Available as Los Alamos Unclassified Report LA-UR 03-7267

A Systematic Approach to Probabilistic Quality of Service Guarantees in Communication Networks *

Sami Ayyorgun Wu-chun Feng

Research And Development in Advanced Network Technology (RADIANT)
Los Alamos National Laboratory
P.O. Box 1663, M.S. D451
Los Alamos, New Mexico 87545

{sami,feng}@lanl.gov

Abstract—We propose a probabilistic definition of burstiness for network traffic characterization, based on *service curves*. This definition can address network traffic with arbitrary *heavy-tailed* distributions, including the *power-tailed* distributions, in performance analyses. It facilitates a simple systematic analysis of provable, probabilistic Quality of Service guarantees. We show that the queue size, output traffic, virtual-delay, aggregate traffic, etc. at various points in a network can easily be characterized within the framework of our definition. We provide preliminary experimental/simulation results showing the applicability of our framework in real networks. We also describe two schemes which make this framework a viable solution to provide probabilistic QoS guarantees in communication networks.

Keywords—Traffic Characterization, Probabilistic QoS Guarantees, Measurement-Based Performance Analysis, Queueing, Network Calculus.

1 Introduction

Service and traffic models which facilitate a tractable performance analysis is crucial for a widespread deployment of network services requiring *provable* Quality of Service (QoS) guarantees, such as real-time multimedia applications, by network service providers. There is a prevailing agreement among network researchers and practitioners that services with provable QoS guarantees are needed, as well as in demand, for broadband networks (most recently, refer to the talk [1]).

The tractability—or, intractability—of performance analyses is greatly determined by the service and traffic models being employed. Service models provide an abstraction to incorporate QoS guarantees into both the analysis and applications. Traffic models, on the other hand, are needed to model the behavior of network traffic in real life, as well as to provide an understanding to alter their behavior as desired for a better network management and performance.

On the service part, a major challenge is to come up with definitions of service models to quantify the concept of "service" in analytically tractable ways. In the mid 90's, a such service model, called the *service curve model*, has been proposed [2, 4], which exhibits analytical tractability similar to that in linear system theory. For various applications of the service curve model, such as multimedia smoothing, refer to [5].

On the traffic part, a major challenge is to come up with traffic models capturing the bursty nature of network traffic, and at the same time facilitating a tractable performance analysis. With the rapid spread of data networks within the last decade, it has become apparent that network traffic exhibits bursty behavior. However, a widely accepted definition of burstiness, be it either deterministic or probabilistic, does not exist in the networking community. Yet, it is known that the burstiness of network traffic has a profound impact on the performance of many network protocols in areas such as congestion control (e.g., TCP), multiple-access (e.g., CSMA), routing (e.g., BGP), and switching and multiplexing in general. Deterministic definitions of burstiness, and hence traffic characterizations, inherently lack statistical gains, leading to an underutilization of network resources; e.g., [3, 4, 6, 7]. The probabilistic characterizations, on the other hand, are either too complicated for tractable performance analyses (e.g., refer to the related work in [8] about *self-similar* traffic models), or lack capability to effectively deal with the network traffic with recently observed behavior (e.g., *long-range dependency*). Almost all the traditional traffic models and/or characterizations (e.g., Markovian models) fall into the later category, as well as some of the new models (e.g., [9, 10, 11]).

We believe that service and traffic models are to be considered together and be defined in accordance, to facilitate a tractable performance analysis. This study attempts to do that by proposing a probabilistic definition of burstiness based on service curves, which facilitate a systematic treatment of probabilistic performance guarantees. Our work is strongly motivated by the work in [9, 10, 11].

In coming up with this definition which we propose here, we have had in mind the properties that a "good" traffic characterization should have in order to facilitate a tractable performance analysis. Some of these properties are listed below: Let $\mathcal{C}$ denote a "good" traffic characterization, then

*Los Alamos National Laboratory, Technical Report LA-UR-03-7267, September 2003.

1. if two traffic flows are characterized according to $\mathcal{C}$, then the aggregate of the flows should also be easily characterized according to $\mathcal{C}$,
2. if a traffic flow characterized according to $\mathcal{C}$ is fed into a network element commonly used in practice, then
   (a) the output flow should also be easily characterized according to $\mathcal{C}$, and
   (b) both the queue-size and the virtual delay should be easily characterized in the same framework as $\mathcal{C}$, as well
3. the characterization $\mathcal{C}$ should be "stationary" in the sense that the characterization of a time-shifted traffic does not change with respect to that of the unshifted traffic.
4. the characterization of a flow according to $\mathcal{C}$ should be measurable, and be easily measured,
5. the characterization $\mathcal{C}$ should capture the behavior of traffic in real networks as closely as possible, without sacrificing neither the tractability nor the reality.

We show that the traffic characterization which we propose here satisfies all of the above properties. The first three properties facilitate a systematic treatment of performance analyses. Together with the last two properties, they make the framework that we examine here a viable solution to provide probabilistic QoS guarantees in communication networks. The characterization that we propose here is capable of addressing the network traffic with arbitrary *heavy-tailed* distributions, including the *power-tailed* distributions, in performance analyses. We provide a few experimental/simulation results to this end.

The rest of the paper is organized as follows: Section 2 gives a background. Section 3 introduces the new burstiness definition. Sections 3.1 and 3.2 provides the implications of the new definition, by itself and over a single network element, respectively. Section 3.3 gives the performance guarantees over a tandem of network elements, according to the new burstiness definition. Section 3.4 provides the mean performance guarantees. Section 3.5 shows the stationarity of the new traffic characterization. Section 3.6 explains how the traffic characterization that we propose here is measurable. Section 3.7 gives experimental/simulation results showing the applicability of the framework that we propose. Section 4 describes two schemes which make the framework that we propose here viable in communication networks. Finally, Section 5 provides conclusions.

2 Background and Convention

We adopt a discrete-time formulation for simplicity. Time is slotted into fixed-length intervals, and marked by the integers. The unit of transmission for communication is referred to as a *packet*, in this study. A *flow* is a discrete random process such that each sample-path is a non-decreasing function defined from the integers to the non-negative integers. The value $R(n)$ of a flow R at time n denotes the total number of packets that has arrived by time n (inclusive) for a *connection*. A *network element* is an input-output device which accepts packets at its input, and delivers them at its output. A network element is said to be *passive* if it does not generate any packet internally. All network elements are assumed to be passive in this study, for simplicity. Packets are assumed to be able to instantaneously arrive and depart at a network element, i.e., a whole packet could arrive instantaneously at time k, and depart later at time n where $n \geqslant k$. Note that a packet could depart in the same interval in which it has arrived; this is sometimes referred to as *cut-through* operation. The capacity of a network element at time n is the total number of packets that it could serve at that time. We denote the set of all the integers by $\mathbb{Z}$. Finally, all functions are assumed to be defined from $\mathbb{Z}$ to $\mathbb{Z}$, unless otherwise noted from here on.

We utilize the following definitions which are previously introduced in the literature (e.g., see [5, 14]).

Definition 1 *Let f and g be any two functions. The* min-+ convolution $f \triangledown g$ *of f and g is defined as*[1]

$$(f \triangledown g)(n) = \min_{k \leqslant n}\{f(k) + g(n-k)\} \qquad \text{for all } n.$$

The convolution[2] $f \triangledown g$ is read as "f min-convolved with g", or as "the min-plus convolution of f with g".

Definition 2 *Let f and g be any two functions. The* min-+ deconvolution $f \,\bar{\triangledown}\, g$ *of f with g is defined as*

$$(f \,\bar{\triangledown}\, g)(n) = \max_{k \geqslant 0}\{f(n+k) - g(k)\} \qquad \text{for all } n.$$

The deconvolution $f \,\bar{\triangledown}\, g$ is read as "f min-deconvolved with g", or as "the min-plus deconvolution of f with g"

Definition 3 *Let f and g be any two functions. The* max-+ convolution $f \vartriangle g$ *of f with g is defined as*

$$(f \vartriangle g)(n) = \max_{k \leqslant n}\{f(k) + g(n-k)\} \qquad \text{for all } n.$$

The convolution $f \vartriangle g$ is read as "f max-convolved with g", or as "the max-plus convolution of f with g".

Definition 4 *Let f and g be any two functions. The* max-+ deconvolution $f \,\bar{\vartriangle}\, g$ *of f with g is defined as*

$$(f \,\bar{\vartriangle}\, g)(n) = \min_{k \geqslant 0}\{f(n+k) - g(k)\} \qquad \text{for all } n.$$

[1] There are two other variants of min-+ convolution, depending on how the subscript of the minimum is delimited in the definition; one for '$0 \leqslant k \leqslant n$', and one for '$k \in \mathbb{Z}$'.

[2] One reason to choose this notation for min-+ convolution over some others, for example '$*$', is that we believe this choice of notation provides a better selection of notations for other related operators introduced in definitions 2, 3, and 4, in a fitting manner.

The deconvolution $f \bar{\Delta} g$ is read as "f max-deconvolved with g", or as "the max-plus deconvolution of f with g".

Definition 5 *An S-server with* service curve S *is a network element that when fed with an input flow R, the corresponding output flow G satisfies*

$$G(n) \geqslant (R \,\bar{\nabla} S)(n) \quad \textit{for all } n$$

for any R. A service curve S is a non-decreasing function defined from the integers to the non-negative integers, where $S(n) = 0$ for all $n \leqslant 0$.

An *S-server with equality* is an S-server that the inequality in definition 5 becomes an equality, i.e., the output flow of an S-server with equality is given by $G(n) = (R \,\bar{\nabla} S)(n)$ for all n.

Note that a *work-conserving*[3] server with the capacity of serving packets with a constant integer rate ρ is an S-server with equality with service curve $S(n) = \max\{0, \rho \cdot n\}$, since the output flow G of a work-conserving server with an input flow R is given[4] by

$$G(n) = \min_{k \leqslant n}\{R(k) + \rho\,(n-k)\} \quad \text{for all } n,$$

which is equivalently represented as

$$G(n) = (R \,\bar{\nabla} S)(n) \quad \text{for all } n.$$

An S-server could be viewed as a generalization of a work-conserving server with constant rate.

The concept of service curves, together with an appropriate traffic characterization, facilitate a systematic treatment of performance guarantees in communication networks. This has been demonstrated for some deterministic regimes, e.g., [3, 4, 5, 14, 16]. In this study, we carry this much desired feature into a probabilistic setting.

3 A Probabilistic Definition of Burstiness Characterization

We propose the following probabilistic definition of burstiness, and study some of its implications.

Definition 6 *A traffic flow R is said to be bursty with service curve S and bounding function f, and denoted as $R \sim (S, f)$, if*

$$\mathrm{P}\Big(R(n) - R(k) > S(n-k) + \sigma, \text{ for some } k < n\Big) \leqslant f(\sigma)$$

for all σ and for all n, where f is defined from the integers to the non-negative real numbers.

[3]A network element is said to be *work-conserving* if it serves packets whenever it has packets to serve, unconditionally of any other criteria.

[4]This is often referred to as Reich's result [12]. See for example [13] for a simple derivation.

We assume that the following properties hold for any bounding function f, without loss of generality:

1. f is non-increasing, since the probability corresponding to a σ in the above definition is non-increasing with σ.

2. $f(\sigma) \leqslant 1$ for all σ, since the probability of an event can not be larger than 1. We assume for mathematical convenience that $f(\sigma) = 1$ for all $\sigma < 0$, unless otherwise noted from here on.

3. $\lim\limits_{\sigma\to\infty} f(\sigma) = 0$, since any cumulative distribution function F satisfies $\lim\limits_{x\to\infty} F(x) = 1$.

In the rest of this section, we examine some of the properties/implications of definition 6. Specifically, we will show that it satisfies all the properties of a "good" traffic characterization that we sought to have as stated in the introduction.

3.1 Implications on Aggregate Flow and Average Rate

We first show that the traffic characterization provided by definition 6 satisfies the property 1 of a "good" traffic characterization. This is given by the following theorem.

Theorem 1 *Given any two flows $R_1 \sim (S_1, f_1)$ and $R_2 \sim (S_2, f_2)$, the aggregate flow $R_1 + R_2$ is bursty with service curve $S_1 + S_2$ and bounding function $f_1 \,\bar{\nabla} f_2$. In other words, $R_1 + R_2 \sim (S_1 + S_2, f_1 \,\bar{\nabla} f_2)$.*

Proof: (See [15].)

By a repeated application of theorem 1, we could easily find the characterization of an aggregate flow with an arbitrary number of aggregation.

The burstiness characterization provided by definition 6 has also an implication on the long-term average rate of a flow. This is given by the following theorem.

Theorem 2 *Given a flow $R \sim (S, f)$, the long-term average rate μ of R satisfies*

$$\begin{aligned}\mu &= \limsup_{(n-k)\to\infty} \frac{\mathrm{E}[R(n) - R(k)]}{n-k} \\ &\leqslant \limsup_{n\to\infty} \frac{S(n) + \sum_0^n f(\sigma)}{n}.\end{aligned}$$

Proof: (See [15].)

Note that if the area under the bounding function f up to n (i.e., $\sum_0^n f(\sigma)$) is $O(n^{1-\epsilon})$ for some $\epsilon > 0$, then the long-term average rate μ of R is upper-bounded by $\limsup_{n\to\infty} \frac{S(n)}{n}$.

In the following section, we study the implications of definition 6 over an S-server, whereby we show some of the other properties of this characterization pursuant to the properties of a traffic characterization that we sought to have as stated in the introduction.

3.2 Implications over an S-server

If a flow $R \sim (S^*, f)$ is fed into an S-server, the distribution of the queue size Q at the server is also similarly upper-bounded by f. This is given more precisely by the following theorem. The queue-size $Q(n)$ is the total number of packets which resides in the server at time n; that is, if R and G denote the aggregates of the flows at the input and at the output of the server, respectively, then $Q(n) = R(n) - G(n)$.

Theorem 3 *If a flow $R \sim (S^*, f)$ is fed into an S-server with service curve S, then the distribution of the queue size Q at the server satisfies*

$$\mathrm{P}(Q(n) > \sigma) \leqslant f\big((S \mathbin{\overline{\Delta}} S^*)(0) + \sigma\big)$$

for all σ and n.

Proof: (See [15].)

Note that if service curve S of the S-server is greater than or equal to S^* of the flow at every point, then the bound on the queue-size distribution becomes equal to $f(\sigma)$.

The output flow of an S-server fed by a flow $R \sim (S^*, f)$ could also be easily characterized according to definition 6. This is given by the following theorem.

Theorem 4 *The output flow G of an S-server with service curve S, fed by a flow $R \sim (S^*, f)$, is bursty with service curve $S^* \mathbin{\underline{\nabla}} S$ and the bounding function f. In other words, $G \sim (S^* \mathbin{\underline{\nabla}} S, f)$.*

Proof: (See [15].)

We might actually need to slightly rectify the above result by replacing the service curve $S^* \mathbin{\underline{\nabla}} S$ in characterizing the output flow by S_o which is given below

$$S_o(n) = \begin{cases} 0 & \text{if } n \leqslant 0 \\ (S^* \mathbin{\underline{\nabla}} S)(n) & \text{else.} \end{cases}$$

We would like to have this rectification for two reasons: (1) A service curve is defined to take on the value zero for non-negative values of its argument. (2) We would only need to have $S^* \mathbin{\underline{\nabla}} S$ for positive values of its argument (this could be noted clearly via the proof of the above result). We have not done this rectification in the body of the theorem in order not to clutter the result.

The burstiness characterization provided by definition 6 has also an implication on the *virtual-delay* at an S-server. The definition of *virtual-delay* is given below, which is previously introduced in the literature.

Definition 7 *The* virtual-delay *$D(n)$ at any time n for an input flow R at a network element is defined as*

$$D(n) = \min\{\delta : \delta \geqslant 0,\, G(n+\delta) \geqslant R(n)\}$$

where G is the corresponding output flow.

The virtual-delay $D(n)$ is basically the delay experienced by the packets arriving at time n, through the network element, if the packets are to be served in the order in which they have arrived.

Theorem 5 *If an input flow $R \sim (S^*, f)$ is fed into an S-server with service curve S, then the distribution of the virtual-delay $D(n)$ at the server satisfies*

$$\mathrm{P}(D(n) > \sigma) \leqslant f\big((S \mathbin{\overline{\Delta}} S^*)(\sigma)\big)$$

for all $\sigma \geqslant 0$, and for all n.

Proof: (See [15].)

We could actually slightly improve the results in theorems 3 and 5. See [15] for details.

All of the theorems in this section facilitate a systematic treatment of performance guarantees over a tandem of network elements, analytically. To further clarify this point, we provide the results in the following section.

3.3 Performance Guarantees Over A Tandem of Network Elements

Performance guarantees over a tandem of network elements, in the framework of definition 6, follow directly from the results in section 3.2. This is given by the following corollary.

Corollary 1 *Let a flow $R_1 \sim (S^*, f)$ be fed into an S-server with service curve S_1, and let the output R_2 of this first server be fed into another S-server with service curve S_2. The following statements hold:*

1. *The output flow R_3 of the S-server with service curve S_2 is bursty with service curve S_3 and bounding function f, where*

   $$S_3(n) = \begin{cases} 0 & \textit{if } n \leqslant 0 \\ (S^* \mathbin{\underline{\nabla}} (S_1 \nabla S_2)) & \textit{else.} \end{cases}$$

   In other words, $R_3 \sim (S_3, f)$.

2. *The total number of packets, $Q_1(n) + Q_2(n)$, stored in the tandem network at any time n satisfies*

   $$\mathrm{P}((Q_1 + Q_2)(n) > \sigma) \leqslant (g \nabla h)(\sigma)$$

   for all σ, where

   $$g(\sigma) = f\big((S_1 \mathbin{\overline{\Delta}} S^*)(0) + \sigma\big)$$

   $$\begin{aligned} h(\sigma) &= f\big((S_2 \mathbin{\overline{\Delta}} (S^* \mathbin{\underline{\nabla}} S_1))(0) + \sigma\big) \\ &= f\big(((S_1 \nabla S_2) \mathbin{\overline{\Delta}} S^*)(0) + \sigma\big)\,. \end{aligned}$$

3. *The total virtual-delay, $D_1(n) + D_2(n+\delta)$, experienced by a packet arriving at any time n at the first*

network element, and at time $n+\delta$ for some $\delta \geqslant 0$ at the second network element, satisfies for any $\delta \geqslant 0$

$$\mathrm{P}(D_1(n) + D_2(n+\delta)(n) > \sigma) \leqslant (g \mathbin{\triangledown} h)(\sigma)$$

for all σ, where

$$g(\sigma) = f\big((S_1 \mathbin{\bar{\Delta}} S^*)(\sigma)\big)$$

$$\begin{aligned} h(\sigma) &= f\big((S_2 \mathbin{\bar{\Delta}} (S^* \mathbin{\underline{\nabla}} S_1))(\sigma)\big) \\ &= f\big(((S_1 \mathbin{\triangledown} S_2) \mathbin{\bar{\Delta}} S^*)(\sigma)\big) . \end{aligned}$$

Proof: (See [15].)

We could actually slightly improve the above results in items 2 and 3 in corollary 1, and further emphasize the message that we are trying to convey by the corollary. It is not difficult to show that the bounds that we would obtain by replacing $(g \mathbin{\triangledown} h)(\sigma)$ in items 2 and 3 in corollary 1 by $h(\sigma)$, also hold; see [15] for a proof. Hence, the bounds in the last two items in corollary 1 could actually be replaced by

$$\min\{h(\sigma), (g \mathbin{\triangledown} h)(\sigma)\}$$

for all σ. Thus, it becomes clear by corollary 1 and the above inequalities being pointed out that we could actually view the tandem network as a single S-server with service curve $S_1 \mathbin{\triangledown} S_2$, and obtain valid characterizations.

By a repeated application of corollary 1 and the discussions in the previous paragraph, we could obtain similar results for any number of network elements in tandem. Thus, it has now become clear that the burstiness characterization 6 does indeed facilitate a systematic treatment of performance guarantees over a tandem of network elements, analytically.

3.4 Average Performance Guarantees

It immediately follows by the results in sections 3.2 and 3.3 that we could also give some average performance guarantees at various points in a network by the burstiness characterization that we propose here (i.e., definition 6). In this section, this is illustrated specifically for average queue-size and virtual-delay at an S-server.

Corollary 2 *If an input flow $R \sim (S^*, f)$ is fed into an S-server with service curve S, then the average queue size $Q(n)$ at the server at any time n is upper-bounded as*

$$\mathrm{E}[Q(n)] \leqslant \sum_{\sigma=0}^{\infty} f\big((S \mathbin{\bar{\Delta}} S^*)(0) + \sigma\big) .$$

Proof: The proof follows immediately by theorem 3;

$$\begin{aligned} \mathrm{E}[Q(n)] &= \sum_{\sigma=0}^{\infty} \mathrm{P}(Q(n) > \sigma) \\ &\leqslant \sum_{\sigma=0}^{\infty} f\big((S \mathbin{\bar{\Delta}} S^*)(0) + \sigma\big) . \end{aligned}$$ ∎

Note that given a flow $R \sim (S^*, f)$ where f is the tightest possible bounding function, if the area under the bounding function $f\big((S \mathbin{\bar{\Delta}} S^*)(0) + \sigma\big)$ is infinite, then when flow R is fed into an S-server with equality with service curve S (note the definition given right after the definition of an S-server), the mean queue size at any time n also becomes infinite.

Similarly, mean virtual-delay is also upper-bounded. This is given below by the following corollary.

Corollary 3 *If an input flow $R \sim (S^*, f)$ is fed into an S-server with service curve S, then the mean virtual-delay $D(n)$ at the server at any time n is upper-bounded as*

$$\mathrm{E}[D(n)] \leqslant \sum_{\sigma=0}^{\infty} f\big((S \mathbin{\bar{\Delta}} S^*)(\sigma)\big) .$$

Proof: The proof follows immediately by theorem 5;

$$\begin{aligned} \mathrm{E}[D(n)] &= \sum_{\sigma=0}^{\infty} \mathrm{P}(D(n) > \sigma) \\ &\leqslant \sum_{\sigma=0}^{\infty} f\big((S \mathbin{\bar{\Delta}} S^*)(\sigma)\big) . \end{aligned}$$ ∎

3.5 Stationarity

The traffic characterization that we propose here (i.e., definition 6) has also the "stationarity" property (i.e., property 3) of a "good" traffic characterization that we have stated in the introduction.

This can be shown almost without any effort by the following identity:

$$\begin{aligned} &\mathrm{P}\Big(R(n) - R(k) > S(n-k) + \sigma, \quad \text{for some } k < n\Big) = \\ &\mathrm{P}\Big(R(n+t) - R(k+t) > S(n-k) + \sigma, \\ &\qquad\qquad\qquad \text{for some } k+t < n+t\Big) \end{aligned}$$

which holds for all σ, n, and t. Note that $R(n+t)$ is a time shifted version of $R(n)$ by t amount.

Hence, if R is bursty with service curve S and bounding function f, then so is its time-shifted version R' where $R'(n) = R(n+t)$ for any given t.

3.6 Measurability

The traffic characterization provided by definition 6 is also measurable, and could easily be measured. In other words, it has the property 4 of a "good" traffic characterization that we have stated in the introduction, as well. This is facilitated by theorems 2 and 3.

Given a flow R and a service curve S, we could easily measure an estimate of a tight bounding-function f in characterizing R as $R \sim (S, f)$, by utilizing theorem 3. This could be done as follows: First, we would need to have an S-server with equality with service curve S (note the definition given right after definition 5 in Section 2).

Then, upon feeding flow R into a such S-server, we observe the queue-size Q at the server. The complementary cumulative distribution function (i.e., $1-F(x)$ where $F(x)$ is the CDF) of the queue size Q could be estimated by various statistical methods from the empirical values. Such an estimate $\hat{f}$ provides an estimate for a tight bounding-function f in characterizing R as $R \sim (S, f)$. This could be seen more clearly via the proof of the theorem 3; see [15].

Secondly, given a flow R with unknown characterization according to definition 6, various service curves S could be suggested in light of theorem 2 in order to characterize R as $R \sim (S, f)$.

For this, we would first need to have a good estimate of the long-term average rate μ of R. Getting a fairly good approximation of the long-term average rate of a flow is often not difficult by applying a variant of the Law of Large Numbers.

Then by theorem 2, we could pick a service curve S such that $\limsup_{n\to\infty} \frac{S(n)}{n} \geqslant \mu$. Thus, service curves of the from

$$S(n) = \max\{0, \lceil \mu \rceil \, (n - D)\}$$

are good candidates to find a characterization of R as $R \sim (S, f)$. A delay parameter D could be chosen in light of the discussions given in Section 3.4; specifically, by considering (1) the end-to-end delay requirement of flow R, (2) how this delay requirement is to be distributed over each network element on the path of flow R before it reaches its destination.

Finally, by the discussions given in the second paragraph of this section, we could find a tight bounding function f to complete the characterization of R as $R \sim (S, f)$.

Moreover, we would like to note that the above procedure to find an appropriate S could be applied recursively, depending on the decay rate of the estimate of the bounding function f, in light of theorem 2, to find a more fine grained characterization of flow R.

3.7 Capturing The Behavior of Real Network Traffic

To see how applicable the framework that we propose here is in real networks, we have performed some preliminary experiments and simulations. In this section, we present a few results from this experimental/simulation study. A comprehensive report about the experimental/simulation work will be made available in a future study.

The simulation setup that we consider here is shown in Figure 1. The input flow R_{in} is obtained from a real traffic trace. The data was collected on the incoming/outgoing link of the Los Alamos National Laboratory, which connected the Laboratory to the outside world. Hence, it represents a Wide Area Network traffic.

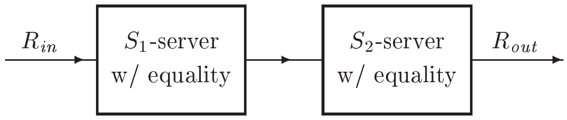

Figure 1: The simulation setup.

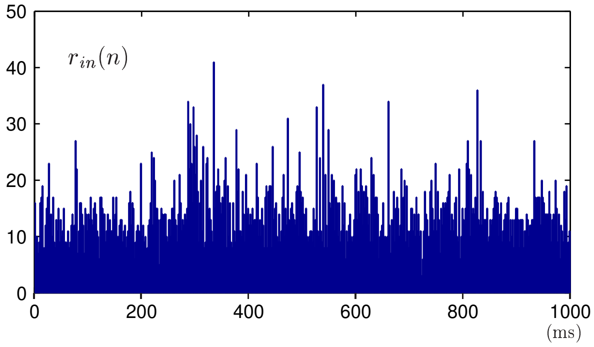

Figure 2: The instantaneous rate r_{in} in the first second.

The underlying technology was FDDI[5].

The flow R_{in} is obtained by first slotting the time into 1 millisecond intervals, and then counting the number of packet arrivals in each interval. The value $r_{in}(n) \triangleq R_{in}(n) - R_{in}(n-1)$ is the total number of packets arrive in interval n. The trace length that we ran in simulations was 1600 seconds (more than 26 minutes). The instantaneous rate r_{in} in the first second is shown in Figure 2.

The estimate of the average rate μ of R_{in} was 11.9974 packets per time slot. Hence, we picked the following service curve

$$S^*(n) = \max\{0, 13\,(n-3)\}$$

in characterizing R_{in}. By feeding R_{in} into an S-server with equality with service curve S^*, we found an estimate of the bounding function f_{in}, as explained in Section 3.6, which is shown in Figure 3.

We feed R_{in} into the tandem of network elements in Figure 1, where

$$S_1(n) = \max\{0, 16\,(n-2)\}$$
$$S_2(n) = \max\{0, 14\,(n-1)\}\,,$$

and obtained the output flow R_{out}.

Corollary 1 shows that $R_{out} \sim (S_3, f_{in})$ where $S_3(n) = (S^* \veebar (S_1 \triangledown S_2))(n)$. In order to verify this result, we have feed R_{out} into another S-server with equality with service curve S_3, and obtained an estimate of its bounding function f_{out}, again as explained in Section 3.6. That estimate of f_{out} is also shown in Figure 3. We would expect by corollary 1 that $f_{out}(n) = f_{in}(n)$.

To show the difference between the two bounding functions more clearly, we have also plotted them in log-log scale, which is shown in Figure 4.

[5]Fiber Distributed Data Interface.

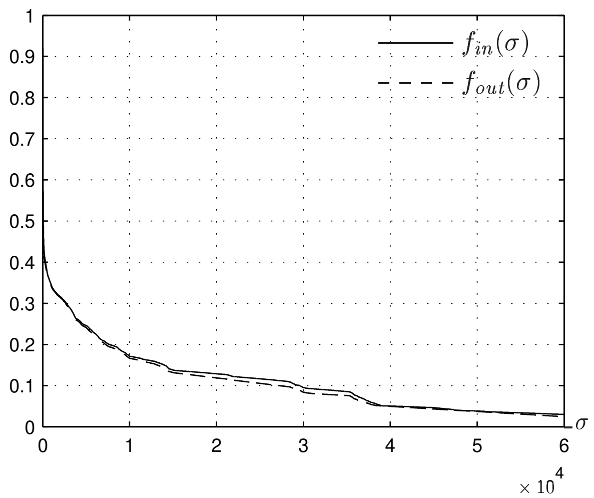

Figure 3: The bounding functions f_{in} and f_{out}.

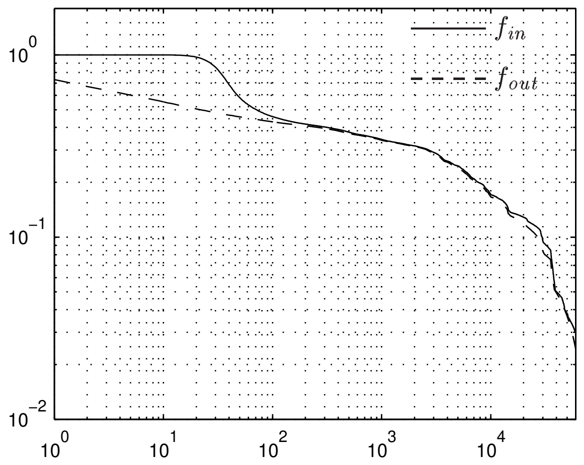

Figure 4: The bounding functions f_{in} and f_{out} in log-log scale.

The proximity of the two bounding functions to each other in these figures is striking. We think that this indicates a strong evidence that the probabilistic framework that we propose in this study is very well applicable to real networks in both capturing the traffic behavior in real network and providing probabilistic QoS guarantees efficiently. An extensive experimental and simulation work is in progress, which might provide further support to our claim. This work will be maid available in a future study.

Finally, we would like to note that the queue-size distributions that we have observed in this experimental study are heavy-tailed, as it could be seen from Figure 4. This suggests that the collected traffic trace is also heavy-tailed.

4 Two Schemes to Support The Probabilistic QoS Guarantees: A Utility of Definition 6

The probabilistic QoS guarantees described in this study are viable in communication networks, in two different schemes. In this section, we explain these schemes briefly.

4.1 An A Priori Scheme

Notice that by definition 6 we are actually characterizing a traffic source, rather than a particular realization (i.e., a sample-path) of a flow. Recall that R is a random process. Hence, we could find a characterization of a traffic source—be it either audio, video, or data—which injects traffic into the network, a priori. This could easily be done for each traffic source by following the remarks given in Section 3.6 which utilizes a comprehensive statistical analysis. Once the characterization of a traffic is known and declared at its entry point into a network, the probabilistic QoS guarantees could then be followed easily via the results given in the previous sections. (We would also need to have similar results for at least *switches* and *multiplexers*. Such results will be made available in a future study.)

A similar scheme as described above could also be considered for the deterministic framework initiated by [3] (refer to [3], [4], [14], or [5] for details). In this case, the *arrival curve*[6] characterization for each traffic source is performed a priori, and declared at the entry points of a network. However, due to the worst-case deterministic nature of this characterization, this scheme leads to overly conservative provisioning of QoS guarantees,

[6]A flow R is said to conform to an *arrival curve* A if the following condition is satisfied;

$$R(n+k) - R(k) \leqslant A(n) \qquad \text{for all } n \geqslant 0, \text{ and for all } k,$$

where A is a non-decreasing function defined from the integers to the non-negative integers that $A(n) = 0$ for all $n \leqslant 0$. See [3], [4], [5], or [14] for details.

and hence very poor utilization of the network resources. With the probabilistic framework that we propose here, this fatal drawback is eliminated; hence the scheme that we suggest here is viable.

4.2 A Probabilistic Leaky-bucket Scheme

Another scheme which makes the framework that we propose here viable is analogous to the *leaky-bucket*[7] scheme in the deterministic framework initiated by [3]. For each traffic in a leaky-bucket scheme, there is a token buffer with a size σ, to which tokens arrive with a rate ρ. A such leaky-bucket attached to a buffer to which a traffic arrives is used to regulate the traffic (at least at its entry point into a network) to make it comply with the (σ, ρ) *model*[8]. Once the characterization of a traffic is known at its entry point into the network, the deterministic QoS guarantees could be followed according to the deterministic framework (see [3], [4], [5], or [14] for details). In this scheme, depending on the size of the buffer to which the traffic arrives, packet losses could be incurred.

In this traditional leaky-bucket scheme, essentially, a synthetic traffic (i.e., the token traffic leaving the bucket) is generated, which conforms to some given (σ, ρ) parameters. Later, the real traffic is "multiplied" with this synthetic traffic in the sense that the amount of packets which leaves the buffer B_R to which the real traffic arrives is equal to the minimum of the number of the packets at the head of B_R and the number of tokens in the token buffer.

Recently, we have discovered how to generate a synthetic traffic $\hat{R}$ which conforms to given (S, f) parameters, probabilistically. The details of how to generate a such synthetic traffic will be made available in a future study.

This discovery facilitates another scheme which makes the framework that we propose here viable. We could regulate a flow R with unknown characterization according to definition 6 at its entry point into a network as follows: First, note that it is not difficult to choose an appropriate service curve S in characterizing R according to definition 6, in light of the discussions given in Section 3.6. Similarly, a suitable bounding function f could be suggested based on the service requirements of flow R, which we could express in our framework, specifically by considering the results exemplified in Section 3.4. Once we pick an S and an f, we could generate a synthetic traffic $\hat{R}$ which conforms to the (S, f) parameters. This synthetic traffic fills a token buffer. More specifically, it deposits $\hat{r}(n) \triangleq \hat{R}(n) - \hat{R}(n-1)$ many tokens into the token buffer, at any time n. The flow R in the mean time fills a buffer B_R. Let $b_R(n)$ denote the number of the packets at the head of the buffer B_R at time n. Exactly $\min\{b_R(n), \hat{r}(n)\}$ many packets leave the buffer B_R at any time n. At the end of each time n, the token buffer is flushed. It is not difficult to show that the departures from the buffer B_R conforms to (S, f).

In this way, since we would know *a* characterization of any traffic entering into the network, we could follow the probabilistic QoS guarantees according to our framework that we propose in this study.

[7] Refer to [2], [3], [4], [5], or [14] for explanations.

[8] In (σ, ρ) *model*, a traffic is said to conform to (σ, ρ) parameters, if it has an arrival curve A such that $A(n) = \rho \cdot n + \sigma$ for all $n \geqslant 1$.

5 Conclusions

In this study, we have proposed a probabilistic definition of burstiness for network traffic characterization, based on service curves. We have shown that the traffic characterization that we propose here satisfies all the properties of a "good" characterization given in the introduction.

Specifically, we have shown that the new characterization facilitates a systematic treatment of probabilistic QoS guarantees, analytically. The characterization of a flow according to this definition is "stationary". It is also measurable by constructing an S-server with equality, and observing the queue size distribution upon feeding a flow into that server. Moreover, it is a realistic characterization in the sense that it can address network traffic with arbitrary *heavy-tailed* distributions, including the *power-tailed* distributions (e.g., Pareto), which have been recently observed in networking research.

We have provided preliminary experimental/simulation results which support that the framework that we propose here is very well applicable in real networks.

We have also described two schemes which make this framework a viable solution to provide probabilistic QoS guarantees in communication networks.

The performance bounds provided by this characterization are also tighter than what is currently available by other similar approaches, e.g., [9, 10, 11]); see [15] for details. The proposed characterization here does not also have the restriction on a bounding function (i.e., f) as specified in [9]. The new characterization is also defined more generally by basing it on service curves.

Future work includes (1) examining the QoS guarantees over switches and multiplexers, in this framework, (2) conducting extensive experimental and simulation study to see how applicable the framework that we propose here is in real networks, (3) synthesizing traffic for any given (S, f) parameters.

It would be also interesting to examine a similar traffic characterization, where the new characterization would also have a lower-bound, as well as an upper-bound, on the probability given in definition 6.

More or less a deterministic version of this study could also be followed in [16].

Acknowledgement

The authors thank Sarut Vanichpun for his help in getting the figures in Section 3.7.

References

[1] A. K. Parekh. *Breaking the Great Internet Deadlock OR Why there isn't any Network QoS and what to do about it?*, The 11th International Workshop on Quality of Service (IWQoS 2003), Invited Talk. `http://iwqos03.cs.berkeley.edu/invited-talks/parekh.pdf`

[2] A. K. Parekh, R. G. Gallagher. *A generalized processor sharing approach to flow control in integrated services networks: the single-node case*, IEEE/ACM Transaction on Networking, vol. 1, pp. 344–357, 1993.

[3] R. L. Cruz. *A Calculus For Network Delay, Part I: Network Elements In Isolation*, IEEE Transactions on Information Theory, vol. 37, no. 1, pp. 114–131, Jan. 1991.

[4] R. L. Cruz. *Quality of Service Guarantees in Virtual Circuit Switched Networks*, IEEE Journal of Selected Areas in Communication, 13(6): 1048–1056, 1995.

[5] J.-Y. Le Boudec, P. Thiran. *Network Calculus: A Theory of Deterministic Queuing Systems for the Internet*, Lecture Notes in Computer Science-2050, Springer Verlag, January 2002.

[6] D. Ferrari, D. Verma. *A Scheme for Real-Time Channel Establishment in Wide-Area Networks*, IEEE Journal on Selected Areas in Communications, vol. 8, pp. 368–379, April 1990.

[7] Z. Wang, J. Crowcroft. *Analysis of Burstiness and Jitter in Real-Time Communications*, The Proceedings of SIGCOMM, pp. 13–19, 1993.

[8] K. Park, W. Willinger (editors). *Self-Similar Network Traffic and Performance Evaluation*, John Wiley & Sons, 2000.

[9] D. Starobinski, M. Sidi. *Stochastically Bounded Burstiness for Communication Networks*, IEEE Trans. on Information Theory, vol. 46, no. 1, Jan. 2000. (Also appeared in the Proceedings of Infocom 1999, pp. 36–42.)

[10] O. Yaron, M. Sidi. *Performance and stability of communication networks via robust exponential bounds*, IEEE/ACM Transactions on Networking, 1(3): 372–385, 1993.

[11] C.S.-Chang. *Stability, queue length and delay of deterministic and stochastic queuing networks*, IEEE Trans. on Automatic Control, vol. 39, pp. 913–931, 1994.

[12] E. Reich. *On the Integrodifferential Equation of Takasc, I*, Ann. Math. Stat., vol. 29, pp. 563–570, 1958.

[13] S. Ayyorgun, W.-C. Feng. *Notes on Burstiness and Bursty Traffic Characterization*, Los Alamos National Laboratory, Technical Report.

[14] C.-S. Chang. *Performance Guarantees in Communication Networks*, Springer Verlag, April 2000.

[15] S. Ayyorgun, W.-C. Feng. *A Probabilistic Definition of Burstiness Characterization: A Systematic Approach*, Los Alamos National Laboratory, Technical Report LA-UR-03-3668, May 2003.

[16] S. Ayyorgun, W.-C. Feng. *A Deterministic Definition of Burstiness For Network Traffic Characterization*, Los Alamos National Laboratory, Technical Report LA-UR-03-4477, June 2003.