

A Queueing-Theoretic Foundation of Available Bandwidth Estimation: Single-Hop Analysis

Xiliang Liu, *Member, IEEE*, Kaliappa Ravindran, and Dmitri Loguinov, *Member, IEEE*

Abstract—Most existing available-bandwidth measurement techniques are justified using a *constant-rate fluid cross-traffic model*. To achieve a better understanding of the performance of current bandwidth measurement techniques in general traffic conditions, this paper presents a queueing-theoretic foundation of single-hop packet-train bandwidth estimation under bursty arrivals of discrete cross-traffic packets. We analyze the statistical mean of the packet-train output dispersion and its mathematical relationship to the input dispersion, which we call the *probing-response curve*. This analysis allows us to prove that the single-hop response curve in bursty cross-traffic *deviates* from that obtained under fluid cross traffic of the same average intensity and to demonstrate that this may lead to significant measurement bias in certain estimation techniques based on fluid models. We conclude the paper by showing, both analytically and experimentally, that the response-curve deviation vanishes as the packet-train length or probing packet size increases, where the vanishing rate is decided by the burstiness of cross-traffic.

Index Terms—Active measurement, bandwidth estimation, packet-pair sampling.

I. INTRODUCTION

AVAILABLE bandwidth of a network path has long been the interest of measurement studies because of its importance to many Internet applications such as overlay routing, server selection, congestion control, and network diagnosis. Several measurement techniques have been developed over the last few years, among which TOPP [17], pathload [8], PathChirp [23], IGI/PTR [6], and Spruce [24] are the major representatives. Most of the current proposals are based on packet-pair or packet-train probing, where bursts of equally spaced packets of uniform size are injected into the path of interest, and the available bandwidth information is inferred from the relationship between input/output inter-packet dispersions.

According to a commonly accepted notion, the available bandwidth of a network hop is its *residual* capacity after transmitting cross traffic. Since at any time instant, the hop is either idle or transmitting packets at its capacity C , the instantaneous link utilization $U(t)$ can be viewed as an ON-OFF function of time, i.e., $U(t) = 0$ if the link is idle at time t and $U(t) = 1$ if the link is busy transmitting a packet at time t . The average

utilization $U(t, t + \delta)$ of the hop during time interval $[t, t + \delta]$ is then given by

$$U(t, t + \delta) = \frac{1}{\delta} \int_t^{t+\delta} U(x) dx. \quad (1)$$

The *hop available bandwidth* $B_\delta(t)$ in $[t, t + \delta]$ is the average unutilized capacity of the hop within that interval, i.e.,

$$B_\delta(t) = C(1 - U(t, t + \delta)). \quad (2)$$

The *available bandwidth of a network path* is the minimum available bandwidth of all traversed links. The link carrying the minimum available bandwidth is called the *tight link*.

Note that $B_\delta(t)$ varies over time t as well as over a wide range of observation intervals δ . These dynamics make it an elusive target to measure. To combat this difficulty, most existing bandwidth-measurement approaches use a constant-rate fluid cross-traffic model to justify the design of their estimation techniques. Under such fluid¹ cross-traffic, $B_\delta(t)$ becomes a constant for all t , and all δ and its relationship to the probing input and output becomes easy to identify.

Although the experimental performance of recent proposals as documented is encouraging, the rationale they are anchored upon is not fully justified in general cross-traffic conditions. To better understand the behavior and performance of existing techniques, this paper presents a queueing-theoretic analytical framework that allows an in-depth analysis of the asymptotic behavior of single-hop packet-train bandwidth estimation under bursty arrivals of discrete cross-traffic packets. Our analysis addresses two fundamental issues. First, given a cross-traffic arrival process and fixed packet-train parameters (i.e., packet size and train length), we demonstrate how the probing output relates to the probing input. We investigate the output rate and dispersion of individual packet-trains as well as their asymptotic average as the number of packet-train samples approaches infinity. We examine the functional dependency between the input and the asymptotic average of the output in the entire input range. We call this relationship the *probing-response curve* and show how the available bandwidth information is embedded in it. Second, we investigate how the response curve evolves with respect to the changes in packet train parameters and cross-traffic burstiness.

Both questions are of central importance for the design of available-bandwidth estimation methods. The answer to the first question provides a theoretical foundation that extends the previous rationale based on fluid cross-traffic models. The answer to the second question offers an insight into parameter tuning strategies in the design of future measurement techniques. Even though published research has produced a great deal of intuition

Manuscript received August 15, 2005; revised June 13, 2006; approved by IEEE TRANSACTIONS ON NETWORKING Editor D. Veitch. This work was supported by NSF Grants CCR-0306246, ANI-0312461, CNS-0434940, and CNS-0519442.

X. Liu is with Bloomberg L.P., New York, NY 10022 USA (e-mail: liuxiliang@gmail.com).

K. Ravindran is with the Computer Science Department, City College of New York, New York, NY 10031 USA (e-mail: ravi@cs.cuny.cuny.edu).

D. Loguinov is with the Computer Science Department, Texas A&M University, College Station, TX 77843 USA (e-mail: dmitri@cs.tamu.edu).

Digital Object Identifier 10.1109/TNET.2007.896235

¹We use the terms “constant-rate fluid” and “fluid” interchangeably.

and empirical findings related to these questions, a mathematically precise explanation of the bandwidth sampling process was not available until now.

While the eventual goal of our analysis is to understand packet-train bandwidth estimation in multihop network paths, single-hop results are indispensable in reaching this goal. Moreover, the single-hop case on its own is an interesting and complex problem calling for an elaborate discussion, which is the focus of this paper. We extend the discussion to multihop paths in a separate paper [14].

Under a theoretically and practically mild assumption, we derive several important properties of the gap (and rate) response curve. Our results show that the gap response curve in constant-rate fluid cross traffic is the tight lower bound of that in bursty cross traffic with the same average intensity. We show that there is an input dispersion range where the real curve positively deviates from its fluid-based prediction. Most existing proposals were designed without being aware of the response deviation phenomenon, which sometimes makes them subject to significant measurement bias.

Our analysis also discovers the source of this deviation and arrives at its closed-form expression in the packet-pair case. We show that the amplitude of the response deviation is exclusively decided by the packet-train parameters and the available bandwidth distribution and that it vanishes as probing packet size or packet-train length increases. We also present an experimental approach to compute with high accuracy the response curves from a given cross-traffic trace. This allows us to empirically validate our theoretical results, qualitatively observe the relationship between the response deviation and packet-train parameters in certain cross-traffic conditions, and evaluate the asymptotic performance of various available-bandwidth estimators.

The rest of the paper is organized as follows. In Section II, we summarize the current measurement proposals and the fluid models they are based upon or related to. In Section III, we present our analytical framework of packet-train bandwidth estimation. Using this framework, we analyze major properties of the response curves and the response deviation phenomenon in Section IV. We provide numerical results of the response deviation and examine its relationship to several deciding factors in Section V. We explain the implications of our findings on some of the current proposals in Section VI. Finally, we present concluding remarks in Section VII.

II. BACKGROUND

IP-layer bandwidth estimation using packet-pairs originates from the seminal work by Bolot [3], Jacobson [7], Keshav [10]. However, due to a lack of consensus on what available bandwidth was and how to measure it, most of the original research efforts in this area went into the measurement of bottleneck capacity [4], [11], [21]. The recent surge of available-bandwidth estimation proposals stems from the fluid model developed in bottleneck capacity estimation research [4], [17]. In what follows, we first briefly introduce this fluid model and then show how the current techniques are related to it.

A. Fluid Model

Consider a single-hop path with capacity C and assume that cross traffic is a fluid with constant arrival rate $\lambda < C$. This fluid

assumption means that for any δ time interval, the amount of cross traffic arriving at the link is $\delta\lambda$. The available bandwidth of the path is $A = C - \lambda$, regardless of the observation time instant or the observation time interval. Consider a probing train of n packets with equal interpacket dispersion g_I and packet size s that passes through the path. The output dispersion g_O can be expressed by the following piecewise linear function of g_I :

$$g_O = \max\left(g_I, \frac{s + \lambda g_I}{C}\right) = \begin{cases} g_I, & g_I \geq \frac{s}{C-\lambda} \\ \frac{s + \lambda g_I}{C}, & g_I \leq \frac{s}{C-\lambda}. \end{cases} \quad (3)$$

Using packet-pairs (i.e., $n = 2$), the practical meaning of (3) is as follows. The first packet p_1 arrives into the hop at time a_1 and experiences zero queueing delay due to the fluid nature of cross traffic. Hence, it departs from the hop at time $a_1 + s/C$. The second packet p_2 arrives into the hop at time $a_1 + g_I$. Before the hop can serve p_2 , the amount of data it has to transmit during the time interval $[a_1, a_1 + g_I]$ is $s + \lambda g_I$. If this is done before p_2 arrives, i.e., $g_I \geq s/(C - \lambda)$, then p_2 also experiences zero queueing delay, and we obtain $g_O = g_I$. Otherwise, the hop undergoes a busy period between the departures of the two packets and the output dispersion is $g_O = (s + \lambda g_I)/C$. A similar argument applies to packet-trains of any length.

Model (3), which we term the *single-hop fluid gap response curve*, has several variants that exhibit more direct associations with available-bandwidth estimation. One such variant is the rate response curve depicting the functional relationship between the input probing rate $r_I = s/g_I$ and the output rate $r_O = s/g_O$:

$$r_O = \min\left(r_I, \frac{C r_I}{r_I + \lambda}\right) = \begin{cases} r_I, & r_I \leq C - \lambda \\ \frac{C r_I}{r_I + \lambda}, & r_I \geq C - \lambda. \end{cases} \quad (4)$$

Since the rate response curve is not linear, Melander *et al.* proposed in [17] to use a transformed version of (4), which depicts the relation between r_I/r_O and r_I :

$$\frac{r_I}{r_O} = \max\left(1, \frac{r_I + \lambda}{C}\right) = \begin{cases} 1, & r_I \leq C - \lambda \\ \frac{r_I + \lambda}{C}, & r_I \geq C - \lambda. \end{cases} \quad (5)$$

We next explain how various types of bandwidth information is embedded in fluid response curves. First, note that each of the three fluid curves contains two segments and that the input rate at the turning point between the two segments is equal to the available bandwidth of the path, i.e., $A = C - \lambda$. Second, observe that the segment in the input rate range (A, ∞) carries information about hop capacity C and cross-traffic arrival rate λ . We next discuss how current techniques are designed to extract available bandwidth from fluid response curves.

B. Measurement Techniques

Recently, Jain and Dovrolis presented a classification of existing measurement techniques [9]. Techniques that use a single input rate are called *direct* and those using multiple input rates are called *iterative*.

Several direct probing methods are Delphi [22], IGI [6], Spruce [24], and the work in [5]. These approaches use one point of the response curve where the input rate r_I is higher than the available bandwidth. They assume that the hop capacity C is known or can be measured separately (e.g., using existing capacity estimation tools such as path rate). Hence, every packet-train sample of these methods can generate an

estimate of the cross-traffic arrival rate during the sampling interval $[a_1, a_1 + (n-1)g_I]$, where a_1 is the arrival time of the first packet in a train to the tight link. These tools obtain a final estimate of λ by averaging multiple probing samples, and the available bandwidth is computed by subtracting the estimated cross-traffic arrival rate from the known link capacity C . Direct probing methods differ among each other in the input probing rate they choose, packet-train parameters, and the assumptions they make on cross-traffic.

Representative iterative methods are TOPP [17], pathload [8], and PTR [6]. Iterative methods do not assume the knowledge of link capacity. They send packet-trains at multiple input probing rates, either to locate the turning point in the response curve (e.g., pathload and PTR) or to extract both C and λ from the linear segment in the input rate range (A, ∞) (e.g., TOPP). Note that compared to PTR, pathload locates the turning point by detecting a trend of increasing one-way delay among the probing packets in a train, instead of comparing the output rate of the packet-train to its input rate. Hence, even though pathload is related to the fluid response curves presented previously, it is not directly based upon them.

C. Discussion

We make several observations regarding the current measurement techniques to motivate our subsequent analysis. First, note that in bursty cross traffic, both the available bandwidth and cross-traffic arrival rate may exhibit a great deal of statistical variability. Most existing techniques produce a single numerical result for each run, which is interpreted as the average available bandwidth within the measurement duration.² Second, current techniques assume that cross-traffic burstiness only causes measurement variability, which can be smoothed out by averaging multiple probing samples. This means that when a large number of packet-train samples are used, the fluid model becomes a valid first-order approximation of the real stochastic process. This assumption can be formalized as the following equality:

$$E[g_O] = \max\left(g_I, \frac{s + \lambda g_I}{C}\right) \quad (6)$$

where term $E[g_O]$ is the statistical mean (or asymptotic average) of output dispersions when a large number of packet-train samples are used. The term λ should be viewed as the long-term average arrival rate of cross-traffic since a large number of packet-train samples naturally extends the measurement duration to a long time period.

However, there has been no analytical investigation regarding the validity of (6) in general traffic conditions. A positive answer to this question would lay a solid ground for the design of available-bandwidth measurement tools and provide them with an assurance of asymptotic accuracy. On the other hand, a negative answer would shed new light on the fundamental limits and tradeoffs in probing-based measurement and give rise to new insights into parameter tuning under diverse application requirements. To tackle this question, we present the necessary analytical framework in the next section.

²Note that pathload is an exception in the sense that it produces a variation range of the δ -interval available bandwidth within the measurement duration, where $\delta = (n-1)g_I$ is the sampling interval of each packet-train.

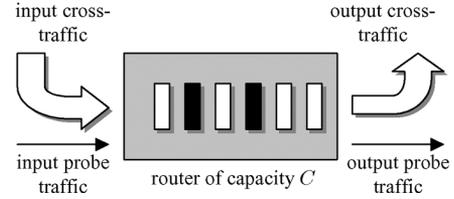


Fig. 1. Single-hop probing model.

III. ANALYTICAL FRAMEWORK OF PACKET-TRAIN PROBING

This paper is an extension of our previous work [12]. Due to limited space, the proofs of several lemmas and theorems are presented in [13] and are omitted from this paper.

We present our analytical framework in four steps. First, we define a set of random processes that model cross-traffic arrival and hop available bandwidth. Second, we introduce a concept we call *probing intrusion residual* to characterize the interaction between probing packets and cross traffic. Using these definitions, we derive in the third step a mathematical relationship between input and output packet-train dispersions. This result calls for a certain understanding of the underlying processes sampled by packet-train arrivals, which is the task of our final step.

A. Cross-Traffic Arrival

Our analysis focuses on the single-hop probing model in Fig. 1. We assume infinite buffer space inside the router, a work-conserving FIFO queueing discipline, and *simple* traffic arrival (i.e., at most one packet arrives at any time instant). We next identify a set of random processes that play a crucial role in modeling packet-train probing.

Definition 1: Cross-traffic is driven by the *packet-counting process* $\{N(t), 0 \leq t < \infty\}$ and the *packet-size process* $\{S_n, 1 \leq n < \infty\}$. The cumulative traffic arrival $\{V(t), 0 \leq t < \infty\}$ is a random process counting the total volume of data³ received by the hop up to time instant t :

$$V(t) = \sum_{n=1}^{N(t)} S_n.$$

Note that $V(t)$ and $N(t)$ are right continuous, meaning that the packet arriving at t is counted in $V(t)$.

Definition 2: We define $\{Y_\delta(t), 0 \leq t < \infty\}$ as the average cross-traffic arrival rate in the interval $(t, t + \delta]$:

$$Y_\delta(t) = \frac{V(t + \delta) - V(t)}{\delta}$$

and call it the δ -interval cross-traffic intensity process.

Definition 3: At time instant t , the *hop-workload process* $\{W(t), 0 \leq t < \infty\}$ is the sum of service times of all packets in the queue and the remaining service time of the packet in service.

Definition 4: We define $\{D_\delta(t), 0 \leq t < \infty\}$ to be the difference between hop workload at times t and $t + \delta$:

$$D_\delta(t) = W(t + \delta) - W(t)$$

and call it the δ -interval workload-difference process.

³In this paper, packet size and data are measured in bits.

Definition 5: The hop utilization process $\{U(t), 0 \leq t < \infty\}$ is an ON-OFF process associated with $\{W(t)\}$:

$$U(t) = \begin{cases} 1, & W(t) > 0 \\ 0, & W(t) = 0 \end{cases} \quad (7)$$

and the δ -interval hop idle process

$$I(t, t + \delta) = I_\delta(t) = \delta - \int_t^{t+\delta} U(x) dx \quad (8)$$

is the total amount of idle time of the forwarding hop in $[t, t + \delta]$. We further call time interval $[t, t + \delta]$ hop busy period if $I_\delta(t) = 0$ and hop idle period if $I_\delta(t) = \delta$.

Definition 6: We define $\{B_\delta(t), 0 \leq t < \infty\}$ as the residual bandwidth in the time interval $[t, t + \delta]$:

$$B_\delta(t) = C \left(1 - \frac{1}{\delta} \int_t^{t+\delta} U(x) dx \right) = \frac{I_\delta(t)C}{\delta} \quad (9)$$

and call it the δ -interval available bandwidth process.

The following lemma describes the relationship among the three important processes $Y_\delta(t)$, $D_\delta(t)$ and $B_\delta(t)$.

Lemma 1: For all positive t and δ , the following holds:

$$B_\delta(t) = C - Y_\delta(t) + \frac{D_\delta(t)C}{\delta}. \quad (10)$$

Bandwidth estimation is essentially interaction of probe packets with *sample-paths* of the processes we just defined. We next examine certain properties of this interaction.

B. Probing Intrusion of Packet Trains

We use quadruple $\langle a_1, g_I, s, n \rangle$ to denote a probing train of n packets p_1, p_2, \dots, p_n , where a_1 is the arrival time of the first packet p_1 at the hop, g_I is the interpacket dispersion at the sender, s is the probe packet size, and n is the train length. Arrival times of probing packets to the hop are denoted by $a_i = a_1 + (i - 1)g_I, i = 1, 2, \dots, n$. Departure times of probing packets from the hop are denoted by $d_i, i = 1, 2, \dots, n$. We define the *output gap* of a packet train as the *average* dispersion between adjacent packets in the train:

$$g_O = \frac{d_n - d_1}{n - 1}. \quad (11)$$

In terms of rate, the input and output probing rates are $r_I = s/g_I$ and $r_O = s/g_O$.

We use $\tilde{W}(t)$ and $\tilde{I}_\delta(t)$ to respectively denote sample paths of workload and hop idle time processes associated with the superposition of cross traffic and probing traffic. Note that this superposition only increases hop workload, i.e., for all t , $\tilde{W}(t) \geq W(t)$. We next define more useful notation to help us examine this intrusion behavior of packet train probing.

Definition 7: The *intrusion range* of probing traffic into $W(t)$ is the set $\{t : \tilde{W}(t) > W(t)\}$. The *intrusion residual function* is $W_d(t) = \tilde{W}(t) - W(t)$.

Function $W_d(t)$ helps us understand the intrusion behavior of the probing traffic into $W(t)$. Before the arrival of probing packets, $W_d(t) = 0$. Upon every arrival of a probe packet, $W_d(t)$ gets an immediate increment of s/C , where s is the probing packet size as before. In $W(t)$'s busy periods without additional probing packet arrival, $W_d(t)$ remains unchanged. In

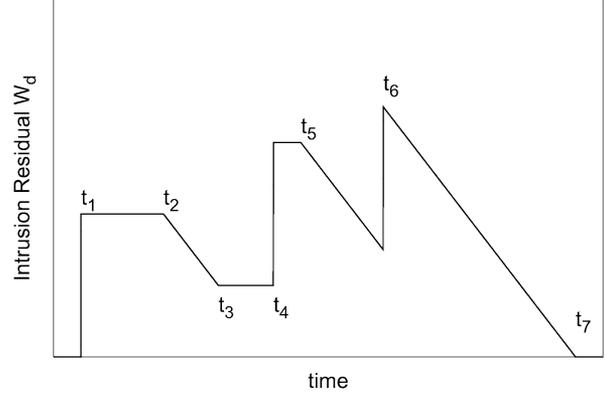


Fig. 2. Illustration of intrusion residual function.

$W(t)$'s idle periods without additional arrival of probe packets, $W_d(t)$ decreases linearly with slope -1 . Function $W_d(t)$ is monotonically nonincreasing between every two adjacent probing packet arrivals. Fig. 2 illustrates this behavior, where (t_1, t_2) and (t_3, t_5) are two busy periods in $W(t)$, and (t_2, t_3) and (t_5, t_7) are two idle periods in $W(t)$. Times t_1, t_4 and t_6 are the instants of probing packet arrivals. Time t_7 is the end point of the intrusion range.

Based on the above observations and assuming a single probe packet of size s arrives to the hop at time t_0 , $W_d(t)$ can be expressed as follows:

$$W_d(t) = \begin{cases} 0, & t < t_0 \\ \max\left(0, \frac{s}{C} - I(t_0, t)\right), & t \geq t_0. \end{cases} \quad (12)$$

When the hop is probed by a packet train $\langle a_1, g_I, s, n \rangle$, we are often interested in computing function

$$R_i(a_1) = W_d(a_i^-) = W_d(a_1 + (i - 1)g_I^-) \quad (13)$$

for $i = 1, 2, \dots, n$, where $f(a^-)$ denotes the left-sided limit $\lim_{x \rightarrow a^-} f(x)$. Metric $R_i(a_1)$ ⁴ is the intrusion residual *caused* by the first $i - 1$ packets in the probing train $\langle a_1, g_I, s, n \rangle$ and *experienced* by packet p_i . In other words, the queuing delay of p_i in the hop is given by

$$\tilde{W}(a_i^-) = W(a_i) + W_d(a_i^-) = W(a_i) + R_i(a_1). \quad (14)$$

The total sojourn time of p_i at the hop is the sum of its service time and its queuing delay

$$d_i - a_i = W(a_i) + R_i(a_1) + \frac{s}{C}. \quad (15)$$

As a direct result of (12), R_i can be recursively computed as follows:

$$R_i = \begin{cases} 0, & i = 1 \\ \max\left(0, \frac{s}{C} + R_{i-1} - I(a_{i-1}, a_i)\right), & i > 1. \end{cases} \quad (16)$$

As shown in (14), the introduction of intrusion residual separates the queuing delay of a probing packet into two portions with different statistical nature. This is one of our key results that make in-depth analysis of packet-train bandwidth estimation possible.

⁴When a_1 is irrelevant, we often write $R_i(a_1)$ as R_i .

C. Output Dispersion of Individual Packet-Trains

Our next lemma expresses the output dispersion of a packet-train from two different angles. This result is the corner stone of our later response-curve analysis.

Lemma 2: Let $\delta = (n-1)g_I$. When a hop with workload process $W(t)$ is probed by a packet train $\langle a_1, g_I, s, n \rangle$, the output gap g_O can be expressed as

$$\begin{aligned} g_O &= \frac{Y_\delta(a_1)g_I}{C} + \frac{s}{C} + \frac{\tilde{I}_\delta(a_1)}{n-1} \\ &= g_I + \frac{D_\delta(a_1)}{n-1} + \frac{R_n(a_1)}{n-1}. \end{aligned} \quad (17)$$

The most interesting feature of Lemma 2 is that its result is unconditional, in the sense that it neither relies on any assumption on the cross-traffic arrival pattern nor imposes any restriction on packet-train parameters. In addition, Lemma 2 develops an avenue towards analytical understanding of the response curve through (especially first-order) statistical properties of each individual terms in (17).

Also note that the four terms $Y_\delta(a_1)$, $D_\delta(a_1)$, $\tilde{I}_\delta(a_1)$, and $R_n(a_1)$ in (17) are generated by sampling four underlying continuous-time sample paths $Y_\delta(t)$, $D_\delta(t)$, $\tilde{I}_\delta(t)$, and $R_n(t)$ at time instant a_1 . Hence, before examining the random samples generated by packet-train probing, we first have to understand several important statistical properties of the underlying continuous-time processes being sampled.

D. Properties of Underlying Processes

We first make an assumption on cross-traffic arrival.

Assumption 1: There exists a constant λ less than hop capacity C such that $V(t)/t \rightarrow \lambda$ as $t \rightarrow \infty$.

This assumption has a series of implications. First, it states that the cross-traffic has a long-term average arrival rate λ . This further leads to the following result.

Lemma 3: The limiting time-average of any δ -interval cross-traffic intensity sample path is equal to λ

$$E[Y_\delta(t)] = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t Y_\delta(u) du = \lambda \quad \forall \delta > 0. \quad (18)$$

Proof: First, notice that

$$\frac{1}{t} \int_0^t Y_\delta(u) du = \frac{\int_t^{t+\delta} V(u) du}{\delta t} - \frac{\int_0^\delta V(u) du}{\delta t}. \quad (19)$$

Computing the limits, we get

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t Y_\delta(u) du = \lim_{t \rightarrow \infty} \frac{\int_t^{t+\delta} V(u) du}{\delta t} = 0. \quad (20)$$

Since $V(t)$ is a nondecreasing function, we can write

$$\delta V(t) \leq \int_t^{t+\delta} V(u) du \leq \delta V(t + \delta). \quad (21)$$

Finally, note that both $\delta V(t)$ and $\delta V(t + \delta)$ have the same limit when divided by δt :

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{V(t)}{t} &= \lim_{t \rightarrow \infty} \frac{\delta V(t)}{\delta t} \leq \lim_{t \rightarrow \infty} \frac{\int_t^{t+\delta} V(u) du}{\delta t} \\ &\leq \lim_{t \rightarrow \infty} \frac{\delta V(t + \delta)}{\delta t} = \lim_{t \rightarrow \infty} \frac{V(t + \delta)}{t + \delta} \frac{t + \delta}{t} \\ &= \lim_{t \rightarrow \infty} \frac{V(t)}{t}. \end{aligned} \quad (22)$$

Combining (20) and (22), we have for $\forall \delta > 0$

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t Y_\delta(u) du = \lim_{t \rightarrow \infty} \frac{V(t)}{t} = \lambda \quad (23)$$

which leads to the statement of the lemma. \blacksquare

Throughout this paper, we adopt sample-path arguments and use the notation of probability expectation as a shorthand representation for sample-path limiting time-average.⁵ Note that the first equality in Lemma 3 is by definition and has nothing to do with ergodicity.

The second important implication of Assumption 1 is the workload stability of the forwarding hop.

Lemma 4: Given Assumption 1, the forwarding hop exhibits workload stability, i.e., $\lim_{t \rightarrow \infty} W(t)/t = 0$.

Proof: See Appendix I. \blacksquare

An immediate consequence of workload stability is the zero-mean nature of $D_\delta(t)$, formally stated as follows.

Lemma 5: If $W(t)/t = o(1)$, the limiting time average of any δ -interval workload-difference sample path is zero:

$$E[D_\delta(t)] = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t D_\delta(u) du = 0 \quad \forall \delta > 0. \quad (24)$$

The next two results concern the available bandwidth process $B_\delta(t)$. Although $B_\delta(t)$ does not explicitly appear in (17), it is functionally related to two processes $\tilde{I}_\delta(t)$ and $R_n(t)$. Furthermore, the process $B_\delta(t)$ is important on its own right as it is the target of our measurement.

Lemma 6: Under the assumptions of this paper, δ -interval available bandwidth converges to $C - \lambda$ as the observation interval becomes large:

$$\lim_{\delta \rightarrow \infty} B_\delta(t) = C - \lambda \quad \forall t > 0. \quad (25)$$

This result implies that a ‘‘good’’ measurement technique should produce $C - \lambda$ as its final estimate given a sufficiently long measurement duration.

Lemma 7: The limiting time-average of any δ -interval available bandwidth process is $C - \lambda$, i.e.,

$$E[B_\delta(t)] = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t B_\delta(u) du = C - \lambda \quad \forall \delta > 0. \quad (26)$$

This result shows that first-order statistics of an available bandwidth sample path do not depend on the observation interval δ and equal the average long-term available bandwidth.

⁵The limiting time-average of a sample path is the expectation of its limiting frequency distribution [19, pp. 45–50]. Hence, it is also called the ‘‘sample-path mean.’’

On the other hand, note that any higher order sample-path statistics of $B_\delta(t)$ have a strong dependence on δ . We define a function $P_\delta(x)$ to describe the available bandwidth distribution along sample path $B_\delta(t)$:

$$P_\delta(x) = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \mathbf{1}(B_\delta(u) \leq x) du. \quad (27)$$

In general, we can assume that the frequency distribution function $P_\delta(x)$ converges to the following step function as $\delta \rightarrow \infty$:

$$P_\infty(x) = \begin{cases} 0, & x < C - \lambda \\ 1, & x \geq C - \lambda. \end{cases} \quad (28)$$

We finish the section by deriving useful bounds on the remaining two terms $\tilde{I}_\delta(a_1)$ and $R_n(a_1)$ in (17) and leave their detailed investigation for the next section. From (16), noticing that $I(a_{i-1}, a_i)$ is no less than zero and applying mathematical induction to i , we get $0 \leq R_n \leq (n-1)s/C$. Combining with Lemma 2, we have.

Corollary 1: Let $\delta = g_I(n-1)$. Then, the following inequalities hold:

$$\frac{D_\delta(a_1)}{n-1} + g_I \leq g_O \leq \frac{D_\delta(a_1)}{n-1} + g_I + \frac{s}{C}, \quad (29)$$

where the second inequality is tight iff $I_\delta(a_1) = 0$.

Our next lemma provides a bound for the term $\tilde{I}_\delta(a_1)$.

Lemma 8: Let $\delta = (n-1)g_I$. Then, we have

$$\begin{cases} 0 \leq \frac{\tilde{I}_\delta(a_1)}{n-1} \leq g_I - \frac{s}{C}, & g_I > \frac{s}{C} \\ \frac{\tilde{I}_\delta(a_1)}{n-1} = 0, & g_I \leq \frac{s}{C}. \end{cases} \quad (30)$$

Collecting Lemmas 2 and 8 leads to the following result.

Corollary 2: When $W(t)$ is probed by packet train $\langle a_1, g_I, s, n \rangle$, the following holds:

$$\begin{cases} g_O = \frac{Y_\delta(a_1)g_I}{C} + \frac{s}{C}, & g_I \leq \frac{s}{C} \\ \frac{Y_\delta(a_1)g_I}{C} + \frac{s}{C} \leq g_O \leq \frac{Y_\delta(a_1)g_I}{C} + g_I, & g_I > \frac{s}{C}. \end{cases} \quad (31)$$

With the basic analytical framework established in this section, we are now in a position to derive the probing-response curve of a single-hop path.

IV. PROBING-RESPONSE CURVES

The probing-response curve depends on a number of factors such as packet-train parameters, intertrain delay pattern, and cross-traffic characteristics. We assume a Poisson arrival of probe trains to the hop, because the asymptotic average of Poisson samples converges to the limiting time average of the sample path being sampled. This property is known as PASTA (Poisson arrivals see time averages) [25]. The average rate of Poisson sampling is assumed to be small enough so that dependency between adjacent trains can be neglected.

We use $\{\{T_m\}, g_I, s, n\}$ to denote a probing train series driven by a Poisson arrival process $\Lambda(t) = \max\{m \geq 0 : T_m \leq t\}$. We use $g_O^{(k)}$ to denote the output gap of the k th probing train $\langle T_k, g_I, s, n \rangle$ in the series, i.e., $g_O^{(k)} = (d_n^{(k)} - d_1^{(k)})/(n-1)$.

The limiting average of the discrete-time sample-path $g_O^{(k)}$ is given by

$$E[g_O] = \lim_{m \rightarrow \infty} \frac{1}{m} \sum_{k=1}^m g_O^{(k)}. \quad (32)$$

We next derive several bounds on the response curve, show examples of applying these bounds, and study packet-train parameter effect on the deviation of the response curve from its fluid bound.

A. Bounds

We first obtain the upper and lower bounds on the gap response curve.

Theorem 1: When $W(t)$ is probed by a Poisson packet-train series $\{\{T_m\}, g_I \leq s/C, s, n\}$, the following holds:

$$E[g_O] = \frac{g_I \lambda}{C} + \frac{s}{C}. \quad (33)$$

Proof: Let $\delta = (n-1)g_I$. Then, using Corollary 2, condition $g_I \leq \frac{s}{C}$ implies

$$E[g_O] = E \left[\frac{g_I Y_\delta(T_m) + s}{C} \right] = \frac{g_I E[Y_\delta(T_m)] + s}{C}. \quad (34)$$

Since $\{T_m\}$ is driven by Poisson arrivals, we have the following due to the PASTA property:

$$E[Y_\delta(T_m)] = E[Y_\delta(t)]. \quad (35)$$

Combining (34), (35), and Lemma 3, we get (33). ■

Rearranging the result of Theorem 1, we get

$$\lambda = \frac{E[g_O]C - s}{g_I} = E \left[\frac{g_O C - s}{g_I} \right] \quad (36)$$

which explains when and why the term $(g_O C - s)/g_I$ can form an unbiased estimator of cross-traffic intensity. We note that this traffic intensity formula is used by several techniques and plays an important role in available bandwidth estimation.

Theorem 2: When $W(t)$ is probed by Poisson packet-train series $\{\{T_m\}, g_I > s/C, s, n\}$, the following holds:

$$\max \left(\frac{g_I \lambda + s}{C}, g_I \right) \leq E[g_O] \leq \min \left(g_I + \frac{g_I \lambda}{C}, g_I + \frac{s}{C} \right).$$

Proof: Notice that from Corollary 2, when $g_I > s/C$

$$\begin{aligned} E[g_O] &\geq \frac{g_I E[Y_\delta(T_m)] + s}{C} = \frac{g_I E[Y_\delta(t)] + s}{C} \\ &= \frac{g_I \lambda + s}{C}. \end{aligned} \quad (37)$$

Similarly, from Corollary 1, PASTA, and Lemma 5, we have

$$E[g_O] \geq g_I + \frac{E[D_\delta(T_m)]}{n-1} = g_I + \frac{E[D_\delta(t)]}{n-1} = g_I. \quad (38)$$

Collecting (37) and (38), we get

$$\max \left(\frac{g_I \lambda + s}{C}, g_I \right) \leq E[g_O]. \quad (39)$$

For the upper bound, from Corollary 2, PASTA, and Lemma 3, we get

$$\begin{aligned} E[g_O] &\leq g_I \left(1 + \frac{E[Y_\delta(T_m)]}{C} \right) \\ &= g_I \left(1 + \frac{E[Y_\delta(t)]}{C} \right) = g_I \left(1 + \frac{\lambda}{C} \right). \end{aligned} \quad (40)$$

Then, from Corollary 1, PASTA, and Lemma 5, we obtain

$$\begin{aligned} E[g_O] &\leq \frac{E[D_\delta(T_m)]}{n-1} + \frac{s}{C} + g_I \\ &= \frac{E[D_\delta(t)]}{n-1} + \frac{s}{C} + g_I = g_I + \frac{s}{C}. \end{aligned} \quad (41)$$

Combining (40) and (41), we get

$$E[g_O] \leq \min \left(g_I \left(1 + \frac{\lambda}{C} \right), g_I + \frac{s}{C} \right). \quad (42)$$

Combining (39) and (42), the theorem follows. \blacksquare

Theorem 2 provides lower and upper bounds on $E[g_O]$ when $g_I > s/C$. Combining this result with the one in Theorem 1 for $g_I \leq s/C$, we get a lower bound on $E[g_O]$ in the entire probing range $0 < g_I < \infty$ as follows:⁶

$$\begin{aligned} L(E[g_O]) &= \begin{cases} \max \left(\frac{g_I \lambda + s}{C}, g_I \right), & g_I > \frac{s}{C} \\ \frac{s + g_I \lambda}{C}, & g_I \leq \frac{s}{C} \end{cases} \\ &= \begin{cases} g_I, & g_I > \frac{s}{C - \lambda} \\ \frac{s + g_I \lambda}{C}, & g_I \leq \frac{s}{C - \lambda}. \end{cases} \end{aligned} \quad (43)$$

This is exactly model (6) we are trying to validate. However, Theorem 2 shows that (6) is a *lower bound* of $E[g_O]$, which does not necessarily equal to $E[g_O]$. Likewise, we can obtain from the two theorems the entire upper bound of $E[g_O]$ as follows:

$$\begin{aligned} U(E[g_O]) &= \begin{cases} \min \left(g_I + \frac{g_I \lambda}{C}, g_I + \frac{s}{C} \right), & g_I > \frac{s}{C} \\ \frac{s + g_I \lambda}{C}, & g_I \leq \frac{s}{C} \end{cases} \\ &= \begin{cases} \frac{s}{C} + \frac{g_I \lambda}{C}, & g_I \leq \frac{s}{C} \\ g_I + \frac{g_I \lambda}{C}, & \frac{s}{C} \leq g_I \leq \frac{s}{\lambda} \\ g_I + \frac{s}{C}, & g_I \geq \frac{s}{\lambda}. \end{cases} \end{aligned} \quad (44)$$

The real gap response curve is contained between these two bounds. We define the response deviation $\beta(g_I, s, n)$ as the difference between the real gap response curve and the lower bound given by (43). It can be expressed by the following using Theorem 2, Lemma 2, and PASTA (where $\delta = (n-1)g_I$):

$$\beta(g_I, s, n) = \begin{cases} \frac{1}{n-1} E[\tilde{I}_\delta(t)], & g_I \leq \frac{s}{C-\lambda} \\ \frac{1}{n-1} E[R_n(t)], & g_I \geq \frac{s}{C-\lambda}. \end{cases} \quad (45)$$

We next provide a closed-form expression for the packet-pair (i.e., $n = 2$) response deviation to gain additional insight into the nature of this phenomenon.

⁶Functions $L(f)$ and $U(f)$ denote lower and upper bound on f , respectively.

B. Packet-Pair Closed-Form Expression

Notice that both $R_n(t)$ and $\tilde{I}_\delta(t)$ can be expressed as *deterministic* functions whose arguments are packet-train parameters and the following $(n-1)$ -dimensional vector:

$$\vec{B}_\delta^{(n-1)}(t) = \begin{pmatrix} B_\delta(t) \\ B_\delta(t + \delta) \\ \vdots \\ B_\delta(t + (n-2)\delta) \end{pmatrix}. \quad (46)$$

In particular, when $n = 2$, δ is simply g_I and the vector (46) degenerates to a scalar, which allows us to obtain simple expressions for $\tilde{I}_\delta(t)$ and $R_2(t)$:

$$\tilde{I}_\delta(t) = \max \left(0, \frac{B_\delta(t)\delta - s}{C} \right) \quad (47)$$

$$R_2(t) = \max \left(0, \frac{s - B_\delta(t)\delta}{C} \right). \quad (48)$$

Consequently, we have the following result regarding the packet-pair probing-response curve.

Theorem 3: Assuming that $W(t)$ is probed by Poisson packet-pair series $\{\{T_m\}, g_I, s, 2\}$, let $\delta = g_I$ and denote by $P_\delta(x)$ the frequency distribution function of sample-path process $B_\delta(t)$. Then, the following holds:

$$\begin{aligned} E[g_O] &= \frac{g_I \lambda + s}{C} + \int_{s/\delta}^C \frac{x\delta - s}{C} dP_\delta(x) \\ &= g_I + \int_0^{s/\delta} \frac{s - x\delta}{C} dP_\delta(x). \end{aligned} \quad (49)$$

It immediately follows that the packet-pair response deviation is

$$\beta(g_I, s, 2) = \begin{cases} \frac{1}{C} \int_{s/\delta}^C (x\delta - s) dP_\delta(x), & g_I < \frac{s}{C-\lambda} \\ \frac{1}{C} \int_0^{s/\delta} (s - x\delta) dP_\delta(x), & g_I \geq \frac{s}{C-\lambda}. \end{cases} \quad (50)$$

The response deviation phenomenon is one of the previously unknown factors that can cause measurement bias for available bandwidth estimation techniques based on (6). Even though a closed-form expression of the response deviation for packet-trains is complex in general, it is clear that the amount of deviation is exclusively decided by the packet-train parameters g_I , s , n and the sample-path frequency distribution of the available bandwidth vector (46).

Next, we show the full picture of response curves for both gap-based and rate-based versions.

C. Full Picture

We now investigate the relationship between the response deviation given in (45) and the input gap g_I while keeping all other parameters fixed. We first examine packet-pair probing.

Theorem 4: When $W(t)$ is probed by Poisson packet-pair series $\{\{T_m\}, g_I, s, 2\}$, the response deviation $\beta(g_I, s, 2)$ equals zero when input gap $g_I \in (0, s/C]$; it is a monotonically increasing function of g_I in the input gap range $(s/C, s/(C-\lambda)]$; and it is a monotonically decreasing function of g_I in the input

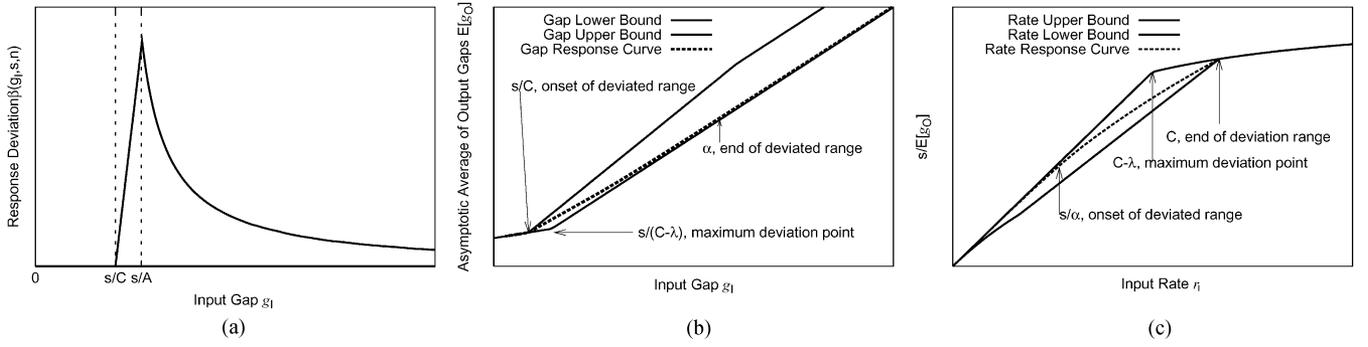


Fig. 3. Illustration of various properties of response curves in the entire input range. (a) Response deviation, (b) gap-response curve, and (c) rate-response curve.

gap range $(s/(C - \lambda), \infty)$. Furthermore, the response deviation $\beta(g_I, s, 2)$ monotonically converges to 0 as g_I approaches infinity. Finally, in the whole input-gap range $(0, \infty)$, the response deviation is a continuous function of g_I .

Packet-pair response deviation has very nice functional properties in terms of continuity and monotonicity. The deviation $\beta(g_I, s, 2)$ is a hill-shaped function with respect to g_I as shown in Fig. 3(a), where it reaches its maximum when $g_I = s/(C - \lambda)$. We also conjecture that the packet-train response deviation $\beta(g_I, s, n)$ is also continuous and has similar monotonicity properties described in Theorem 4. In Section V, we experimentally observe this fact.

In summary, the response deviation is significant only in the middle part of the probing range. We call this range the *deviation probing range*. The full picture of the gap response curve is illustrated in Fig. 3(b). The entire probing range $(0, \infty)$ is divided into three portions. Interval $(0, s/C]$ is the *nondeviation* region where traffic intensity formula (36) forms an unbiased estimator of λ . Interval $(s/C, \alpha)$ is a *deviation* region where $E[g_O]$ is larger than (6), but smaller than the upper bound in (44). Finally, interval (α, ∞) is the second *nondeviation* probing range where $E[g_O] = g_I$. Theoretically, α can be infinitely large and this range often does not exist. Practically, however, a sufficiently small response deviation can be assumed to be zero. Input dispersion $g_I = s/(C - \lambda) = A$ is the point where the response deviation is maximized and deviation offset point α is never equal to A . Further note that the upper bound on the gap-response curve given in (44) is actually not tight.

It is often more informative to look at the rate version of the response curve rather than the gap version, because the former has a direct association with our measurement targets: traffic intensity and available bandwidth. Transforming (43) into the corresponding rate version, we get the rate upper bound⁷

$$U\left(\frac{s}{E[g_O]}\right) = \begin{cases} r_I, & 0 < r_I \leq C - \lambda \\ C \frac{r_I}{r_I + \lambda}, & r_I > C - \lambda. \end{cases} \quad (51)$$

Transforming (44) gives us the rate lower bound as

$$L\left(\frac{s}{E[g_O]}\right) = \begin{cases} \frac{r_I}{r_I + C} C, & 0 < r_I \leq \lambda \\ \frac{r_I}{\lambda + C} C, & \lambda < r_I \leq C \\ \frac{r_I}{r_I + \lambda} C, & C < r_I. \end{cases} \quad (52)$$

⁷Note that existing tools compute the average output rate as $s/E[g_O]$ instead of $E[s/g_O]$. This is because the rate curve is derived from the gap curve and $E[s/g_O]$ cannot be determined from $E[g_O]$.

As illustrated in Fig. 3(c), along the vertical direction, the rate response curve appears between the two bounds given above. Along the horizontal direction, the curve shows one negatively deviating probing region sandwiched by two nondeviation regions.

D. Impact of Packet Train Parameters

First, we examine the impact of probing packet size on response deviation. At any fixed input rate r , let $s \rightarrow \infty$. This causes the sampling interval $\delta = (n - 1)s/r$ to approach infinity at rate proportional to s . The following theorem states a sufficient condition for the response deviation to vanish as s increases.

Theorem 5: At any input rate r , the response deviation vanishes as the packet size s increases if the following condition holds:

$$P_\delta(r) - P_\infty(r) = o\left(\frac{1}{\delta}\right) \quad (53)$$

where $P_\delta(x)$ is the frequency distribution function of $B_\delta(t)$ and $P_\infty(x)$ is the step function given in (28).

Proof: We first prove the case when $n = 2$ and $r < C - \lambda$. In this situation, let $\delta = g_I$ and observe that

$$\beta\left(\frac{s}{r}, s, 2\right) = \int_0^r \frac{s - x\delta}{C} dP_\delta(x) = \frac{\delta}{C} \int_0^r P_\delta(x) dx. \quad (54)$$

Hence, a sufficient and necessary condition for packet-pair response deviation at input rate $r < A$ to vanish as $s \rightarrow \infty$ is

$$\lim_{\delta \rightarrow \infty} \delta \int_0^r P_\delta(x) dx = 0 \quad (55)$$

which can be satisfied when (53) holds because of the following inequality, where we should recall that $P_\infty(r) = 0$:

$$0 \leq \delta \int_0^r P_\delta(x) dx \leq \delta r (P_\delta(r) - P_\infty(r)). \quad (56)$$

Similarly, for any input rate $r \in (C - \lambda, C]$, a sufficient and necessary condition for packet-pair response deviation to vanish is

$$\lim_{\delta \rightarrow \infty} \delta \left(C - r - \int_r^C P_\delta(x) dx \right) = 0 \quad (57)$$

which can be satisfied when (53) holds because of the following inequality, where $P_\infty(r) = 1$:

$$0 \leq \delta \left(C - r - \int_r^C P_\delta(x) dx \right) \leq \delta(C - r) (P_\infty(r) - P_\delta(r)). \quad (58)$$

For the case of packet-train probing where $n > 2$, we refer the reader to Theorem 8 in [13] for a detailed proof. ■

Without getting into technicalities, we point out that (53) is also a necessary condition for the response deviation to vanish as s increases, given that $P_\delta(x)$ is continuous. Note that many cross-traffic arrivals produce a regenerative workload (and consequently a regenerative link utilization process) in the forwarding hop (e.g., Poisson and ON-OFF traffic used in the experiments of this paper). In this case, we can apply the regenerative central limit theorem [26, p. 124] to show that the frequency distribution function of $B_\delta(t)$ converges exponentially to the step function (28), which is much faster than the convergence speed required by (53). Therefore, larger probing packet size implies less response deviation in these cases.

Theorem 6: When hop utilization process $\{U(t)\}$ is regenerative [26 p. 89], $P_\delta(r) - P_\infty(r)$ is an asymptotically exponential function of δ , i.e., there exists a positive constant k such that

$$P_\delta(r) - P_\infty(r) = \Theta(e^{-k\delta}). \quad (59)$$

Proof: See Appendix II. ■

To examine the impact of packet-train length, we first show that when $n \rightarrow \infty$, the output gap converges to the fluid prediction almost surely.

Theorem 7: Given Assumption 1 and an arbitrary input gap g_I , the output gap converges with probability 1 to the fluid model (3) as n increases.

Proof: We first consider the case when $g_I > s/(C - \lambda)$. As $n \rightarrow \infty$, the aggregated traffic has a long-term rate $\lambda + s/g_I < C$. Hence, according to Lemma 4, our queueing system is stable and produces the following:

$$\lim_{n \rightarrow \infty} \frac{\tilde{W}(a_1 + (n-1)g_I)}{n-1} = 0. \quad (60)$$

Further note that

$$0 \leq R_n(a_1) \leq \tilde{W}(a_1 + (n-1)g_I). \quad (61)$$

Dividing by $n-1$ and taking the limit of (61), we get

$$\lim_{n \rightarrow \infty} \frac{R_n(a_1)}{n-1} = 0. \quad (62)$$

Combining (17) and (62), we have

$$\lim_{n \rightarrow \infty} g_O = g_I. \quad (63)$$

Next, consider the case when $g_I < s/(C - \lambda)$. As $n \rightarrow \infty$, the aggregated traffic has a long-term rate $\lambda + s/g_I > C$, leading to an unstable queue that grows unbounded:

$$\lim_{n \rightarrow \infty} \frac{\tilde{W}(a_1 + (n-1)g_I)}{a_1 + (n-1)g_I} = \lambda + s/g_I - C. \quad (64)$$

Omitting certain details, we can show that there exists a constant M such that

$$\lim_{n \rightarrow \infty} \tilde{I}(a_1, a_1 + (n-1)g_I) \leq M. \quad (65)$$

Combining (17), (18), and (65), we get

$$\lim_{n \rightarrow \infty} g_O = \frac{\lambda g_I + s}{C}. \quad (66)$$

Collecting both cases, we have proved the theorem. ■

Note that, in general, Theorem 7 does not necessarily imply a vanishing response deviation as $n \rightarrow \infty$, which requires proving a mean convergence rather than an almost-surely convergence. However, the two convergence modes are equivalent when taking into account a practical factor that the output gap is uniformly bounded by a constant for any packet-train length n . This is because, in practice, the cross-traffic arrival rate is bounded by the capacities of incoming links at a given router. Suppose that the sum of all incoming link capacities at the hop is C_+ . Then, $Y_\delta(t)$ is distributed in a finite interval $[0, C_+]$. Consequently, the output gap g_O is also distributed in a finite interval $[s/C, \max(g_I, (s + g_I C_+)/C)]$. Given this constraint, almost surely convergence of Theorem 7 implies a mean convergence, hence a vanishing response deviation when n tends to infinity.

E. Discussion

We now briefly mention how sensitive our results are with respect to the assumptions made in this paper. First, this paper assumed infinite buffer space in the hop. Hence, our results are valid when buffer space is sufficiently large and packet loss can be neglected. Otherwise, equality $A = C - \lambda$ becomes invalid. Analysis of the impact of buffer size on bandwidth estimation requires future work.

Second, we assumed a Poisson intertrain probing pattern. This can be relaxed to more general ASTA (Arrivals see time averages) [16] sampling. Recent theoretical progress [2] showed that in nonintrusive probing, ASTA was not unique to Poisson arrivals, but was shared by a large class of other sampling processes. Recall that our work assumed sufficiently large intertrain delays (which is called *rare probing* in [2]). This supports nonintrusiveness of our probing process and makes our results applicable to a large number of non-Poissonian sampling patterns.

Finally, we made a sample-path assumption on cross-traffic and avoided the cross-traffic stationarity assumption, which was commonly agreed upon in prior work. Our results are applicable, but not limited to, stationary cross-traffic. For additional comments on this issue, we refer the reader to [13].

Next, we present our experimental methodology for computing the probing-response curve and study the deviation function experimentally.

V. NUMERICAL RESULTS

In this section, we first introduce an offline algorithm that can compute with high accuracy the probing-response curves from a given cross-traffic arrival trace. We then apply this method to several types of cross traffic with different characteristics to examine the quantitative relationship between the response deviation and the packet-train parameters.

A. Computing Response Curves

Our offline algorithm computes single-hop response curves based on cross-traffic packet arrival trace, packet-train parameters (s, n) , and hop capacity C . Traffic traces provide arrival time and packet size for every cross-traffic packet at the hop. Given a sufficiently long trace, the frequency distributions of the associated sample paths (such as $Y_\delta(t)$, $D_\delta(t)$, and $B_\delta(t)$) in that time interval become good approximations of their limiting frequency distributions. Our offline algorithm approximates the sample-path mean of $g_O^{(k)}$ for any given input spacing g_I . Next, we briefly explain the structure of this algorithm.

We use Υ_0^τ to denote a cross-traffic trace in the time interval $[0, \tau]$. Given Υ_0^τ and hop capacity C , the hop workload sample path $W(t)$ in the interval $[0, \tau]$, denoted as $W(t)_0^\tau$, can be computed. The following corollary states basic properties of the workload sample path.

Corollary 3: Hop workload sample path consists of alternating busy periods and idle periods. Any busy period comprises piecewise linear segments with slope -1 .

Taking advantage of these functional properties and using a proper data structure, we can represent $W(t)_0^\tau$ without losing any of its information. Furthermore, we are able to retrieve $Y_\delta(t)$, $D_\delta(t)$, and $B_\delta(t)$ for any t in $[0, \tau - \delta]$. In other words, we keep the full information about $Y_\delta(t)_0^{\tau-\delta}$, $D_\delta(t)_0^{\tau-\delta}$, and $B_\delta(t)_0^{\tau-\delta}$ in the data structure of $W(t)_0^\tau$.

Instead of approximating $E[g_O^{(k)}]$ using a finite number of output dispersion samples, we approximate $E[g_O(t)]$, the corresponding continuous-time sample-path mean, using the time average of $g_O(t)$ in a finite interval. Note that due to the ASTA assumption, $E[g_O^{(k)}] = E[g_O(t)]$. Hence, a good approximation of the latter sample-path mean also serves as a good approximation for the former. The continuous-time sample path $g_O(t)$ also has certain “nice” properties as we state in the next theorem. The proof is in constructive terms, which provides a concrete idea of how our offline algorithm is designed.

Definition 8: Event points are the time instants at which the workload sample path switches from a busy period to an idle period or undergoes a sudden increment due to packet arrival. An *intervent interval* is the interval between two adjacent event points.

Theorem 8: The sample path $g_O(t)$ consists of piecewise linear segments with possible slopes 0, 1 and -1 . For any two time instants $0 < t_1 < t_2$, $g_O(t)$ is continuous in the interval (t_1, t_2) given that 1) t_1 and t_2 fall into the same intervent interval of $W(t)$ and 2) $t_1 + (n-1)g_I$ and $t_2 + (n-1)g_I$ fall into the same intervent interval of $W(t)$.

Proof: See Appendix I in [15]. ■

Our algorithm computes $g_O(t)_0^{\tau-(n-1)g_I}$, the sample path $g_O(t)$ in the time interval $[0, \tau-(n-1)g_I]$, based on Υ_0^τ , C , s , n , and g_I . The computation makes use of the second formula in (17), where R_n is computed recursively using (16). Furthermore, taking advantage of Theorem 8, we can represent the sample-path information of $g_O(t)_0^{\tau-(n-1)g_I}$ to its full precision in a proper data structure. We then compute the following:⁸

$$\frac{1}{\tau - \delta} \int_0^{\tau-(n-1)g_I} g_O(t) dt \quad (67)$$

⁸Theorem 8 also allows an efficient computation of (67) with high accuracy.

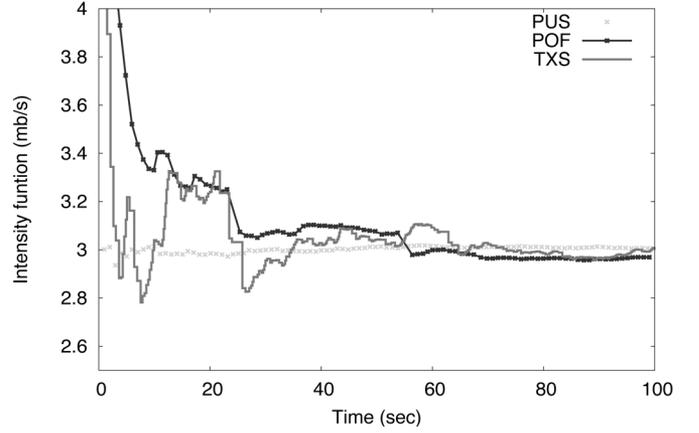


Fig. 4. Intensity function $\mathcal{I}(t) = V(t)/t$ for the three traffic traces.

and use it as an approximation to

$$E[g_O(t)] = \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \int_0^\tau g_O(u) du. \quad (68)$$

It is clear that the precision of this approximation is mainly decided by τ . Thus, we can pick a large enough τ so that its further increase would make little difference. It could sometimes be impractical to have such a long trace; however, note that even when (67) is not a good approximation of $E[g_O(t)]$, it still represents a correct result in a hypothetical periodic cross-traffic that repeats itself after every τ time units. This is due to the fact that in periodic cross traffic, the sample path $g_O(t)$ has a limiting time average equal to its time average in one period.

B. Traffic Traces

We compute the probing-response curves using three different cross-traffic types: Poisson traffic with packet sizes (in bytes) uniformly distributed in $[1, 1500]$ (PUS), Pareto ON-OFF traffic (POF), and a real traffic trace TXS-1148742649 (TXS) from the National Laboratory for Applied Network Research (NLNR). Hop capacity C is fixed at 10 mb/s. The cross-traffic packet size is 750 bytes for POF. The average sending rate is 500 packets per second for PUS. The mean duration of POF ON-OFF periods is 10 and 5 ms, respectively. The Pareto shape parameter α for the duration of both ON-OFF periods is set to 1.9. In POF ON periods, the source sends CBR traffic at 750 packets per second. Given these settings, Both PUS and POF cross-traffic satisfy Assumption 1 and have a long-term arrival rate λ equal to 3 mb/s. To facilitate comparison, we scale the packet interarrival times in the TXS trace by a common factor so that the average traffic arrival rate within the trace duration is also 3 mb/s. We use random-number generators to produce two packet-arrival traces for PUS and POF. These traces record the time instants of all packet arrivals and their sizes within a period of 100 s.

In Fig. 4, we plot the traffic intensity function $\mathcal{I}(t) = V(t)/t$ for the three traffic traces. As shown in the figure, all traffic types have the same long-term arrival rate 3 mb/s. However, there are prominent differences in their convergence delays. Traffic trace PUS converges very fast, in 10 s, reaching within 1% of $\lambda = 3$ mb/s. The convergence delay of POF and TXS, however, is much longer and equals about 60 s for the 1.5%-neighborhood

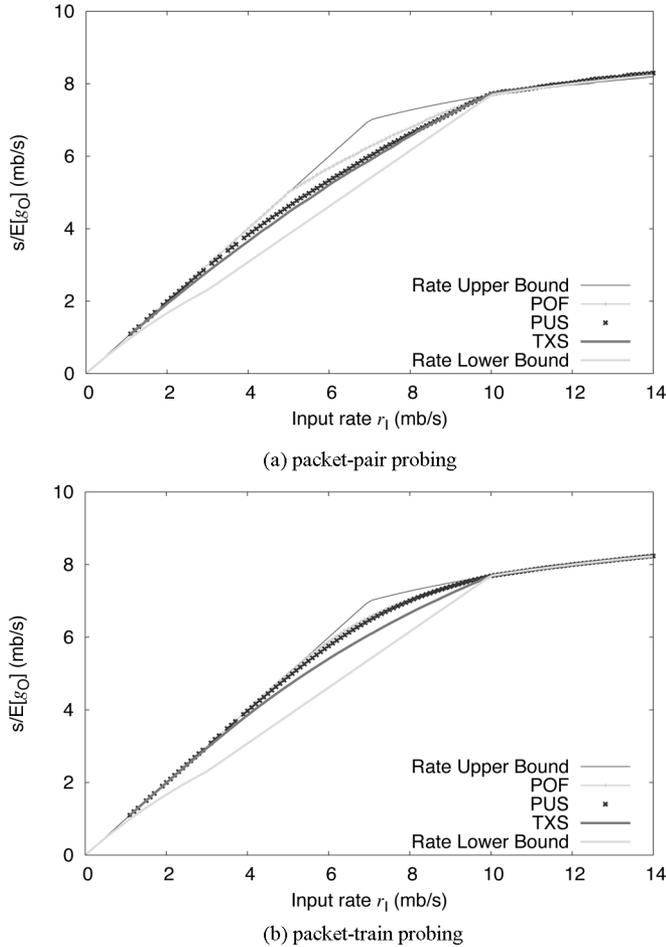


Fig. 5. Rate response curve for the three cross-traffic traces. (a) Probing pairs and (b) 16-packet trains (probing packet size is 750 bytes).

of 3 mb/s. Based on these cross-traffic characteristics, we choose $\tau = 20$ s for PUS and $\tau = 60$ s for POF and TXS.

In what follows, we first compute response curves for several fixed packet-train parameters. We then provide numerical results for the response deviation for a range of packet-train parameters to demonstrate their quantitative relationship. For each response curve, we compute the mean output dispersion $E[g_O]$ at 140 equally spaced input rates from 1 to 14 mb/s. For each input rate, we apply our offline algorithm to compute the numerical integration of (67) as an approximation of $E[g_O]$.

C. Results and Discussion

Fig. 5(a) shows rate response curves for the three traces when the hop is probed using 750-byte packet pairs. Notice in the figure that all three curves substantially deviate from the fluid upper bound. The curve of TXS appears lower than those of PUS and POF, indicating that TXS has the most deviation from the fluid bound of the three traces. It is also interesting to note that POF is much closer to the upper bound than PUS, which means that the former suffers *less* response deviation than the latter. This indicates that for fixed packet-train parameters, cross traffic of more burstiness does not necessarily imply larger response deviation. We explain the reasons for this in a short while.

Fig. 5(b) shows the rate response curves for the three traces when the hop is probed using probing trains of 16 packets. Com-

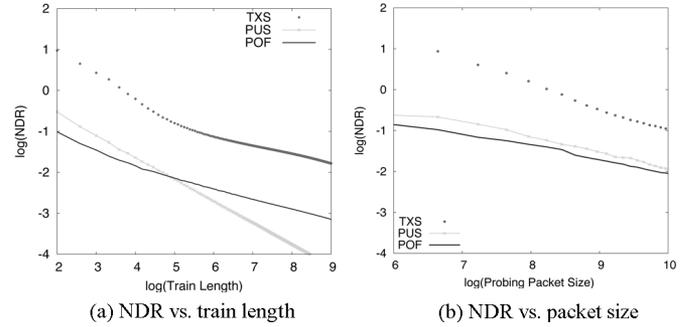


Fig. 6. Log-scale NDR for the three cross-traffic traces. (a) Probing train length from 2 to 512 and (b) probing packet size from 50 bytes to 1500 bytes.

pared with the previous figure, all three response curves are closer to the fluid upper bound. However, one observation is that PUS approaches the fluid bound much quicker than the curves of the other two traces. This shows that, as the probing train length increases, the response deviation diminishes at a rate that depends on the burstiness of cross traffic.

Since we constantly observe that the response curves suffer the largest deviation when the input rate equals to the available bandwidth, we define a metric called NDR (normalized deviation ratio) to characterize the amount of deviation in a rate response curve. Let \mathbf{r} be the output rate $s/E[g_O]$ when the input rate is $A = C - \lambda$. We define

$$\text{NDR} = (A - \mathbf{r}) \left(\mathbf{r} - \frac{AC}{C + \lambda} \right)^{-1} \quad (69)$$

which is the distance of the actual curve to its upper bound divided by the distance to its lower bound, given that the input probing rate is equal to the available bandwidth A . The NDR metric takes values in $[0, \infty)$, where larger NDR values indicate more deviation in the response curve. We next investigate the relationship between the NDR and packet-train parameters.

For all three traces, we computed the NDR using probing packet sizes between 50 and 1500 bytes with a 50-byte step and probing train lengths between 2 and 512 packets with a two-packet step. Thus, in total, we have $256 \times 30 = 7680$ different packet-train parameter settings for each of the three traces. For each parameter setting, we calculate the output rate \mathbf{r} in (69) using our offline algorithm.

Fig. 6(a) shows the NDR for the three traces using $s = 750$ bytes. In all cases, the NDR decreases as the probing train length increases and this relationship appears to be a power-law function as confirmed by our log-log plot. Fig. 6(b) shows the NDR for a fixed train-length of 16 packets and varying probing packet size from 50 to 1500 bytes. We again observe a power-law decrease of NDR with respect to the increase in the probing packet size. Modeling the relationship between NDR and packet-train parameters s and n using function $\text{NDR} = k/(s^{\alpha_1} n^{\alpha_2})$, where k , α_1 , and α_2 are cross-traffic related parameters, we get

$$\log(\text{NDR}) = \log(k) - \alpha_1 \log(s) - \alpha_2 \log(n). \quad (70)$$

We plot 3-D charts of $\text{NDR}(s, n)$ and their contours on log-log scale for all three traces to gain more insight into this relationship. Fig. 7 shows two of them, i.e., NDR planes for PUS and TXS. We use 3-D-fitting to find the parameters of the

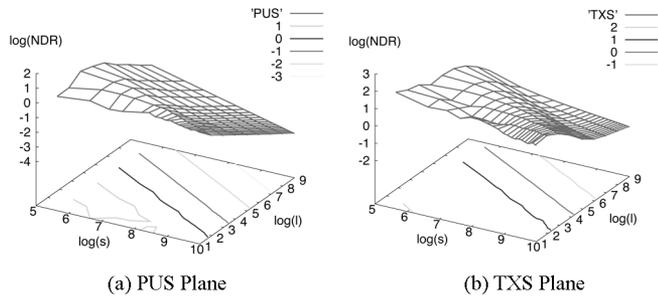


Fig. 7. NDR(s, n) planes and contours for PUS and TXS. (a) PUS plane and (b) TXS plane.

TABLE I
3-D-FITTING RESULTS FOR NDR PLANES

Method	α_1	α_2	$\log(k)$
PUS	0.524 ± 0.008	0.539 ± 0.004	6.111 ± 0.077
POF	0.413 ± 0.007	0.338 ± 0.003	4.000 ± 0.074
TXS	0.341 ± 0.016	0.379 ± 0.007	5.019 ± 0.143

three planes, where all least-square fitting errors are less than 5%, indicating that the power-law function (70) is a reasonable model for NDR. Curve-fitting results are given in Table I, which shows that traffic with more burstiness has smaller values of α_1 and α_2 . This explains why the response deviation in POF and TXS is harder to overcome than that in PUS.

The experimental results we obtained agree with our analytical findings very well. Furthermore, they show that with fixed packet-train parameters, more cross-traffic burstiness does not necessarily imply more response deviation. However, in the former case, this deviation is more difficult to overcome by increasing the probing packet size or packet-train length.

To understand this phenomenon, recall that the notion of traffic burstiness in this paper relates to how fast the traffic becomes “smooth” with respect to the increase of observation intervals rather than how “smooth” the traffic appears in a given fixed observation interval. Hence, it is normal that for a given observation interval, POF has smaller variance than Poisson traffic and appears “smoother,” which leads to less response deviation when packet-trains are constructed to sample the traffic in such an observation interval. As the train length or packet size increases, the observation interval increases and Poisson traffic becomes smooth quicker than POF. Therefore, the response deviation also vanishes quicker.

Even though we do not offer a precise interpretation for the power-law relation between the NDR metric and packet-train parameters, we believe that it is related to the evolving trend of available bandwidth frequency distribution with respect to the increase of observation interval. This view is supported by the fact that the response deviation is exclusively decided by the packet-train parameters and the available bandwidth distribution. This implies that there is no other factor that can decide the NDR metric.

VI. IMPLICATIONS

Among the five representative proposals TOPP, IGI/PTR, Spruce, pathload, and pathChirp, the first three directly fall under the umbrella of our work, while the last two techniques

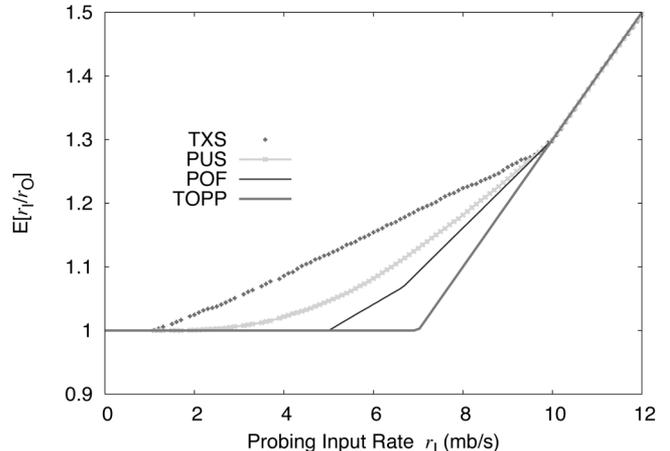


Fig. 8. TOPP-transformed rate response curves.

TABLE II
TOPP RESULTS (IN mb/s) USING DEVIATION SEGMENTS IN FIG. 8 (CORRECT VALUES: $C = 10$ mb/s, $A = 7$ mb/s)

Method	Estimated C	Estimated λ	Estimated A
PUS	32.51	29.24	3.28
POF	23.38	18.36	5.02
TXS	29.59	28.16	1.43

have quite a few tunable parameters and their behavior is beyond the scope of this paper.

A. TOPP

Fig. 8 shows the rate response curves for the three traces when the hop is probed using 1500-byte packet-pairs (as suggested in [18]). The curves are transformed to depict the relationship between $E[r_I/r_O]$ and r_I . TOPP applies segmented linear regression on this transformed curve to obtain the hop capacity and cross-traffic intensity information. In the order of closeness to TOPP’s expected piece-wise linear curve (i.e., fluid lower bound) appear the response curves of POF, PCS, and TXS. TOPP uses the second linear segment of the measured curve assuming that it contains hop information. However, in practice, unless the deviation is very small and undetectable, the second segment may not be linear (such as in the figure) and may not coincide with the fluid bound, which may mislead TOPP into believing that the curve in the deviation range is the second linear segment of the fluid model. In Fig. 8, all deviation ranges are very clear and will be incorrectly acted upon by TOPP. Table II shows the results of a linear regression applied to deviating response curves in Fig. 8 using the basic algorithm in TOPP. As the table shows, the available bandwidth is significantly *underestimated*, especially for the real traffic trace TXS. Both the hop capacity and cross-traffic intensity are significantly *overestimated*. Hence, to assure asymptotic accuracy, TOPP has to apply additional techniques to bypass the segment in the deviating input range.

B. PTR

PTR uses output rate $s/E[g_O]$ at the point of transition between the two linear segments in the response curve as an estimate of available bandwidth. As we have established, this point corresponds to the input rate at which the stochastic curve starts deviating from the fluid bound and does not usually represent

the available bandwidth of the path. Since this point is always smaller than or equal to available bandwidth, PTR is a *negatively biased* available-bandwidth estimator in all single-hop paths.

We examine the rate response curves for the three cross-traffic traces using packet-train parameters $s = 750$ bytes and $n = 64$ packets, similar to the parameters used by PTR [6]. We find that, for PUS and POF, these parameters are sufficient for reducing the response deviation to a negligible level and producing accurate estimation of available bandwidth from the output rate at the deviation onset point. In both cases, the measurement bias is around 0.5 mb/s (i.e., 7% of the actual available bandwidth). For the real traffic trace TXS, the measurement bias becomes 3 mb/s, which is substantially higher and might be non-negligible for certain applications. Our analysis also leads to recommendations that conflict with those stated in [6]—using larger packet size (e.g., 1500 bytes) should reduce measurement bias and not cause overestimation.

C. Spruce

Spruce uses (36) with input probing rate $r_I = C$ to estimate cross-traffic intensity. Thus, it is unbiased according to Theorem 1, regardless of the packet-train parameters s and n . However, by extending the analysis in this paper to multihop paths, it has been shown in [14] that cross-traffic interference from nontight hops can often cause significant amount of negative bias (often more than the actual available bandwidth) into Spruce's estimator. We skip the details and refer the reader to [14] for a more thorough discussion of this issue.

VII. CONCLUDING REMARKS

This paper focused on developing a theoretical understanding of single-hop bandwidth estimation in nonfluid cross-traffic conditions. Our main contributions include a queueing-theoretic framework of packet-train bandwidth estimation, a thorough investigation of the single-hop response deviation phenomenon, and an experimental methodology that computes the response curve with high accuracy from a given cross-traffic trace.

While we identified the *response deviation* phenomenon as one potential contributing source of measurement bias, there are certainly other important issues related to the performance of measurement techniques such as multihop effects, timing errors, and layer-2 effects [20]. Our future work involves extending this analysis to multihop paths and understanding the behavior of current measurement techniques in arbitrary network paths.

APPENDIX I PROOF OF LEMMA 4

Proof: The proof is by contradiction. Suppose that $W(t)/t$ does not approach 0 as $t \rightarrow \infty$. Then, there exists a $\gamma \in (0, \lambda)$ and an increasing sequence of time points $\{\tau_n, n \geq 1\}$ with $\tau_n \rightarrow \infty$ as $n \rightarrow \infty$ such that $W(\tau_n) \geq \tau_n \gamma$ for all $n \geq 1$.

We define another sequence of time points $\{\alpha_n, n \geq 1\}$, where $\alpha_n = \sup\{s : s < \tau_n, W(s) = 0\}$. Then it follows that $W(\alpha_n -) = 0$ and that for any time instant $s \in [\alpha_n, \tau_n]$, $W(s) > 0$, which means that interval $[\alpha_n, \tau_n]$ is a hop busy period. Also, it is obvious that as $n \rightarrow \infty$, $\alpha_n \rightarrow \infty$.

From Lemma 1, we can easily get the following:

$$W(\alpha_n) = W(\tau_n) + V(\alpha_n) - V(\tau_n) + C(\tau_n - \alpha_n). \quad (71)$$

Let $\epsilon = \gamma/3$, for sufficiently large n , we have

$$(\lambda - \epsilon)\alpha_n \leq V(\alpha_n) \leq (\lambda + \epsilon)\alpha_n \quad (72)$$

and

$$(\lambda - \epsilon)\tau_n \leq V(\tau_n) \leq (\lambda + \epsilon)\tau_n. \quad (73)$$

Combining (71), (72), and (73), we get

$$\begin{aligned} W(\alpha_n) &\geq \gamma\tau_n + (\lambda - \epsilon)\alpha_n - (\lambda + \epsilon)\tau_n + C(\tau_n - \alpha_n) \\ &= \gamma\tau_n + (C - \lambda)(\tau_n - \alpha_n) - \epsilon(\alpha_n + \tau_n) \\ &\geq \gamma\tau_n - \epsilon(\alpha_n + \tau_n) \geq \gamma\tau_n/3 \geq \gamma\alpha_n/3. \end{aligned} \quad (74)$$

Now define a third sequence of time points $\{b_n, n \geq 1\}$ such that $\alpha_n - b_n$ is less than a constant κ for all n and $[b_n, \alpha_n]$ is a hop idle period. Then, it follows from (74) that

$$\lim_{n \rightarrow \infty} \left(\frac{V(\alpha_n)}{\alpha_n} - \frac{V(b_n)}{b_n} \right) \geq \gamma/3. \quad (75)$$

This contradicts Assumption 1, which implies the following:

$$\lim_{n \rightarrow \infty} \frac{V(\alpha_n)}{\alpha_n} = \lim_{n \rightarrow \infty} \frac{V(b_n)}{b_n} = \lambda. \quad (76)$$

Therefore, the term $W(t)/t$ must approach 0 as $t \rightarrow \infty$. ■

APPENDIX II PROOF OF THEOREM 6

Proof: When the hop utilization process $\{U(t)\}$ is regenerative, the process $\{C(1 - U(t))\}$ is also regenerative with the same stopping times and regeneration cycles. Further note that the δ -interval available bandwidth $B_\delta(t)$ is a time-average of the regenerative process $\{C(1 - U(t))\}$. According to the regenerative central limit theorem [26, p. 124], the frequency distribution $P_\delta(x)$ converges to a Gaussian distribution $N(C - \lambda, \sigma^2/\delta)$ as δ approaches infinity, where σ is a constant. This implies that the mean of the Gaussian distribution remains $C - \lambda$ for all δ , while the variance is inversely proportional to δ . Therefore, for sufficiently large δ , we have

$$P_\delta(r) = \frac{1}{2} \left(1 + \mathbf{erf} \left(\frac{(r - C + \lambda)\sqrt{\delta}}{\sigma\sqrt{2}} \right) \right) \quad (77)$$

where \mathbf{erf} is the *Gauss error function*.

According to the asymptotic series of $\mathbf{erf}(x)$ [1, pp. 297–309], we have

$$\mathbf{erf}(x) = \begin{cases} \Theta \left(\frac{-1}{xe^{x^2}} \right) - 1, & x < 0 \\ \Theta \left(\frac{-1}{xe^{x^2}} \right) + 1, & x > 0. \end{cases} \quad (78)$$

Combining (78) with (77), we have

$$P_\delta(r) = \begin{cases} \Theta \left(\frac{1}{\sqrt{\delta}e^{k\delta}} \right), & r < C - \lambda \\ 1 - \Theta \left(\frac{1}{\sqrt{\delta}e^{k\delta}} \right), & r > C - \lambda \end{cases} \quad (79)$$

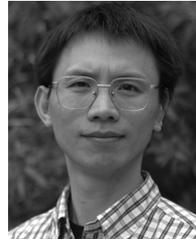
where k is a positive constant given below:

$$k = \frac{(r - \lambda)^2}{2\sigma^2}. \quad (80)$$

Subtracting (79) from (28), the theorem follows. ■

REFERENCES

- [1] M. Abramowitz and I. A. Stegun, Eds., *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, ser. Applied Mathematics Series, 55. Washington, DC: Nat. Bureau of Standards, 1972.
- [2] F. Baccelli, S. Macliraju, D. Veitch, and J. Bolot, "The role of PASTA in network measurement," in *Proc. ACM SIGCOMM*, Aug. 2006, pp. 231–242.
- [3] J. Bolot, "Characterizing end-to-end packet delay and loss in the internet," in *Proc. ACM SIGCOMM*, Aug. 1993, pp. 289–298.
- [4] C. Dovrolis, P. Ramanathan, and D. Moore, "What do packet dispersion techniques measure?," in *Proc. IEEE INFOCOM*, Apr. 2001, pp. 905–914.
- [5] G. He and J. Hou, "On exploiting long range dependence of network traffic in measuring cross traffic on an end-to-end basis," in *Proc. IEEE INFOCOM*, Mar. 2003, pp. 1858–1868.
- [6] N. Hu and P. Steenkiste, "Evaluation and characterization of available bandwidth probing techniques," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 6, pp. 879–894, Aug. 2003.
- [7] V. Jacobson, "Congestion avoidance and control," *Proc. ACM SIGCOMM*, pp. 314–329, Aug. 1988.
- [8] M. Jain and C. Dovrolis, "End-to-end available bandwidth: Measurement methodology, dynamics, and relation with TCP throughput," *Proc. ACM SIGCOMM*, pp. 295–308, Aug. 2002.
- [9] M. Jain and C. Dovrolis, "Ten fallacies and pitfalls in end-to-end available bandwidth estimation," in *Proc. ACM IMC*, Oct. 2004, pp. 272–277.
- [10] S. Keshav, "A control-theoretic approach to flow control," *Proc. ACM SIGCOMM*, pp. 3–15, Sep. 1991.
- [11] K. Lai and M. Baker, "Measuring bandwidth," in *Proc. IEEE INFOCOM*, Mar. 1999, pp. 235–245.
- [12] X. Liu, K. Ravindran, B. Liu, and D. Loguinov, "Single-hop probing asymptotics in available bandwidth estimation: Sample-path analysis," *Proc. ACM IMC*, pp. 300–313, Oct. 2004.
- [13] X. Liu, K. Ravindran, B. Liu, and D. Loguinov, "Single-hop probing asymptotics in available bandwidth estimation: Sample-path analysis." City University of New York, New York, Tech. Rep. TR-2004012, Aug. 2004. [Online]. Available: <http://www.cs.gc.cuny.edu/tr/files/TR-2004012.pdf>
- [14] X. Liu, K. Ravindran, and D. Loguinov, "Multi-hop probing asymptotics in available bandwidth estimation: Stochastic analysis," *Proc. ACM IMC*, pp. 173–186, Oct. 2005.
- [15] X. Liu, K. Ravindran, and D. Loguinov, "What signals do packet-pair dispersions carry?," in *Proc. IEEE INFOCOM*, Mar. 2005, pp. 281–292.
- [16] B. Melamed and D. Yao, "The ASTA property," in *Advances in Queueing: Theory, Methods and Open Problems*, J. H. Dshalalow, Ed. Boca Raton, FL: CRC Press, 1995, pp. 195–224.
- [17] B. Melander, M. Bjorkman, and P. Gunningberg, "A new end-to-end probing and analysis method for estimating bandwidth bottlenecks," in *Proc. IEEE GLOBECOM Global Internet Symp.*, Nov. 2000, pp. 415–420.
- [18] B. Melander, M. Bjorkman, and P. Gunningberg, "Regression-based available bandwidth measurements," *Proc. SPECTS*, Jul. 2002.
- [19] E. Muhammad and S. Stidham, *Sample-Path Analysis of Queueing Systems*. Norwell, MA: Kluwer Academic, 1999.
- [20] A. Pasztor and D. Veitch, "The packet size dependence of packet pair like methods," in *Proc. IEEE/IFIP Int. Workshop on Quality of Service (IWQoS)*, May 2002, pp. 204–213.
- [21] V. Paxson, "End-to-end Internet packet dynamics," *IEEE/ACM Trans. Netw.*, vol. 7, no. 3, pp. 277–292, Jun. 1999.
- [22] V. Ribeiro, M. Coates, R. Riedi, S. Sarvotham, B. Hendricks, and R. Baraniuk, "Multifractal cross traffic estimation," in *Proc. ITC Specialist Seminar on IP Traffic Measurement*, Monterey, CA, Sep. 2000, pp. 15-1–15-10.
- [23] V. Ribeiro, R. Riedi, R. Baraniuk, J. Navratil, and L. Cottrell, "Pathchirp; Efficient available bandwidth estimation for network paths," presented at the Passive Active Measurement Workshop, San Diego, CA, Apr. 2003.
- [24] J. Strauss, D. Katabi, and F. Kaashoek, "A measurement study of available bandwidth estimation tools," in *Proc. ACM IMC*, Oct. 2003, pp. 39–44.
- [25] R. Wolff, "Poisson arrivals see time averages," *Oper. Res.*, vol. 30, no. 2, pp. 223–231, 1982.
- [26] R. Wolff, *Stochastic Modeling and the Theory of Queues*. Englewood Cliffs, NJ: Prentice-Hall, 1989.



Xiliang Liu received the B.S. (hons.) degree in computer science from Zhejiang University, Hangzhou, China, in 1994, the M.S. degree in information science from the Institute of Automation, Chinese Academy of Sciences, Beijing, China, in 1997, and the Ph.D. degree in computer science from the City University of New York, New York, in 2005.

He currently works for Bloomberg L.P., New York. His research interests include Internet measurement and monitoring, overlay networks, bandwidth estimation, and stochastic modeling and

analysis of networked systems.

Dr. Liu has been a member of the Association for Computing Machinery (ACM) since 2002.



Kaliappa Ravindran received the Ph.D. degree in computer science from the University of British Columbia, Vancouver, BC, Canada.

He has previously held faculty positions at the Kansas State University, Manhattan, and at the Indian Institute of Science, Bangalore. He had worked in Canadian communication industries for a short period before moving to the United States. He is currently a faculty member of computer science at the City University of New York, located in the City College campus. His research interests span

the areas of service-level management of distributed networks, compositional design of network protocols, system-level support for information assurance, distributed collaborative systems, and Internet architectures. His recent industry project relationships have been with IBM, AT&T, Philips, ITT, and HP. Besides industries, some of his research has been supported by grants and contracts from federal government agencies.



Dmitri Loguinov (S'99–M'03) received the B.S. (hons.) degree in computer science from Moscow State University, Moscow, Russia, in 1995 and the Ph.D. degree in computer science from the City University of New York, New York, in 2002.

Since September 2002, he has been an Assistant Professor of computer science with Texas A&M University, College Station. His research interests include peer-to-peer networks, Internet video streaming, congestion control, image and video coding, and Internet traffic measurement and

modeling.