \geq R(i-w), \quad \forall i > w \tag{2}$$

Now as C_j's are normally-distributed random variables, the sum of n C_j's will also be normally distributed, and is described by the n-times autoconvolution of C_j's CDF $F(x)$, denoted by $F^{(n)}(x)$.

Hence to guarantee (2) with a probability of Δ, we need to ensure that

$$F^{(n)}(R(n-w)) \geq \Delta \tag{3}$$

Thus the earliest time for playback to begin can be obtained from

$$w = \min\left\{v \mid F^{(n)}(R(n-v)) \geq \Delta, \forall n \geq v\right\} \tag{4}$$

B. Performance Evaluation

To evaluate the latency reduction achievable by the previously discussed hybrid download-streaming algorithm, we conducted trace-driven simulations using traffic traces gathered from the PlanetLab to obtain the playback latencies for three algorithms: (a) *pure download* – begin playback only after video file is downloaded completely; (b) *hybrid download-streaming*; and (c) *lower bound* of the download time.

We obtain the lower bound from a priori knowledge of the traffic traces, i.e., all the C_i's are assumed to be known a priori. Obviously this algorithm is not realizable in practice and is thus included for comparison only.

We run 10 different experiments each with a different traffic trace collected from PlanetLab. The movie length and its bit rate range from 500 seconds to 1,000 seconds and 200 kbps to 300 kbps respectively. In each experiment, there are 5 to 10 servers serving disjoint subsets of the video file and the mean aggregate bandwidth available is lower than the video bit-rate (i.e., conventional streaming is not feasible).

Fig. 6 plots the playback latency versus different traffic traces for the three algorithms. As expected, the playback latency for pure download is very long – longer than the video duration, due to the limited bandwidth available. This represents the upper bound for the playback latency.

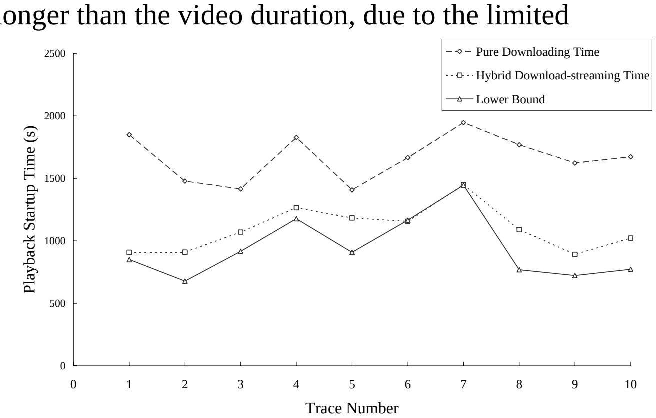

Fig. 6 Playback start time comparing to different schemes

By contrast, the hybrid download-streaming algorithm performs very well, closely tracking the lower bound. Considering the fact that the lower bound algorithm requires a prior knowledge of the available bandwidth and thus is not realizable, the hybrid download-streaming algorithm demonstrates the improvements achievable through using multiple sources together with the bandwidth model introduced in Section IV.

VI. Summary and Future Work

This work is a first step in exploring the feasibility and performance gain achievable through the modeling of aggregate available bandwidth from multiple senders. The experiments conducted in the Internet strongly support the bandwidth model and the application to hybrid download-streaming is also very promising. In addition to this application, the proposed multi-source bandwidth model can also be applied to adaptive video streaming [5], and potentially many other applications such as peer-to-peer applications where having multiple sources is the norm rather than the exception.

References

[1] T. Nguyen and A. Zakhor, "Distributed Video Streaming over the Internet" *SPIE Conference on Multimedia Computing and Networking*, San Jose, California, January 2002.

[2] V. Agarwal and R. Rejaie, "Adaptive Multi-source Streaming in Heterogeneous Peer-to-Peer Networks" *SPIE Conference on Multimedia Computing and Networking*, San jose, California, January 2005.

[3] E. Setton, Yi Liang, B. Girod, "Adaptive Multiple Description Video Streaming over Multiple Channels with Active Probing" *International Conference on Multimedia and Expo*, Baltimore, Maryland, vol. 1, pp. I-509-12, July, 2003.

[4] Planetlab Homepage : http://www.planet-lab.org/

[5] S.C. Hui, Jack Y.B. Lee, "Playback-Adaptive Multi-Source Video Streaming" *Fourth International Conference on Intelligent Multimedia Computing and Networking*, Salt Lake City, Utah, USA, July 2005.

[6] V. Paxson, "Fast approximation of self-similar network traffic", Tech. Rep., Lawrence Berkeley Laboratory and EECS Division, University of California, Berkeley. April 1995.

[7] K. Park, G. Kim, and M. Crovella, "On the Relation Between File Sizes, Transport Protocols, and Self-Similar Network Traffic," *International Conference on Network Protocols*, Columbus, Ohio, USA, p 171-180, Oct. 1996.

[8] K. Park, G. Kim, and M. Crovella, "On the Effect of Traffic Self-Similarity on Network Performance," *SPIE International Conference on Performance and Control of Network System*, pp. 296–310, November, 1997,.

[9] T. Tuan and K. Park, "Congestion Control for Self-Similar Network Traffic," Department of Computer Science., Purdue University, CSD-TR 98-014, May 1998,

[10] P. R. Morin, "The Impact of Self-Similarity on Network Performance Analysis," Ph.D. dissertation, Carleton University, Dec. 1995.

[11] R. Morris and Dong Lin, "Variance of aggregated Web traffic," *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies*, Tel-Aviv, Israel, vol. 1, pp.360-366, March 2000.

[12] Iperf Homepage : http://dast.nlanr.net/Projects/Iperf/

[13] L. S. Lam, Jack Y. B. Lee, S. C. Liew, and W. Wang, "A Transparent Rate Adaptation Algorithm for Streaming Video over the Internet," *18th International Conference on Advanced Information Networking and Applications*, Fukuoka, Japan, March 2004.

[14] Hahn and Shapiro, *Statistical Models in Engineering*, Wiley, 1994.