5-2006

Modeling End-to-End Available Bandwidth

Wanida Putthividhya
Iowa State University

Wallapak Tavanapong
Iowa State University

Modeling End-to-End Available Bandwidth

Abstract
The ability to accurately predict available bandwidth of an end-to-end path is of importance for improving quality of services for end users. For instance, end users can receive high quality video if senders can select to stream video data on an end-to-end path with available bandwidth that fits the required transmission rate. Modeling available bandwidth of an end-to-end path as a stochastic process is one of the approaches that can be used to predict the future available bandwidth of the path. The effectiveness of this approach depends on how well the chosen stochastic process can model the available bandwidth. The result of our analysis on the data published by Stanford Linear Accelerator Center (SLAC) indicates that Brownian Motion processes are not good for modeling available bandwidth of an end-to-end path.

Keywords
Available bandwidth, self-similar process, Brownian motion process, long-range dependence

Disciplines
OS and Networks

Modeling End-to-End Available Bandwidth

Wanida Putthividhya and Wallapak Tavanapong

TR #06-11
May 2006

Keywords: Available bandwidth, self-similar process, Brownian motion process, long-range dependence

Department of Computer Science
226 Atanasoff Hall
Iowa State University
Ames, Iowa 50011-1040, USA

Modeling End-to-End Available Bandwidth

Wanida Putthividhya Wallapak Tavanapong
Department of Computer Science
Iowa State University, Ames, IA 50011-1040, USA
Email:{wanidap,tavanapo}@cs.iastate.edu

Abstract

The ability to accurately predict available bandwidth of an end-to-end path is of importance for improving quality of services for end users. For instance, end users can receive high quality video if senders can select to stream video data on an end-to-end path with available bandwidth that fits the required transmission rate. Modeling available bandwidth of an end-to-end path as a stochastic process is one of the approaches that can be used to predict the future available bandwidth of the path. The effectiveness of this approach depends on how well the chosen stochastic process can model the available bandwidth. The result of our analysis on the data published by Stanford Linear Accelerator Center (SLAC) indicates that Brownian Motion processes are not good for modeling available bandwidth of an end-to-end path.

1 Introduction

End-to-end available bandwidth has been one of the most studied Quality-of-Service (QoS) metrics during the past decade. This is due to the increasing popularity of bandwidth-sensitive applications such as video streaming, video conferencing, distance learning, network gaming, just to name a few. Knowing the available bandwidth of an end-to-end path in advance is beneficial for bandwidth-sensitive applications in several aspects. For instance, a video streaming server can select the path with the least bandwidth fluctuation or choose to stream a low quality version of the requested video when the maximum available bandwidth of all paths to the client is less than the original playback bitrate. We use the term path and end-to-end path interchangeably hereafter.

Modeling the available bandwidth of a path using a known stochastic process is one possible method for estimating future available bandwidth of the path without explicit support from network routers. In other words, the stochastic process acting as the model reveals the probabilistic trend of the available bandwidth of the path. A stochastic process is a set of random variables indexed by times. Modeling the available bandwidth of a path as a stochastic process $\{X(t), t \geq 0\}$ is to consider the available bandwidth of the path at time t as one possible value of a random variable $X(t)$ and determine a class of known stochastic processes whose properties can be satisfied by the process $\{X(t), t \geq 0\}$. To verify whether the process $\{X(t), t \geq 0\}$ falls into a class of known processes, the process must satisfy all the properties of the class. We can then estimate the statistics about the available bandwidth of the path using the statistics of the known process. The effectiveness of the prediction depends on whether or not the chosen stochastic process is suitable to model the available bandwidth of a path.

Modeling the available bandwidth of a path is very challenging since the available bandwidth of the path usually varies from time to time. To the best of our knowledge, this is the first work attempting to model the available bandwidth of one path over time as a process. Our study is different from the related work in the literature which can be categorized into three main categories.

The related work in the first group [1–6] studied tools for estimating available bandwidth of a path at the measurement time. However, they cannot reveal the trend of the available bandwidth of the path in the future. The related work in the second category [7] studied modeling of the aggregate available bandwidth across several paths. The reference [7] reported experimental results showing that the aggregate available bandwidth across sufficiently large numbers of paths is normally distributed. The result of this recent work implies nothing about the behavior of the available bandwidth of a single path over time, which is the focus of our paper. The related work in the third category [8–15] studied modeling and characteristics of major network traffic types, for example, FTP, TELNET, HTTP (WWW), and variable-bit-rate (VBR) videos. The reference [10] showed that the amount of transmission bandwidth consumed by World Wide Web (WWW) traffic at a stub network is self-similar. The self-similar process can also be used to model the amount of transmission bandwidth used during the transmission of a long VBR video file [8, 11]. The existing work in this category and our work are different since we study modeling of the available bandwidth of a path.

Since the available bandwidth of the path and the amount of transmission bandwidth consumed by the cross traffic sharing the path are related, our hypothesis is that we can model the available bandwidth of a path as a **self-similar process** $\{Y(t)^H, t \geq 0\}$, where $H \in (0, 1)$ [16]. The Hurst parameter H measures the degree of self-similarity of the process. The closer H is to one, the higher the self-similarity level. Because of the complexity of verifying whether a process is self-similar and estimating the Hurst parameter of the process, we started with a simpler model. We simplified our hypothesis to whether or not the available bandwidth of a path can be modeled as a **Brownian motion process**[1] $\{Y(t)^{0.5}, t \geq 0\}$. Brownian motion process is a self-similar process with H = 0.5. Fig. 1 exhibits the relationship among processes in the self-similarity family. By testing this hypothesis, we would like to answer two questions: *(i)* Is Brownian motion process a good model for the available bandwidth of an end-to-end path? and *(ii)* If the answer to the first question is "no", is it still reasonable to model the available bandwidth of a path as a self-similar process?

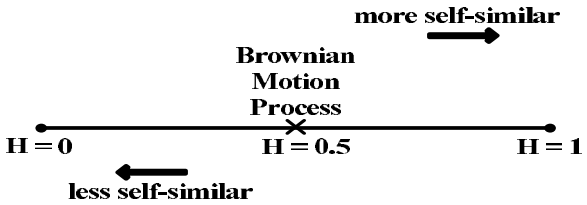

Figure 1: Self-Similarity Family

In this study, we test the hypothesis using available bandwidth data published by Stanford Linear Accelerator Center (SLAC) [17]. Using their own bandwidth measurement tool [4], they periodically collected available bandwidth of 34 different end-to-end paths starting from June to September 2004. Our results show that the available bandwidth of an end-to-end path over time is not a Brownian motion process. We do not have enough evidence at this time either to say whether it is reasonable to model available bandwidth of an end-to-end path as a self-similar process. Future work is needed to answer the question.

The rest of this paper is organized as follows. We present a brief literature survey of recent work on bandwidth estimation in Section 2. We provide background knowledge on Brownian motion process in Section 3. We explain our methodology in Section 4. Section 5 provides the results and discussions on the results. Finally, we give a conclusion and description of our future work in Section 6.

[1] Also known as Wiener process

2 Related Work

Bandwidth estimation has received considerable research attention in recent years [1–15, 18–23]. Related work in the literature can be divided into three main categories: *(i)* the studies of tools for estimating available bandwidth of an end-to-end path at a particular time [1–6], *(ii)* the studies of modeling of aggregate available bandwidth across sufficiently large numbers of paths [7], and *(iii)* the studies of modeling of network traffic [8–15].

2.1 Tools for Estimating Current Available Bandwidth

The existing work in the first category studied tools for estimating available bandwidth of an end-to-end path at the measurement time. They shared a common idea of using trains of probing packets with pre-determined gaps among them to test the current available bandwidth on a path. The packet trains are injected into a path of interest. The receiving side then determines the available bandwidth of the path according to the gaps between arriving packets and the original packet gaps. These tools yield only the available bandwidth of the path at the measurement time. *Trains of Packet Pairs* (TOPP) [3] and *Available Bandwidth Estimator* (ABwE) [4] are examples of such tools. Both tools inject trains of probing packets into the path of interest, but differ in methods to estimate the available bandwidth from the gaps between packets arriving at the destination host.

TOPP [3] consists of two phases: the probing phase and the analysis phase. In the probing phase, the sending host injects a number of trains of n equal sized packet pairs into a path of interest. The *intra-pair spacing* in each train corresponds to a single *offered rate*. The offered rates of consecutive trains differ by Δo. Fig. 2 illustrates an example of the probe traffic generated by TOPP. The probe traffic consists of three trains of packet pairs. Each train contains only one packet pair ($n = 1$). Each packet is of size b bytes. The three trains are sent with increasing offered rates (decreasing intra-pair gaps), i.e., $t_a > t_b > t_c$. The spacing between consecutive pairs (*inter-pair gap*) in the figure is set to be ΔT^p.

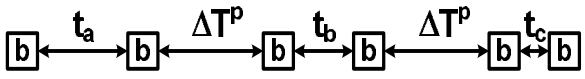

Figure 2: An example of probe traffic generated by TOPP tool.

The analysis phase of TOPP relies on the principle of the bottleneck spacing effect illustrated in Fig. 3. When two packets with time separation ΔS arrive at a link with service time $Q_b > \Delta S$ as shown in Fig. 3(a), then as the packets leave the link the time separation between them will be $\Delta R = Q_b$. On the other hand, if the service time of the link is $Q_b \leq \Delta S$ as shown in Fig. 3(b), the time separation between the two packets as they leave the link will remain the same, i.e., $\Delta R = \Delta S$. The available bandwidth experienced across the link can then be estimated as $b/\Delta R$, where b is the packet size. Upon receiving the i^{th} train of packet pairs, the receiving TOPP computes the available bandwidth estimate (f_i) as $b/\Delta R_i$, where ΔR_i represents the mean of intra-pair gaps within the i^{th} train upon receiving. After receiving every train of packet pairs, the receiving TOPP will get a set of ordered pairs $< o_i, f_i >, \forall i$.

With one congestible link (link shared by other end-to-end traffic) along a path, the linear equation $Y = \beta_1 X + \beta_0$ expressing the relationship between the offered rate (o) and the corresponding available bandwidth estimate (f) is as

$$\frac{o}{f} = \frac{1}{l}\, o + (1 - \frac{s}{l}),$$

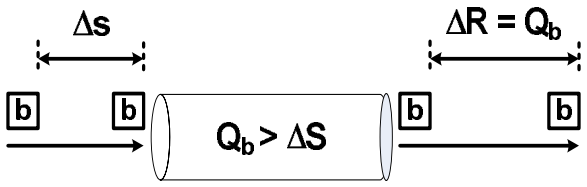

(a) Service time of the link (Q_b) is greater than the original intra-packet gap (ΔS)

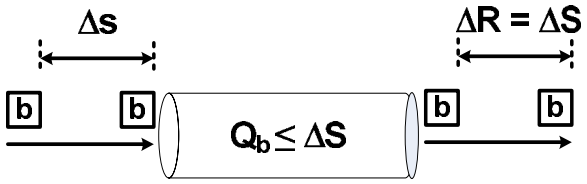

(b) Service time of the link (Q_b) is less than or equal to the original intra-packet gap (ΔS)

Figure 3: The principle of the bottleneck spacing effect.

where l represents the bottleneck link capacity of the path and s represents the true available bandwidth in which we try to estimate. Given two or more $< o_i, f_i >$ pairs, it is possible to estimate the unknown bottleneck bandwidth (l) and available bandwidth (s) of the path using *simple linear regression*. Given that $x_i = o_i$ and $y_i = \frac{o_i}{f_i}$, the coefficient $\beta_1 = \frac{1}{l}$ and the offset $\beta_0 = 1 - \frac{s}{l}$ can be computed as follows [24].

$$\begin{aligned} \beta_1 &= \frac{\sum(x_i - \bar{x})(y_i - \bar{y})}{\sum(x_i - \bar{x})^2} \\ \beta_0 &= \bar{y} - \beta_1 \bar{x} \end{aligned}$$

The bottleneck bandwidth (l) and the available bandwidth (s) of the path are then $1/\beta_1$ and $l \times (1 - \beta_0)$, respectively.

With more than one congestible links along a path, the plot between the offered rate (o) and the ratio between the offered rate and the corresponding available bandwidth estimate (o/f) becomes *piece-wise linear*. The plot contains a number of linear segments, each corresponding to one congestible link along the path. For each linear segment, TOPP employs the concept of simple linear regression discussed earlier to estimate the available bandwidth (s) of the corresponding congestible link. The available bandwidth of the entire path is the minimum among the available bandwidth of the congestible links along the path.

The reference [3] presented the performance evaluation of TOPP using ns network simulator [25]. The results showed that TOPP could estimate the available bandwidth of end-to-end paths with a high accuracy when the simulated topology contained a single congestible link, two congestible links, and more than two congestible links. The accuracy of TOPP decreases as the number of congestible links increases.

To estimate available bandwidth of an end-to-end path, ABwE sends a train of several (typically 20) packet pairs to the destination. Packets have the same size and the delay between packets of each packet pair is same. Like TOPP, ABwE determines from the set of the time separation within every pair of packets upon receiving the train of packets. The minimum time separation within a packet pair (Td_{init}) determines the most probable bottleneck bandwidth of the path. The current version of ABwE used the mean of the time separation within every packet pair (Td) to compute the available bandwidth (C) of the path. Let L_{PP} and L_{CT} be the sizes of a probing packet and a

cross traffic packet, respectively. The available bandwidth of the path can be estimated as follows [4].

$$\begin{aligned} C &= (L_{PP} + QDF \times L_{CT})/Td \\ QDF &= (Td - Td_{init})/NTT_{class} \end{aligned}$$

Queueing Delay Factor (QDF) of the path represents the expected number of cross traffic packets in between a pair of probing packets while traveling through the path. If there is no cross traffic sharing the path or the amount of cross traffic is so small that it can be negligible ($Td = Td_{init}$), the QDF of the path becomes zero. If the amount of the cross traffic is significant, the QDF of the path is greater than zero. Nominal Transmission Time (NTT_{class}) is the time needed for transmission of one packet with the size of Maximum Transfer Unit (MTU) [4].

The reference [4] ran ABwE to monitor the available bandwidth of end-to-end paths from Stanford to 12 remote hosts in US, Europe, and Asia for about 3 months. For each path, the ABwE results were compared with Iperf TCP throughputs. The relationship between the two set of results can be explained by the linear equation $Y = 1.13X$, where Y represents the results from ABwE and X represents the results from Iperf TCP throughputs.

2.2 Modeling of Aggregated Available Bandwidth

The existing work in the second category studied modeling of aggregate available bandwidth across several paths. The behavior of the aggregate available bandwidth across several paths does not imply anything about the behavior of the available bandwidth of an individual path over time which is the aim of this paper.

The reference [7] showed that the aggregate available bandwidth across a sufficiently large number of paths is normally distributed. This work also illustrated the application of the results through a *hybrid download-streaming algorithm* for video delivery with probabilistic performance guarantees. It was found that if one receiver receives a media file from multiple senders, the aggregate data transfer rate will exhibit a normal distribution. From this discovery, the algorithm determines w—the start of the playback at the receiver. Let C_i be the aggregate data transfer rate at time interval i after the download has started. Assume that the receiver begins the playback at w intervals since the download has begun. Let n represent some time interval since the download has begun. Assume also that R is the playback rate at the receiver. To achieve a continuous playback at the receiver,

$$\sum_{j=1}^{n} C_j \quad \geq \quad R(n-w).$$

Since it has been shown that C_j is normally distributed, the sum of n C_j's is also normally distributed. Assume that the receiver would like to ensure a continuity playback with probability of Δ. That is,

$$Pr\,[\,\sum_{j=1}^{n} C_j \geq R(n-w)\,] \quad \geq \quad \Delta.$$

Let $F^{(n)}$ be the cumulative distribution function of the sum of C_j, $j = 1...n$. The earliest playback time at the receiver is then

$$w \quad = \quad Min\,\{\, v \mid 1 - F^{(n)}(R(n-v)) \geq \Delta,\, \forall n \geq v \,\}.$$

The performance of the hybrid download-streaming algorithm was evaluated using traffic traces collected from PlanetLab [26]. For each traffic trace, the performance

of the hybrid download-streaming algorithm was compared with the *pure-download* algorithm where the entire video was downloaded before the playback and the *lower-bound* algorithm specifying the lower bound of the playback time. The lower-bound algorithm computed the ideal earliest playback time assuming the traffic is known a priori. Therefore, this method cannot actually be implemented in practice. The results showed that the playback time determined by the hybrid download-streaming algorithm was significantly lower than the one determined by the pure-download algorithm and very close to the one determined by the lower-bound algorithm.

2.3 Modeling of Network Traffic

The work in this category focused on characterizing major types of network traffic. For example, the inter-arrival time of FTP and TELNET sessions, the inter-arrival time of FTP and TELNET packets, the distribution of TELNET packet sizes, and modeling of the amount of bandwidth consumed by variable-bit-rate (VBR) video flows. The references [8, 14] are two examples of the work in this category.

Reference [14] is one of the classic studies in the area. The work analyzed main characteristics of FTP and TELNET traffic from traffic traces collected at Lawrence Berkeley Laboratory and Internet access point for the Digital Equipment Corporation. Some important findings are as follows. A Poisson process is not always a good model for the connection arrivals of every traffic type. In particular, while the arrivals of user-initiated TCP sessions such as TELNET and FTP sessions can be well modeled as Poisson processes, the arrivals of machine-generated traffic sessions, for example HTTP (WWW), SMTP, and NNTP, cannot be modeled as Poisson processes. For TELNET, this work found that the distributions of the inter-arrival time of TELNET packets and the TELNET connection sizes follow Pareto distribution (with the shape parameter b between 0.9 and 0.95) and log-normal distribution, respectively. For FTP, the work found that FTP data usually come in bursts. In addition, the burst sizes of FTP data follow Pareto distribution with the shape parameter $0.9 \leq b \leq 1.4$. The authors also studied the possibility of *self-similarity* in overall wide-area traffic. They found the evidence inconclusive even though the traffic clearly exhibits large-scale correlations inconsistent with Poisson processes.

The existing work [8] studied the characteristics of variable-bit-rate (VBR) video traffic from 20 different video sequences of different lengths (15 seconds or a few hundreds frames to 2 hours or about 171,000 frames). One important finding is that the frame sizes (in bytes) from a very long sequence of VBR video frames exhibits *long-range dependence* [2].

Our work is different from these existing studies in this category since we are focusing on modeling the available bandwidth of an end-to-end path, not a particular type of traffic.

3 Brownian Motion Process

In this section, we provide background on Brownian motion process. **Brownian motion process** is named after an English botanist Robert Brown. The process is also known as a **Wiener process** since Robert Weiner (1918) came up with a precise definition of the process. The reference [27] gives the definition of Brownian motion processes as follows.

Definition A process $\{X(t), t \geq 0\}$ is called a Brownian motion process, if

1. $X(0) = 0$.

[2] A process with long-range dependence exhibits a slow decay auto-correlation function.

2. **Stationary**: Given two time points $t_1 < t_2$, the increment $X(t_2) - X(t_1)$ is a normally distributed random variable with mean of $\mu(t_2 - t_1)$ and variance of $\sigma^2(t_2 - t_1)$. The parameters μ and σ^2 are the mean and variance of all the increments. That is,

$$X(t_2) - X(t_1) \sim N(\mu(t_2 - t_1),\ \sigma^2(t_2 - t_1)).$$

3. **Independent Increments**: Given time points $0 < t_1 < t_2 < \ \ldots \ < t_{n-1} < t_n$, the increments

$$X(t_1),\ X(t_2) - X(t_1),\ \ldots\ ,\ X(t_n) - X(t_{n-1})$$

are independent random variables.

It is know that the following lemmas hold for Brownian motion processes.

Lemma 1 $X(t) \sim N(\mu(t),\ \sigma^2(t)),\ \ \forall t > 0.$

Proof: This can be directly drawn from the first two properties of Brownian motion processes.

$$\begin{aligned} X(t) - X(0) &\sim N(\mu(t),\ \sigma^2(t)) \\ X(t) &\sim N(\mu(t),\ \sigma^2(t)) \end{aligned}$$

Lemma 2 *Given* $X(0) = a$ *and* $t_1 < t_2$, $X(t_2) - X(t_1) \sim N(\mu(t_2 - t_1),\ \sigma^2(t_2 - t_1))$.

Proof: The fact that $X(0) = a$ affects neither the value nor the distribution of any increment since the beginning value a is included in $X(t),\ \forall t \geq 0$.

Let $X_0(t_1)$ and $X_0(t_2)$ be the processes at time t_1 and t_2, when $X(0) = 0$. Let $X_a(t_1)$ and $X_a(t_2)$ be the processes at time t_1 and t_2, when $X(0) = a$.

$$\begin{aligned} X_a(t_1) &= a + X_0(t_1) \\ X_a(t_2) &= a + X_0(t_2) \\ So,\ \ X_a(t_2) - X_a(t_1) &= (a + X_0(t_2)) - (a + X_0(t_1)) \\ &= X_0(t_2) - X_0(t_1) \end{aligned}$$

Therefore, the distributions of increments when $X(0) = 0$ and $X(0) = a$ are the same.

Lemma 3 *Given that the current state of the process at time* t_1 *is* b *and the bandwidth increments are normally distributed, i.e.,* $X(t_2) - X(t_1) \sim N(\mu(t_2 - t_1),\ \sigma^2(t_2 - t_1))$, $t_1 < t_2$, *the process at time* t_2 *(*$X(t_2)$*) is a normally distributed random variable with mean of* $b + \mu(t_2 - t_1)$ *and variance of* $\sigma^2(t_2 - t_1)$.

Proof: The proof is based on the idea that if $X \sim N(\mu,\ \sigma^2)$ and c is a constant, then $X + c \sim N(c + \mu,\ \sigma^2)$.

$$\begin{aligned} X(t_2) &= X(t_1) + (X(t_2) - X(t_1)) \\ X(t_2) &= b + (X(t_2) - X(t_1)) \\ So,\ \ X(t_2) &\sim N(b + \mu(t_2 - t_1),\ \sigma^2(t_2 - t_1)) \end{aligned}$$

Lemma 3 can be used to obtain an estimate of the future available bandwidth of a path from the current available bandwidth if the available bandwidth can be modeled as a Brownian motion process. The lemma also shows that Brownian motion processes satisfy the *Markov property* since the conditional probability of the state of the process at time t_2 (future state) depends only on the state of the process at time t_1 (present

state). Actually, Brownian motion process is a Markov process with a continuous state space and a continuous time set [28]. Readers are referred to the Appendix for the definition of Markov processes and the Markov property in detail.

Brownian motion process is also a self-similar process with Hurst parameter $H = 0.5$ [16, 29]. A stochastic process $\{Y(t), t \geq 0\}$ is said to be *"self-similar"* with Hurst parameter H, if it satisfies the condition:

$$\{Y(at)\} \quad \stackrel{d}{=} \quad \{a^H\, Y(t)\},\ \forall a > 0,\ 0 < H < 1,$$

where the equality is in terms of finite dimensional distributions[3] [16]. The Hurst parameter H indicates the level of self-similarity of the process. The closer H is to one, the higher the self-similarity of the process.

Self-similar processes with $H > 0.5$ exhibit long-range dependence. The closer H is to one, the greater the degree of long-range dependence[4]. Long-range dependence can be characterized by the auto-correlation (auto-covariance) plots of the processes. A stationary process is *long-range dependent* if its auto-correlation function is non-summable, i.e., $\sum_k\ \rho(k) = \infty$. Fig. 4 taken from the reference [30] displays the auto-covariance functions of long-range dependent and short-range dependent processes. The y-axis of the plots represents the auto-covariance values whereas the x-axis represents the time differences between any pair of data considered when computing auto-covariance. From the figure, the auto-correlation of the long-range dependent process decays slowly as the time difference increases. In the other words, the current state of the long-range dependent process will have a very long-term influence on future states of the process.

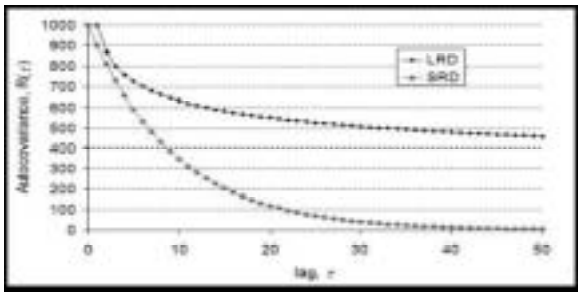

Figure 4: Auto-covariance functions of long-range dependent and short-range dependent processes.

4 Methodology

In this section, we first state our hypothesis about the model of the available bandwidth of an end-to-end path. Then, we discuss the preparation of the available bandwidth data published by Stanford Linear Accelerator Center (SLAC) [17] at Stanford University for verifying the hypothesis. Last, we describe our procedure for testing the hypothesis.

4.1 Model of the Available Bandwidth of an End-to-End Path

Our *initial hypothesis* is that the available bandwidth of a single end-to-end path can be modeled as a self-similar process with the Hurst parameter H. Due to the complexities in estimating H to verify the initial hypothesis, we simplify the initial hypothesis and

[3] $P[X(at) < c] = P[a^H X(t) < c]$.

[4] Also known as long memory

verify this simplified hypothesis in this work. The *simplified hypothesis* is stated as follows.

Hypothesis 1: *The available bandwidth of a single end-to-end path can be modeled as a Brownian motion process* $\{X(t)^H,\ t \geq 0\}$*, where* $H = 0.5$*.*

It is noteworthy that the initial and simplified hypotheses are relevant. Recall that Brownian motion process is a self-similar process with Hurst parameter $H = 0.5$.

4.2 Preparation of Available Bandwidth Data

Stanford Linear Accelerator Center (SLAC) published available bandwidth data of 34 different end-to-end paths from the center at Stanford University to 34 remote hosts in 12 countries in North America, Europe, and Asia. The data was collected for about 101 days from June to September 2004. The available bandwidth on every path was measured at about 1 minute interval using *ABwE*—their own tool for estimating available bandwidth of an end-to-end path at a particular time [4].

Ideally, we should be able to use the data right away. However, the data have some characteristics that may affect the accuracy of the results. In the following, we discuss these characteristics and the corresponding modifications we made to the data set.

First, the data corresponding to a path may contain some outlier periods. These are periods in which the available bandwidth is significantly low. SLAC identified these outlier periods using a modified "plateau" algorithm [31]. We removed outlier periods reported for every path. By removing outlier periods, one data set corresponding to each path is divided into several sets of data, each containing available bandwidth data of the path measured consecutively in time. We considered these newly created sets of available bandwidth data as different processes and verified our assumption using one set at a time. Fig. 5 shows outlier periods corresponding to an end-to-end path detected by the modified "plateau" algorithm. Each black bar on the top of the figure roughly exhibits the length of one outlier period.

Second, the time differences between pairs of consecutive measurements inside a set of data corresponding to an end-to-end path may not be exactly the same. Fig. 6 shows two examples of variations in time differences between pairs of consecutive measurements. To verify the stationary property of a SLAC's data set, the non-uniform interval lengths need to be taken care of. We will explain this issue in detail later in the next section.

4.3 Hypothesis Testing

To test Hypothesis 1, we need to check whether the available bandwidth of an end-to-end path satisfies the stationary and independent increments properties of Brownian motion processes discussed in Section 3. Note that we omit the first property since Lemma 2 has shown that it is not a strong restriction.

Let $X(t)$ be a random variable representing the available bandwidth of a path at time t. Let $Y(t)$ be another random variable representing the difference between the available bandwidth of a path measured at two consecutive measurement times. That is, given measurement times of a path $t_0 < t_1 < \ldots < t_{n-1} < t_n$,

$$Y(t_i) \quad = \quad X(t_{i+1}) - X(t_i),\ i = 0 \ldots n-1.$$

We observe from SLAC's available bandwidth data that the mean of bandwidth increments of a path is either zero or very close to zero. Fig. 7 displays means of bandwidth increments corresponding to end-to-end paths considered by SLAC. The two notable spikes in the figure correspond to the means of bandwidth increments of an end-to-end path from SLAC to the remote host *node1.fnal.gov* where average available bandwidth is about 500Mbps and to the remote host *node1.kek.jp* where average available bandwidth is about 180Mbps, respectively. Therefore, the non-zero means of bandwidth

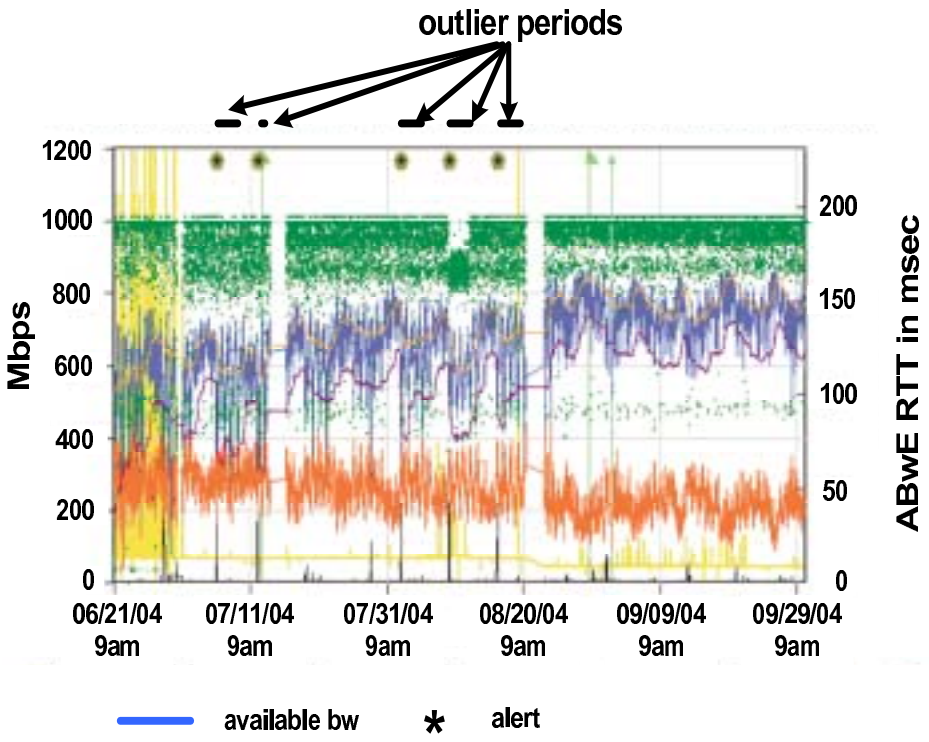

Figure 5: Detected by a modified "plateau" algorithm, periods in which available bandwidth on the path between Stanford University and the server node1.cacr.caltech.edu was significantly low.

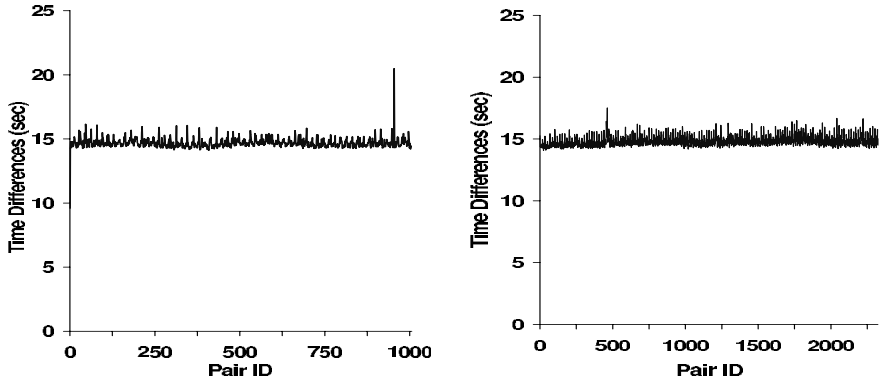

(a) node1.dl.ac.uk (b) node1.cacr.caltech.edu

Figure 6: Variations of time differences between pairs of consecutive measurements.

increments are small enough to be considered negligible. In the rest of this paper, we assume that for a path, $\frac{1}{n}\sum_{t=t_0}^{t_{n-1}} Y(t) = 0$. Our observation about zero-mean bandwidth increments of the data alters the stationary property to be

$$X(t_2) - X(t_1) \sim N(0,\ \sigma^2(t_2 - t_1)), \quad t_1 < t_2.$$

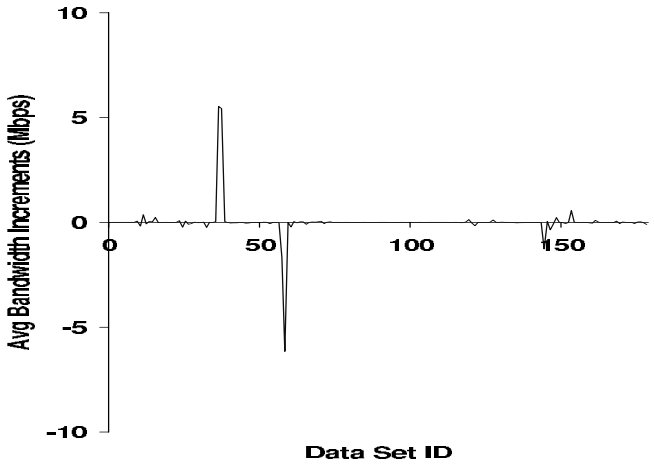

Figure 7: Means of bandwidth increments.

As discussed earlier, SLAC made measurements of available bandwidth of an end-to-end path at every non-uniform interval. This causes $Y(t)$, $t = t_0 \ldots t_{n-1}$, to have different variances. We, then, cannot simply use $y(t)$, $t = t_0 \ldots t_{n-1}$ to verify the stationary property. Therefore, we need to further standardize the random variables $Y(t)$, $t = t_0 \ldots t_{n-1}$. The standardized bandwidth increments at time t, from t_0 to t_{n-1}, will have the same mean of 0 and variance of σ^2. Let $Z(t)$ be a random variable representing the standardized bandwidth increments at time t. That is, given time points $t_0 < t_1 < \ldots < t_{n-1} < t_n$,

$$\begin{aligned} Z(t_i) &= \frac{Y(t_i)}{\sqrt{t_{i+1} - t_i}}, \qquad i = 0 \ldots n-1 \\ &= \frac{X(t_{i+1}) - X(t_i)}{\sqrt{t_{i+1} - t_i}} \ . \end{aligned}$$

Let $z(t)$, $t = t_0 \ldots t_{n-1}$ be the standardized bandwidth increments derived from the available bandwidth data from SLAC. We use $z(t)$ to verify the stationary property.

4.3.1 Verifying Stationary Property

To verify whether the available bandwidth of an end-to-end path satisfies the stationary property, we check whether $z(t_0),\ z(t_1),\ z(t_2),\ \ldots,\ z(t_{n-1})$ follow a normal distribution. Since the standardized bandwidth increments are just the bandwidth increments divided by corresponding time differences, showing that the standardized bandwidth increments are normally distributed is equivalent to showing that the bandwidth increments follow a normal distribution.

There are several normality checking mechanisms. We use the *Quantile-Quantile plot (Q-Q plot)* which is a graphical method for checking the normality of $z(t)$, $t = t_0 \ldots t_{n-1}$. The estimated percentiles of the standardized bandwidth increments are listed on the y-axis of the plot. The percentiles of a normal distribution of interest are listed on the x-axis of the plot. For each path, the q^{th} percentile of the standardized bandwidth increments is plotted against the q^{th} percentile of standard normal. If

an approximately straight line is formed, the standardized bandwidth increments are normally distributed. Otherwise, they are not normally distributed.

4.3.2 Verifying Independent Increments Property

To verify whether the available bandwidth of an end-to-end path satisfies the independent increments property, we compute the auto-correlation among $z(t)$, $t = t_0 \ldots t_{n-1}$. In particular, we compute the auto-correlation between $z(t_i)$ and $z(t_{i+k})$, where $k \geq 0$ and $i = 0 \ldots (n-1-k)$. The variable k is usually called *lag*. The trivial case is when $k = 0$, i.e., the auto-correlation between $z(t_i)$ and itself. Let $cov(k)$ and $\rho(k)$ be the auto-covariance and auto-correlation between $z(t_i)$ and $z(t_{i+k})$, respectively. From the reference [32], we compute the auto-covariance and auto-correlation among $z(t)$, $t = t_0 \ldots t_{n-1}$, as follows.

$$
\begin{aligned}
cov(0) &= \frac{1}{n} \sum_{i=0}^{n-1} (z(t_i) - \bar{z})^2 \\
cov(k) &= \frac{1}{n} \sum_{i=0}^{n-k-1} (z(t_i) - \bar{z})(z(t_{i+k}) - \bar{z}), \quad k > 0 \\
\rho(0) &= cov(0)/cov(0) \\
\rho(k) &= cov(k)/cov(0), \qquad\qquad k > 0
\end{aligned}
$$

Note that $\rho(k) \in [-1, 1]$, $k \geq 0$. If $z(t)$, $t = t_0 \ldots t_{n-1}$, are independent, the auto-correlations should be near zero for all lag values. When presented, the auto-correlations are usually plotted against the corresponding k values. Similar to the case of the stationary property, showing that the standardized bandwidth increments of an end-to-end path are independent is the same as showing that the bandwidth increments of the path are independent.

5 Results and Discussions

We present and discuss our results in the following. The results show whether the standardized bandwidth increments corresponding to an end-to-end path are normally distributed and independent. In turn, the results will indicate whether the available bandwidth of an end-to-end path can be modeled as a Brownian motion process.

Since the available bandwidth of every considered path exhibits the same behavior, we present only the results of some paths in this section. Fig. 8-13 (a-b) demonstrate normality checking and auto-correlation of the standardized bandwidth increments corresponding to end-to-end paths from SLAC to two remote servers in North America (*node1.cacr.caltech.edu* and *node1.ece.rice.edu*), Asia (*node1.jp.apan.net* and *node1.niit.pk*), and Europe (*node1.cesnet.cz* and *node1.desy.de*), respectively.

Fig. 8-13 (b) show that except for lag of 0 which is always 1 by definition, almost all auto-correlations are close to 0. Besides, the auto-correlations in each figure do not exhibit any exact pattern[5] that might lead to the conclusion that the data are related [32]. Therefore, there are no significant auto-correlations among the standardized bandwidth increments of each path. These indicate that the available bandwidth of an end-to-end path satisfies the independent increments property of a Brownian motion process.

In contrast, the available bandwidth of an end-to-end path does not satisfy the stationary property of a Brownian motion process. The normal Q-Q plot of standardized bandwidth increments exhibits a straight line except for both tails of the plot.

[5]One possible pattern that the auto-correlations among non-independent data can exhibit is that the auto-correlations gradually decrease as the lag value increases.

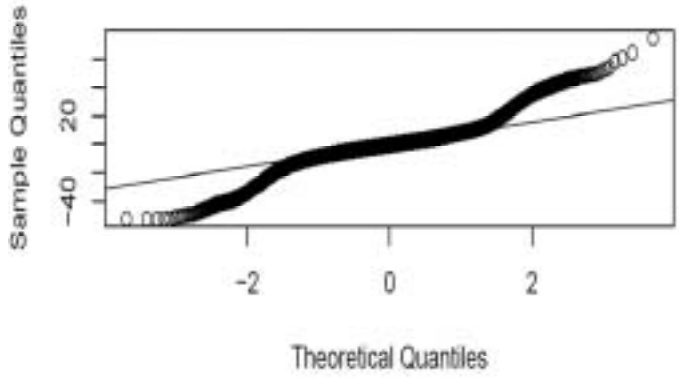

(a) normal Q-Q plot

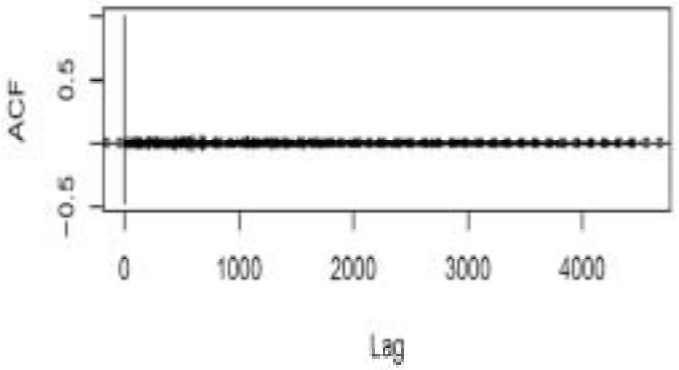

(b) auto-correlation

Figure 8: Distribution and auto-correlation of the standardized bandwidth increments corresponding to path from SLAC to node1.cacr.caltech.edu.

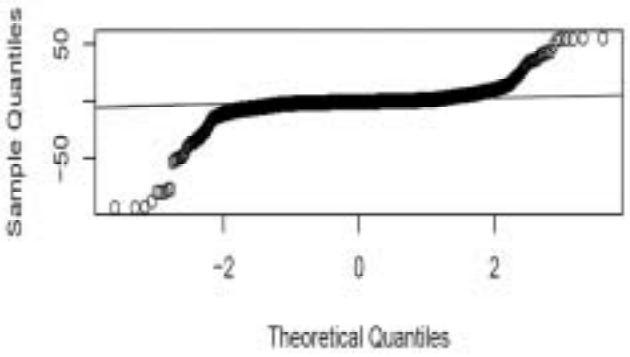

(a) normal Q-Q plot

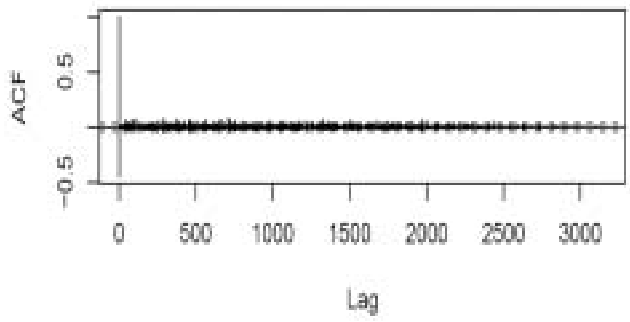

(b) auto-correlation

Figure 9: Distribution and auto-correlation of the standardized bandwidth increments corresponding to path from SLAC to node1.ece.rice.edu.

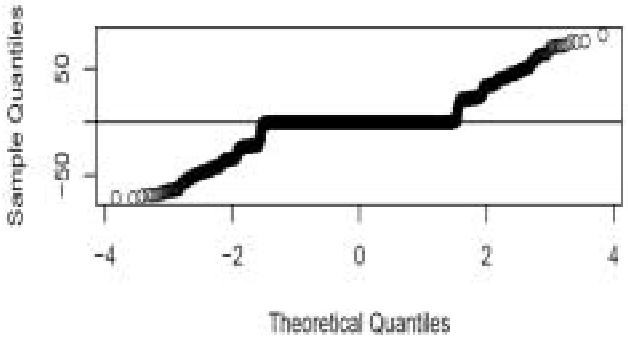

(a) normal Q-Q plot

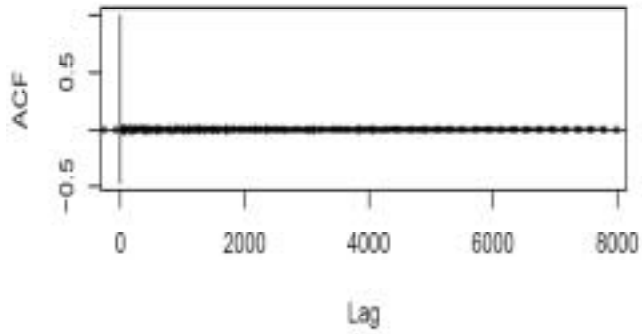

(b) auto-correlation

Figure 10: Distribution and auto-correlation of the standardized bandwidth increments corresponding to path from SLAC to node1.jp.apan.net.

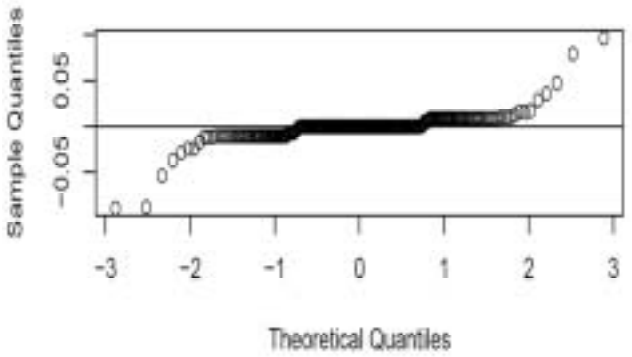

(a) normal Q-Q plot

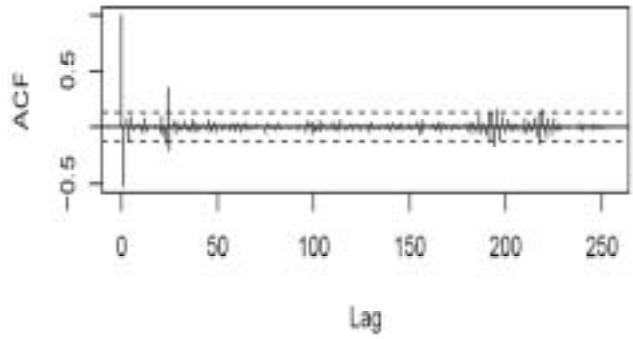

(b) auto-correlation

Figure 11: Distribution and auto-correlation of the standardized bandwidth increments corresponding to path from SLAC to node1.niit.pk.

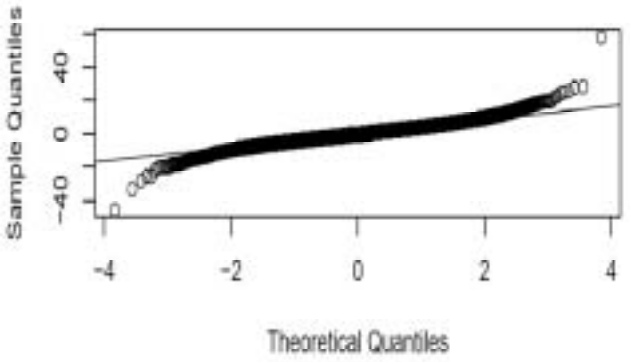

(a) normal Q-Q plot

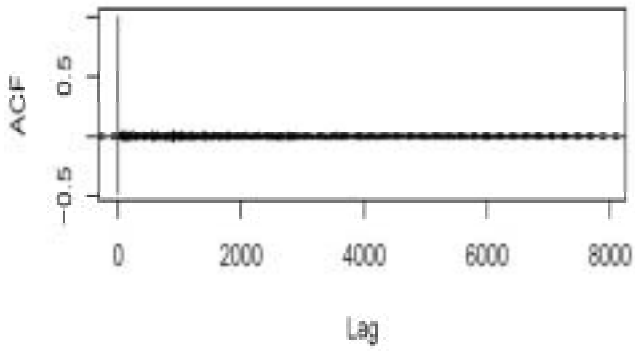

(b) auto-correlation

Figure 12: Distribution and auto-correlation of the standardized bandwidth increments corresponding to path from SLAC to node1.cesnet.cz.

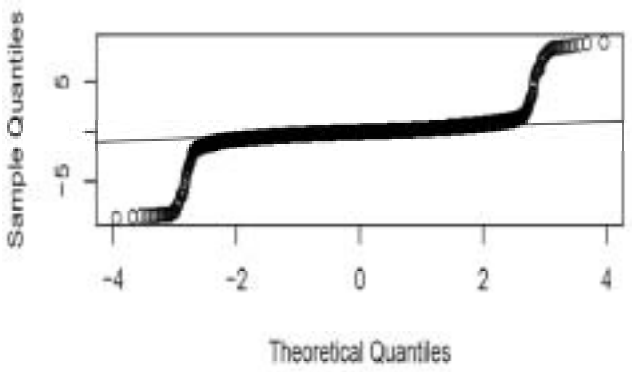

(a) normal Q-Q plot

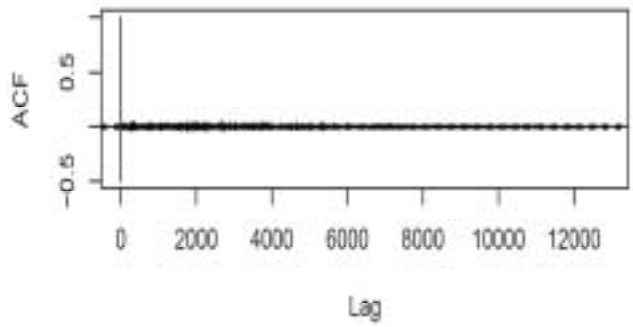

(b) auto-correlation

Figure 13: Distribution and auto-correlation of the standardized bandwidth increments corresponding to path from SLAC to node1.desy.de.

This shows that the bandwidth increments of an end-to-end path are not normally distributed. In fact, the normal Q-Q plots in Fig. 8-13 (a) correspond to those of data with heavy-tailed distribution [33].

Therefore, we conclude from the results that Brownian motion process is not a good model for available bandwidth of an end-to-end path. The main reason is that the stationary property of Brownian motion processes cannot be satisfied.

6 Concluding Remarks

In this paper, we have shown that a Brownian motion process is not a suitable model for available bandwidth of an end-to-end path. Our analysis were based on data published by Stanford Linear Accelerator Center (SLAC). The main reason is that the stationary property of Brownian motion processes cannot be satisfied. This result left us with a future work of finding some other model that is better suited. Other self-similar processes are candidates for the model. In addition, to ensure that our result does not depend on the SLAC method for bandwidth estimation, we will need to verify the model using available bandwidth information obtained from some other bandwidth measuring tools.

Acknowledgments

The authors thank Dr. Arka P. Ghosh from the Department of Statistics, Iowa State University for his help and suggestions on the statistic work in this study. This work is partially supported by the US National Science Foundation under Grant No. 0092914. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

[1] M. Jain and C. Dovrolis. Pathload: A measurement tool for end-to-end available bandwidth. In *Proc. of Passive and Active Measurements (PAM) Workshop*, pages 14–25, Fort Collins, Colorado, USA, March 2002.

[2] N. Hu and P. Steenkiste. Evaluation and characterization of available bandwidth probing techniques. *IEEE Journal on Selected Areas in Communications*, 21(6):879–894, August 2003.

[3] B. Melander, M. Bjorkman, and P. Gunningberg. A new end-to-end probing and analysis method for estimating bandwidth bottlenecks. In *Proc. of IEEE GLOBECOM*, pages 415–420, San Francisco, CA, USA, November 2000.

[4] J. Navratil and R. L. Cottrell. ABwE: A practical approach to available bandwidth estimation. In *Proc. of Passive and Active Measurements (PAM) Workshop*, La Jolla, California, USA, April 2003.

[5] V. Ribeiro, R. Riedi, R. Baraniuk, J. Navratil, and L. Cottrell. pathchirp: Efficient available bandwidth estimation for network paths. In *Proc. of Passive and Active Measurements (PAM) Workshop*, La Jolla, California, USA, April 2003.

[6] J. Strauss, D. Katabi, and F. Kaashoek. A measurement study of available bandwidth estimation tools. In *Proc. of ACM IMC*, pages 39–44, Miami Beach, FL, USA, October 2003.

[7] S. C. Hui and J. Y. B. Lee. Modeling of aggregate available bandwidth in many-to-one data transfer. In *Proc. of Fourth Int'l. Conf. on Intelligent Multimedia Computing and Networking*, Salt Lake City, UT, USA, July 2005.

[8] J. Beran, R. Sherman, M. S. Taqqu, and W. Willinger. Long-range dependence in Variable-Bit-Rate video traffic. *IEEE Transactions on Communications*, 43(2/3/4):1566–1579, February/March/April 1995.

[9] T. Borsos. On the temporal characteristics of video traffic. In *26th Annual IEEE Int'l. Conf. on Local Computer Networks*, pages 500–508, Tampa, Florida, USA, November 2001.

[10] M. E. Crovella and A. Bestavros. Self-similarity in World Wide Web traffic: Evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, December 1997.

[11] M. Garrett and W. Willinger. Analysis, modeling and generation of self-similar VBR video traffic. In *Proc. of ACM SIGCOMM*, pages 269–280, London, UK, August 1994.

[12] K. Park, G. Kim, and M. Crovella. On the relationship between file sizes, transport protocols, and self-similar network traffic. In *Proc. of IEEE ICNP*, pages 171–180, Columbus, Ohio, USA, October 1996.

[13] K. Park, G. Kim, and M. Crovella. On the effect of traffic self-similarity on network performance. In *Proc. SPIE Int'l. Conf. on Performance and Control of Network Systems*, pages 296–310, Dallas, Texas, USA, November 1997.

[14] V. Paxson and S. Floyd. Wide-area traffic: The failure of Poisson modeling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, June 1995.

[15] Ve. Paxson. Fast approximation of self similar network traffic. Technical Report LBL-36750, Lawrence Berkeley Laboratory and EECS Division, University of California Berkeley, CA, USA, April 1995.

[16] R. J. Adler, R. E. Feldman, and M. S. Taqqu. *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*. Birkhauser Boston, Boston, MA, USA, 1998.

[17] SLAC. Stanford linear accelerator center. `http://www.slac.stanford.edu/comp/net/bandwidth-tests/eventanalysis/all_100days_sep04/`, Last accessed April 13, 2006.

[18] R. L. Carter and M. E. Crovella. Measuring bottleneck link speed in packet-switched networks. *Performance Evaluation*, 27-28:297–318, 1996.

[19] C. Dovrolis, P. Ramanathan, and D. Moore. What do packet dispersion techniques measure? In *Proc. of IEEE INFOCOM*, pages 905–914, Anchorage, AK, USA, April 2001.

[20] A. B. Downey. Using pathchar to estimate internet link characteristics. In *Proc. of ACM SIGCOMM*, pages 241–250, Cambridge, MA, USA, September 1999.

[21] K. Harfoush, A. Bestavros, and J. Byers. Measuring bottleneck bandwidth of targeted path segments. In *Proc. of IEEE INFOCOM*, pages 2079–2089, San Francisco, CA, USA, March 2003.

[22] M. Jain and C. Dovrolis. End-to-end available bandwidth: Measurement methodology, dynamics, and relation with tcp throughput. *IEEE/ACM Transactions on Networking*, 11(4):537–549, August 2003.

[23] K. Lai and M. Baker. Measuring bandwidth. In *Proc. of IEEE INFOCOM*, pages 235–245, New York, NY, USA, March 1999.

[24] J. Dallal. Introduction to simple linear regression. `http://www.tufts.edu/~gdallal/slr.htm`, Last accessed April 13, 2006.

[25] NS2. Ns2: Network simulator. `http://www.isi.edu/nsnam/ns/`, Last accessed April 13, 2006.

[26] PlanetLab. Planetlab. `http://www.planet-lab.org/`, Last accessed April 13, 2006.

[27] D. Freedman. *Brownian Motion and Diffusion*. Holden-Day, San Francisco, CA, USA, 1971.

[28] M. Kim. Chapter 4 brownian motion. `http://web.am.qub.ac.uk/users/m.s.kim/chap4.pdf`, Last accessed April 13, 2006.

[29] P. Embrechts and M. Maejima. *Selfsimilar Processes*. Princeton University Press, Princeton, NJ, USA, 2002.

[30] T. Neame. Performance evaluation of a queue fed by a poisson pareto burst process. *Computer Networks: The International Journal of Computer and Telecommunications Networking (Special issue: Advances in modeling and engineering of Longe-Range dependent traffic)*, 40(3):377–397, October, 2002.

[31] A. J. McGregor and H. W. Braun. Automated event detection for active measurement systems. Available at `http://byerley.cs.waikato.ac.nz/~tonym/papers/event.pdf`, Last access April 13, 2006.

[32] Engineering Statistic Handbook. Autocorrelation plot. `http://www.itl.nist.gov/div898/handbook/eda/section3/autocopl.htm`, Last accessed April 13, 2006.

[33] PROPHET StatGuide. Examples of normal probability plots. `http://www.basic.northwestern.edu/statguidefiles/probplots.html#Heavy-tailed%20Data`, Last accessed April 13, 2006.

[34] WIKIPEDIA. Pareto distribution. `http://en.wikipedia.org/wiki/Pareto_distribution`, Last accessed April 13, 2006.

[35] MathWorld. Pareto distribution. `http://mathworld.wolfram.com/ParetoDistribution.html`, Last accessed April 13, 2006.

[36] WIKIPEDIA. Log-normal distribution. `http://en.wikipedia.org/wiki/Log-normal_distribution`, Last accessed April 13, 2006.

Supplementary Materials

Markov Processes

A stochastic process is a Markov process if it satisfies the *Markov property*. In turn, a stochastic process has the Markov property if the conditional probability distribution of future states of the process, given the past and present states, depends only on the present state. Let x(t) be the current state of the process. Let $x(s),\ s < t$ be the past states of the process. Let $x(t+h)$ be the future state of the process.

$$\begin{aligned} Pr[X(t+h) = x(t+h) \mid X(s) = x(s),\ s \leq t] \quad = \\ Pr[X(t+h) = x(t+h) \mid X(t) = x(t)] \end{aligned}$$

Pareto Distribution

The Pareto distribution is named after an Italian economist, Vilfredo Pareto. The Pareto distribution belongs to the heavy-tailed distribution family and has been found in many real-world situations, for example, the distribution of Internet file sizes, the distribution of the sizes of human settlements, the distribution of lengths of jobs assigned to supercomputers, just to name a few [34].

Let $a > 0$ and $b > 0$ are the *location* and *shape* parameters, respectively. The probability density function ($f(x)$) and the cumulative distribution function ($F(x)$) of the Pareto distribution defined over the interval $x \geq b$ are as follows [35].

$$\begin{aligned} f(x) &= \frac{ab^a}{x^{a+1}} \\ F(x) &= 1 - (\frac{b}{x})^a \end{aligned}$$

From the probability density function, the expected value and variance of the Pareto distribution can be derived as follows [34, 35]. If $b \leq 2$, the distribution has an infinite variance. When $b \leq 1$, the distribution has an infinite mean [14, 34].

$$\begin{aligned} E[X] &= \frac{ab}{a-1}, \qquad a > 1 \\ VAR[X] &= \frac{ab^2}{(a-2)(a-1)^2},\ a > 2 \end{aligned}$$

Fig. 14 (a-b) illustrate the plots corresponding to the probability density function of Pareto random variables with different location (a) and shape (b) parameters. Fig. 14 (a) shows the plots corresponding to the probability density function of Pareto random variables with $b = 1.0$ and a varied from 0.25 to 0.75. Likewise, Fig. 14 (b) shows the plots corresponding to the probability density function of Pareto random variables with $b = 1.4$ and a varied from 0.25 to 0.75. Notice that with the same location parameter (a), the shapes of the plots corresponding to probability density function of Pareto random variables with the shape parameters $b = 1.0$ and $b = 1.4$ are different.

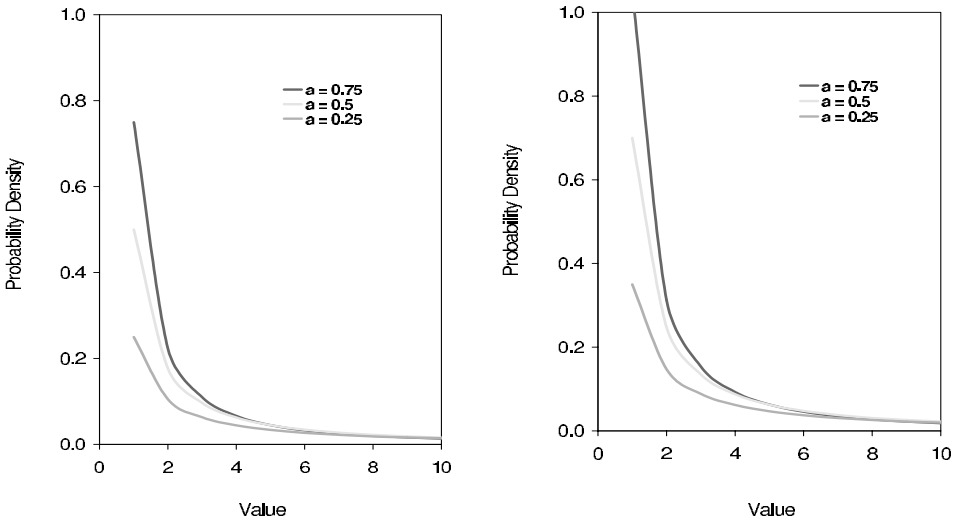

(a) b = 1.0 (b) b = 1.4

Figure 14: Probability density function of Pareto random variables with different location (a) and shape (b) parameters. The location parameters (a) are 0.75, 0.5, and 0.25. The shape parameters (b) are 1.0 and 1.4.

Log-Normal Distribution

The log-normal distribution is the probability distribution of any random variable whose logarithm is normally distributed [36]. Therefore, if X is a random variable with a normal distribution, then $exp(X)$ has a log-normal distribution.

Let X be a random variable with log-normal distribution. Assume also that μ and σ are the mean and the standard deviation of the logarithm of X. The probability density function of the log-normal distribution ($f(x)$) is

$$f(x) = \frac{e^{-((ln((x-\mu)/m))^2/(2\sigma^2))}}{(x-\mu)\sigma\sqrt{2\pi}},\ x \geq \mu;\ m, \sigma > 0,$$

where m is a constant specifying the *scale parameter* of the distribution. μ and σ are also called the *location* and *shape parameters* of the distribution. That is, μ and σ determine the location and the shape of the plot of the probability density function, respectively. The log-normal random variable with $\mu = 0$ and $m = 1$ is called the *standard log-normal random variable*.

From the probability density function, the expected value and the variance of the log-normal distribution can be derived as follows.

$$E[X] = e^{\mu+\frac{1}{2}\sigma^2}$$
$$VAR[X] = (e^{2\mu+\sigma^2})(e^{\sigma^2}-1)$$

Fig. 15 (a-d) illustrates the probability density of four log-normal random variables. The random variables have the same scale parameter ($m = 1$). However, the logarithms of the

random variables have different means (μ) and standard deviations (σ). Fig. 15 (a) and (c) show the plots corresponding to the probability densities of two log-normal random variables with the same mean of the logarithms ($\mu = 0$) but different standard deviations ($\sigma = 0.5$ vs $\sigma = 1$). Similarly, Fig. 15 (b) and (d) show the plots corresponding to the probability densities of two log-normal random variables with the same mean of the logarithms ($\mu = 1$) but different standard deviations ($\sigma = 0.5$ vs $\sigma = 1$). On the other hand, Fig. 15 (a) and (b) illustrate the plots corresponding to the probability densities of two log-normal random variables with the same standard deviation of the logarithms ($\sigma = 0.5$) but different means ($\mu = 0$ vs $\mu = 1$). Likewise, Fig. 15 (c) and (d) show the plots corresponding to the probability densities of two log-normal random variables with the same standard deviation of the logarithms ($\sigma = 1$) but different means ($\mu = 0$ vs $\mu = 1$).

Since the standard deviation of the logarithm (σ) is the shape parameter of the log-normal distribution, the shape of the plot of the probability density alters according to σ. In Fig. 15(a) and (c), the plots of the probability density functions are more right-skewed as σ increases. This is also shown in Fig. 15 (b) and (d).

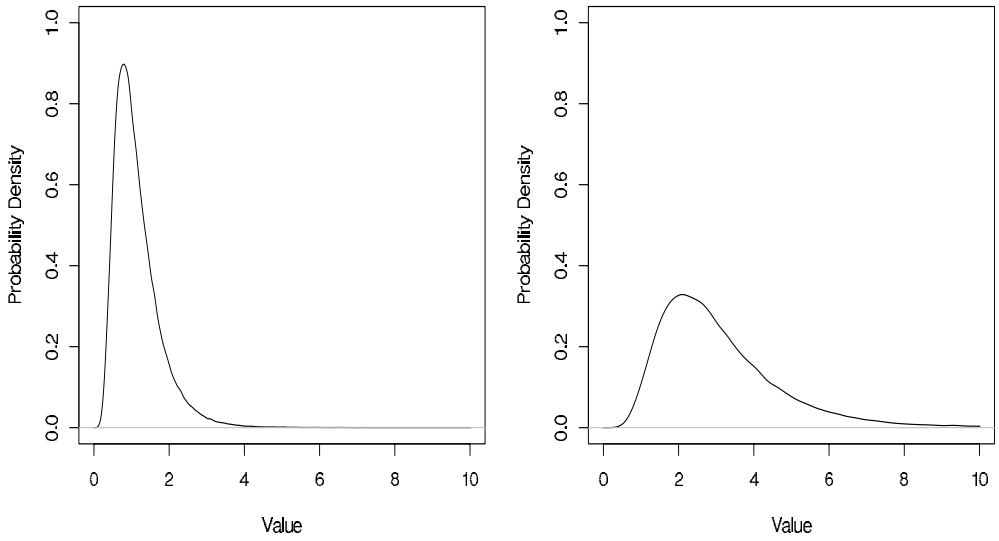

(a) $\mu = 0$, $\sigma = 0.5$

(b) $\mu = 1$, $\sigma = 0.5$

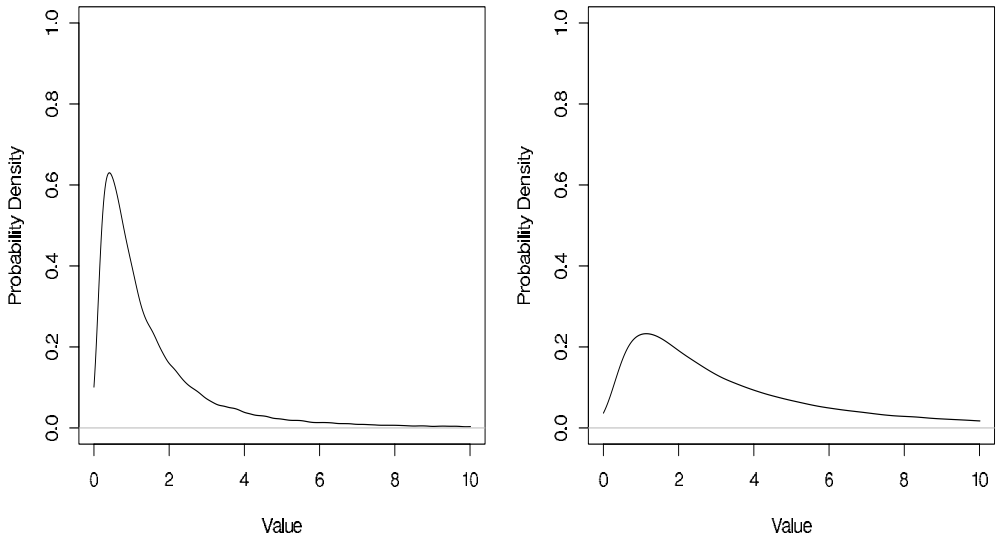

(c) $\mu = 0$, $\sigma = 1$

(d) $\mu = 1$, $\sigma = 1$

Figure 15: Probability density function of log-normal random variables with the same scale parameter ($m = 1$) but different means (μ) and standard deviations (σ) of the logarithm. The means of the logarithms (μ) are 0 and 1. The standard deviations of the logarithms (σ) are 0.5 and 1.0.