A flow-based model for Internet backbone traffic

Chadi Barakat, Patrick Thiran
ICA - DSC - EPFL
{Chadi.Barakat,Patrick.Thiran}@epfl.ch

Gianluca Iannaccone, Christophe Diot
Sprint Labs
{gianluca,cdiot}@sprintlabs.com

Philippe Owezarski
LAAS-CNRS
owe@laas.fr

Technical Report EFPL/DSC/2001/044

Abstract— **We model traffic on an uncongested backbone link of an IP network using Poisson Shot-noise process and M/G/∞ queue. We validate the model by simulation. We analyze the model accuracy with real traffic traces collected on the Sprint IP backbone network. We show that despite its simplicity, our model provides a good approximation of the real traffic observed on OC-3 links. This model is also very easy to use and requires few simple parameters to be input.**

I. Introduction

Modeling the Internet traffic is an important issue. It is unlikely that we will be able to understand the traffic characteristics, to predict network performance (e.g., for QoS guarantees or service level agreement definition) or to design dimensioning tools without analytical models. The successful evolution of the Internet is tightly coupled to the ability to design simple and accurate models.

Modeling has helped in many circumstances. It helped understanding how self-similarity and long-range dependence impact network performance [11], [16]. Models have been used to show that routers require in general more buffer space to absorb a self-similar traffic than Markovian traffic [11]. Traffic models have also been used to generate traffic that captures some characteristics of a real traffic [15], [16]. However, analytical models are most of the time very complex and difficult to use for operational purposes.

Modeling Internet traffic at the packet level has proven to be very difficult since traffic on a link is the result of a high level of multiplexing of numerous flows which behavior is strongly correlated to the transport protocol used and to the application. To solve this problem, a new trend has emerged, which consists in modeling the Internet traffic at the flow level (see [2] and the references therein). Every packet belongs to a flow (e.g., a TCP session, an UDP stream, etc.). Flows arrive at random times and share the available bandwidth in the network according to certain rules (e.g., equal share in case of TCP connections of similar round-trip times). Using Processor Sharing queues, one can get an idea about the response time of a flow and about the distribution of flows running at a certain time in the system.

In this paper we propose a new model whose goal is to characterize the total throughput observed on an IP backbone link. This model uses information at the flow level. The traffic on a backbone can be viewed as the superposition (i.e., multiplexing) of a large number of flows that arrive at random times and that stay active for a random period. In our model, a flow is a generic concept that should be able to capture the characteristics of any kind of data stream. In contrast to other works in the literature (e.g., [2]), we choose to model a link that is not congested (congestion appears elsewhere on the flow path). This assumption is credible on backbone links that are generally over-provisioned (i.e., the level of utilization of a backbone link rarely reaches 60% [8]). This choice is driven by our main objective which is to provide a link dimensioning tool usable on an operational backbone network. We focus on the computation of the moments of the total throughput, namely the average and the variance. The study of the other performance measures for the total throughput (e.g., the correlation) is left for future work.

Our model is very easy to use and requires few simple parameters as input. With some information on the arrival rate of flows, and on the distribution of their size and duration, an ISP can predict the variations of the total throughput on its backbone, and provision links in such a way that a congestion rarely occurs. Moreover, our model opens the door to many future works on a simpler modeling and better understanding of Internet backbone traffic.

In the next two sections we present our flow-based model and its performance analysis. These performances are validated both by simulations (Section IV) and with real traces captured on the Sprint IP backbone (Section V). Conclusions and perspectives on our future work are presented at the end of the paper.

II. The model

Consider data flows that arrive on the link as a Poisson process of constant rate λ (Figure 1). The Poisson assumption can be easily relaxed to more general processes such as MAPs (Markov Arrival Processes) [1], but we will keep working with it for simplicity of the analysis. Poisson might be the right model if we consider recent findings by [5] about the process of flow arrivals in Internet backbone networks. Because of the multiplexing of a large number of flows from many different sources, flows inter-arrival times are closer to those of a Poisson process in a backbone network than in an access network [5]. Even if the backbone traffic is not quite Poisson, the latter property applies to aggregates at the session level [13], [16].

Let T_n ($T_n > 0$ and $n \geq 1$) denote the arrival time of the n-th flow. S_n represents the size of the n-th flow (in bits) and D_n represents its duration (in seconds). We assume that each of the two sequences $\{S_n\}$ and $\{D_n\}$ are iid. However, S_n and D_n are obviously correlated: the time it takes to complete a flow transmission is in average proportional to the size of the flow. Let $X_n(t-T_n)$ denote the throughput of the n-th flow at time t (in bits/s), with $X_n(t-T_n)$ equal to zero for $t < T_n$ and for $t > (T_n+D_n)$. $X_n(t-T_n)$ is a function of S_n, D_n and of the "dynamics" of the flow throughput. For example, for TCP flows, the dynamics of the flow throughput is a function of the dynamics of the window size, which in turn is a function of the round-trip time of the TCP connection, and of the characteristics of the packet loss process [1], [6], [14]. Define $R(t)$ as the total throughput (in bits/s) on the modeled link at time t. We have

$$R(t) = \sum_{n=1}^{\infty} X_n(t - T_n). \tag{1}$$

This model is a Poisson shot-noise process [3], [7] where shots (the shape of the throughput of a flow, see the shaded area in Figure 1 for an example) arrive as a Poisson process of rate λ. In the theory of Poisson shot-noise processes, a shot is the response of a linear system to a Dirac pulse at its input. So the process $R(t)$ is no other than the response of a linear system to a train of Dirac pulses at instants T_n. We use the theory of Poisson shot-noise [3], [7] to state at the end of this section, an expression of the spectral density and of the auto-covariance function of the process $R(t)$.

In the next sections, we will compute the moments of the process $R(t)$ in its stationary regime using techniques from queuing theory. A unique stationary regime exists for finite λ and $\mathbf{E}[D_n]$. In particular, we focus on the computation of the average of $R(t)$ and its variance. We denote by T, S, D, and $X(t-T)$, the time of arrival, the size, the duration, and the throughput at time t of an arbitrary flow, respectively.

III. Performance analysis

A. Average of the total throughput

We suppose that the throughputs of all flows observed on a backbone link are independent of each other. This

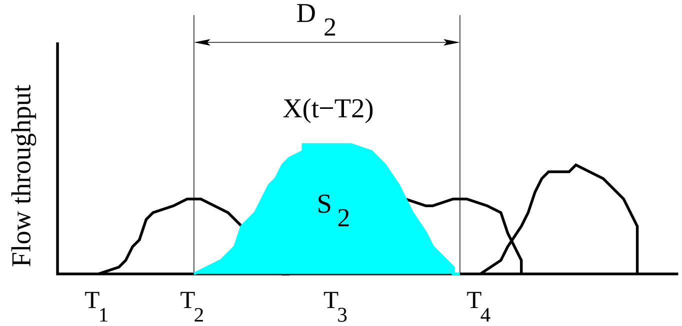

Fig. 1. Traffic modeled as a set of flows

assumption holds since: (i) a backbone link being over-provisioned, it does not experience congestion; (ii) the flows sharing this link have a large number of independent sources and destinations and use many different routes before being multiplexed on the backbone link (i.e., the functions $X_n(u)$ are independent of each other); and (iii) the flow arrival process is Poisson which eliminates any correlation among the instants of arrival. Therefore, we can write using a simple Little argument [10],

$$\mathbf{E}[R(t)] = \lambda \mathbf{E}[D]\, \mathbf{E}^0 [X(t-T)] \, .$$

The superscript 0 indicates that the expectation of $X(t-T)$ is computed under the condition that the flow is running at time t. When computing the expectation, we have to take into account two facts. First, time t is located between the beginning and the end of the flow. This means that we have to compute the expectation of $X(t-T)$ given that the flow has started somewhere between $t-D$ and t (i.e., $t - D < T < t$). The arrival process being Poisson, the starting time of the flow is then uniformly distributed between $t - D$ and t, and it is independent of the starting times of the other active flows [10]. Second, the flow is found running at time t, and time t is randomly chosen. This implies that it is more probable that this flow has a long duration. This second fact is used in queuing theory to compute the distribution of the residual service time of the client in the server when a new client arrives to the queue [10]. Its impact on the expectation will be later clarified.

Let us find the expression of $\mathbf{E}^0[X(t-T)]$, the expectation of the throughput of a flow $X(t-T)$ that we find running at time t. Suppose that this flow has a size S, a duration D and a time of arrival T. As stated above, the time of arrival of the flow is uniformly distributed between $t-D$ and t. Thus, we can write

$$\begin{aligned} &\mathbf{E}^0 [X(t-T)|\, S = s, D = d] \\ &\quad = \mathbf{E}\left[\int_0^D X(u) \frac{\mathrm{d}u}{D} \right| S = s, D = d \left. \vphantom{\int_0^D} \right] \end{aligned}$$

$$= \mathbf{E}\left[\frac{S}{D}\middle| S=s, D=d\right] = \frac{s}{d}.$$

Denote by $f^0_{SD}(s,d)$ the joint distribution of the size and the duration of the flow we find running at time t. As we will see later, this distribution is different from $f_{SD}(s,d)$, the joint distribution of random variables S and D for an arbitrary flow. It follows that,

$$\begin{aligned}
&\mathbf{E}^0\left[X(t-T)\right] \\
&= \int\int f^0_{SD}(s,d)\mathbf{E}^0\left[X(t-T)\middle| S=s, D=d\right] \mathrm{d}s\mathrm{d}d \\
&= \mathbf{E}^0\left[\frac{S}{D}\right].
\end{aligned}$$

Again, this expectation is different from the expectation of the ratio S/D for an arbitrary flow. The average of the total throughput is then equal to

$$\mathbf{E}\left[R(t)\right] = \lambda \mathbf{E}\left[D\right] \mathbf{E}^0\left[\frac{S}{D}\right]. \tag{2}$$

We conclude that the average of $R(t)$ can be computed if we know λ and the two expectations in the right hand side of (2). Fortunately, the second expectation can be further simplified using an argument similar to that used for the computation of the average residual time in queuing theory (e.g., [10, Section 5.2]). This simplification will be better clarified when we present a simple expression for the average of $R(t)$ obtained by the following alternative method.

Alternative method for the computation of $\mathbf{E}\left[R(t)\right]$. The expression of $\mathbf{E}\left[R(t)\right]$ can be obtained by integrating directly the infinite sum in (1),

$$\mathbf{E}\left[R(t)\right] = \lim_{\tau\to\infty} \frac{1}{\tau}\int_0^{\tau} \sum_{n=1}^{\infty} X_n(t-T_n)\mathrm{d}t.$$

By interchanging the order of integration and summation [18], we get

$$\begin{aligned}
\mathbf{E}\left[R(t)\right] &= \lambda \lim_{n\to\infty}\frac{1}{n}\sum_{k=1}^{n}\int_0^{\infty} X_k(t-T_k)\mathrm{d}t \\
&= \lambda \lim_{n\to\infty}\frac{1}{n}\sum_{k=1}^{n} S_k = \lambda \mathbf{E}\left[S\right]. \qquad (3)
\end{aligned}$$

This expression for $\mathbf{E}\left[R(t)\right]$ is very simple (and also intuitive) and it does not need any knowledge of the distribution of flow duration. When comparing this expression with (2), we get

$$\mathbf{E}^0\left[\frac{S}{D}\right] = \frac{\mathbf{E}\left[S\right]}{\mathbf{E}\left[D\right]}. \tag{4}$$

In the rest of this subsection we will justify equality (4). This justification will help us in the computation of the other moments of the total throughput. But before that, we highlight the fact that the simple expression of $\mathbf{E}\left[R(t)\right]$ in (3) is valid for any flow arrival process, and not only for Poisson. The Poisson assumption for arrivals will be used to approximate the higher moments of the total throughput.

For the interpretation of (4), we use a technique issued from queuing theory to compute the distribution of the residual service time when a new client arrives in a queue (see e.g., [10, Section 5.2]). This technique says that it is more probable that the flow we find running at time t has a long duration, simply because the time t is randomly chosen. The distribution of the duration of this particular flow must be scaled with its duration in a way to give more weight to large values of D. Using Equation (5.8) in [10], we write

$$f^0_{SD}(s,d) = f_{SD}(s,d)\frac{d}{\mathbf{E}\left[D\right]}.$$

It follows that,

$$\begin{aligned}
\mathbf{E}^0\left[\frac{S}{D}\right] &= \int\int f^0_{SD}(s,d)\frac{s}{d}\mathrm{d}s\mathrm{d}d \\
&= \int\int f_{SD}(s,d)\frac{d}{\mathbf{E}\left[D\right]}\frac{s}{d}\mathrm{d}s\mathrm{d}d \\
&= \frac{\mathbf{E}\left[S\right]}{\mathbf{E}\left[D\right]}.
\end{aligned}$$

Substituting this expression in (2) yields (3).

B. *Variance of the total throughput*

The variance of the total throughput is the second performance measure an ISP needs to know in order to properly dimension network links. A backbone link has to be provisioned in order to absorb the average of the total throughput as well as its variations.

Let $N(t)$ denote the number of flows running at time t (i.e., active flows). $N(t)$ represents the number of clients in a dual M/G/∞ queue where the arrival of a flow corresponds to the arrival of a client to the queue, and where the end of a flow corresponds to the departure of the corresponding client. Thus, the process $N(t)$ can be completely characterized using known results on M/G/∞ queues (see e.g., [9], [10]). Given the assumption that the throughputs of the different flows are independent of each other, we can write [10],

$$V_R = \mathbf{E}\left[N(t)\right] V^0_X + (\mathbf{E}^0\left[X(t-T)\right])^2 V_N,$$

where V_R, V^0_X and V_N denote respectively the variances of the total throughput, of the throughput of a flow running at time t, and of $N(t)$. In particular, $V_X^0 = \mathbf{E}^0\left[X^2(t-T)\right] - \left(\mathbf{E}^0\left[X(t-T)\right]\right)^2$.

$\mathbf{E}\left[N(t)\right]$ and $\mathbf{E}^0\left[X(t-T)\right]$ are respectively equal to $\lambda\mathbf{E}\left[D\right]$ and $\mathbf{E}\left[S\right]/\mathbf{E}\left[D\right]$. V_N can be computed from the dual M/G/∞ queuing model [9]. Indeed, we have

$$\mathbf{P}\left\{N(t)=k\right\} = \frac{\left(\lambda\mathbf{E}\left[D\right]\right)^k}{k!}e^{-\lambda\mathbf{E}[D]}, \qquad k=0,1,2\ldots.$$

$\lambda\mathbf{E}\left[D\right]$ denotes the load of the queue and the required condition for the stability of the system is for $\lambda\mathbf{E}\left[D\right]$ to be finite. This guarantees that the number of active flows does not grow to infinity. The probability generating function (PGF) of $N(t)$ is then equal to

$$\tilde{N}(z) = e^{\lambda\mathbf{E}[D](z-1)}. \qquad (5)$$

Hence,

$$V_N = \frac{\mathrm{d}^2\tilde{N}}{\mathrm{d}z^2}(z=1) + \mathbf{E}\left[N(t)\right] - \left(\mathbf{E}\left[N(t)\right]\right)^2 = \lambda\mathbf{E}\left[D\right],$$

and,

$$V_R = \lambda\mathbf{E}\left[D\right]\left(\mathbf{E}^0\left[X^2(t-T)\right]\right). \qquad (6)$$

The variance of $R(t)$ is therefore a function of $\lambda\mathbf{E}\left[D\right]$ and of the second moment of the throughput of a flow we found running at time t (i.e., $X(t-T)$). The computation of the second moment of $X(t-T)$ requires some assumptions (or more information) on the dynamics of flow throughput. Later, we will provide approximations of the variance of $R(t)$ for some particular $X(t-T)$ functions. Before that, we generalize our result to the higher moments of $R(t)$.

C. Higher moments of the total throughput

All the moments of the total throughput (and hence its distribution) can be computed from the moments of $X(t-T)$ and the PGF of the random variable $N(t)$ (Equation (5)). We have the following relation between Laplace Stieltjes Transforms (LST) [10],

$$\tilde{R}(w) = \tilde{N}(\tilde{X^0}(w)), \qquad (7)$$

where $\tilde{R}(w) = \mathbf{E}\left[e^{-wR(t)}\right]$, $\tilde{X^0}(w) = \mathbf{E}^0\left[e^{-wX(t-T)}\right]$, and w is a complex number with positive real part. Again, the LST of $X(t-T)$ has to be computed under the condition that the flow is running at time t. Equation (7) tells us that the computation of the k-th moment of the total throughput requires the computation of the k first moments of $X(t-T)$, that is, all moments up to $\mathbf{E}^0\left[X^k(t-T)\right]$. Using an argument similar to that we used for the computation of $\mathbf{E}^0\left[X(t-T)\right]$, we write

$$\mathbf{E}^0\left[X^k(t-T)\right] = \mathbf{E}^0\left[\frac{1}{D}\int_0^D X^k(u)\mathrm{d}u\right]$$

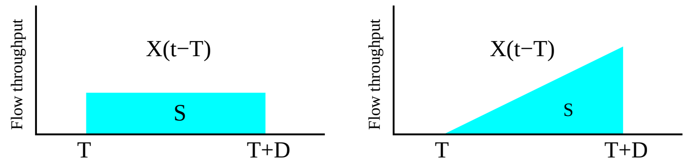

(a) Constant X approximation (b) Linear X approximation

Fig. 2. Simple flow throughput approximations

$$= \frac{1}{\mathbf{E}\left[D\right]}\mathbf{E}\left[\int_0^D X^k(u)\mathrm{d}u\right]. \qquad (8)$$

Hence, the total throughput can be completely characterized if we can compute all the expectations $\mathbf{E}\left[\int_0^D X^k(u)\mathrm{d}u\right]$. For $k \geq 2$, this computation requires more information on the process $X(t-T)$ (i.e., on the dynamics of flow throughput) than the sole knowledge of distributions of S and D for an arbitrary flow. For example, in the case of a TCP flow, we can compute these expectations from the knowledge of the round-trip time of the flow and its packet losses characteristics [1], [6], [14]. In the present work, we assume that the only information we have on a flow is the distributions of S and D. In future work, we will specifically focus our study on TCP flows (since they are the most dominant in the Internet [8]) and introduce other parameters so as to compute more accurately the moments of $X(t-T)$.

D. Some particular cases for the computation of the moments of $R(t)$

If only the distributions of S and D are available, some simple assumptions must be made on the process $X(t-T)$ in order to compute the moments (of order higher than 2) of the total throughput. We consider here two particular cases shown in Figure 2: the constant and the linear flow throughput. We focus on the computation of the variance of $R(t)$, which can be written, from (6) and (8), as follows:

$$V_R = \lambda\mathbf{E}\left[\int_0^D X^2(u)\mathrm{d}u\right]. \qquad (9)$$

D.1 Constant throughput approximation

First, we consider the case where the throughput of a flow is constant and equal to S/D (Figure 2a). Here, we need to compute all expectations $\mathbf{E}\left[S^k/D^{k-1}\right]$ in order to fully characterize the total throughput. The average of the total throughput requires the expression of $\mathbf{E}\left[S\right]$. The variance of the total throughput requires the expression of $\mathbf{E}\left[S^2/D\right]$, and so on. In particular, the variance of $R(t)$ is equal to

$$V_R = \lambda \mathbf{E}\left[\frac{S^2}{D}\right]. \quad (10)$$

The previous assumption is the simplest one. Expression (10) only captures the variation of the total throughput caused by the variation of $N(t)$ and by the variation of the ratio S/D. However, in most cases (e.g., TCP flows [6]), the throughput of a flow increases from low values at the beginning up to its stationary regime. The assumption that the throughput of a flow is constant underestimates the moments of $X(t-T)$, and hence the moments of the process $R(t)$: knowing only S and D, we cannot capture the dynamics of the throughput of a flow. However, we can make other assumptions that lead to a better approximation of the flow throughput.

D.2 Linear throughput approximation

A better assumption is to consider that the throughput of a flow linearly increases with time (Figure 2b). This assumption is inspired from the dynamics of short TCP transfers that form about 70% of the flows on the Sprint IP backbone [8]. For a flow with size S and duration D, the throughput is assumed to increase linearly from zero to $2S/D$, with a mean equal to S/D. At a time t between T and $T+D$, we can write $X(t-T) = (2S/D^2)(t-T)$. Hence, for $k \geq 1$,

$$\mathbf{E}^0\left[X^k(t-T)\right] = \frac{2^k}{k+1}\frac{1}{\mathbf{E}[D]}\mathbf{E}\left[\frac{S^k}{D^{k-1}}\right].$$

Again, we need to compute all the expectations $\mathbf{E}\left[S^k/D^{k-1}\right]$ in order to fully characterize the total throughput. For the variance of $R(t)$ we have in the linear case,

$$V_R = \frac{4\lambda}{3}\mathbf{E}\left[\frac{S^2}{D}\right]. \quad (11)$$

We notice that the variance is larger in the linear case than that obtained when we assume that the throughput of a flow is constant and equal to S/D (Equation (10)).

The linear throughput assumption is not the only one and we can always consider other approximations of throughput dynamics using S and D (log, square root, exponential, etc.). We can also mix various flow throughput models using the notion of classes we will later introduce. We keep these evolutions for future work.

E. *Moments of $R(t)$ and averaging interval*

In reality, the instantaneous total throughput is computed by averaging the number of packets that cross a link during short time intervals. Thus, except for the first moment, the moments of $R(t)$ strongly depend on the given averaging interval: the longer the averaging interval, the smoother the total throughput. Thus, before making any assumption on the dynamics of the throughput during a flow, the ISP using our model must define the averaging interval (s)he wants to use. In other words, (s)he must define the smallest interval below which variations of the total throughput are not important. One can take as averaging interval the maximum queuing time at the input of the link. Then, we have to develop a model for the dynamics of flow throughput that is able to capture the remaining variations of the measured total throughput. For example, for an averaging interval equal to δ_a seconds, the model for the flow throughput must be able to capture frequencies up to $1/\delta_a$. On the other hand, the throughput of a flow has an upper bound on its frequency. For example, we know that the throughput of a TCP connection implementing delayed acknowledgements remains stable during two round-trip times in the steady state [14]. Denote by $1/\delta_f$ the maximum frequency of the flow throughput. Hence, the model for the flow throughput must capture frequencies up to $\min(1/\delta_a, 1/\delta_f)$, otherwise some error will occur in the approximation of the moments of $R(t)$. Later, we will see how our previous approximations underestimate the moments of $R(t)$ since they do not capture the high frequencies of the flow throughput variations. We will also see that a better match can be obtained if we slightly increase the averaging interval.

F. *Distribution of the total throughput*

The number of active flows at a given time in a backbone is large [8]. Hence, by the Central Limit Theorem, $R(t)$ can be supposed to have a normal distribution with average $\mathbf{E}[R(t)] = \lambda\mathbf{E}[S]$ and with variance V_R. The probability density function of $R(t)$ can be written as,

$$f_R(r) = \frac{1}{\sqrt{2\pi V_R}} e^{-\frac{(r-E[R])^2}{2V_R}}.$$

Let us mention that, if one knows the LST of $X(t-T)$, one can use Equation (7) to find the LST of $R(t)$, and then invert it back to get the exact expression of the distribution of $R(t)$.

G. *Extension to the case of multiple classes of flows*

The previous model can be easily extended to the case where the link is crossed by multiple classes of flows. This will be interesting when the flows have different size and duration distributions, and therefore require a mixture of models for flow throughput dynamics. For example, on the

Sprint backbone, approximately 60% of the flows are short lived TCP flows that would be better modeled by our linear case (or probably an exponential flow throughput model given the exponential slow start phase of TCP [6]), and almost 10% are long-lived flows that could be better approximated as rectangles (i.e., constant flow throughput).

Suppose that we have I classes of flows. Flows of each class arrive at the link following a Poisson process of rate $\lambda^{(i)}$, with $1 \leq i \leq I$. The flows of a class have sizes described by random variable $S^{(i)}$ and have durations described by random variable $D^{(i)}$. Our previous results can be extended to this case in the following way,

$$\lambda = \sum_{i=1}^{I} \lambda^{(i)}, \qquad \mathbf{E}\left[S^k\right] = \sum_{i=1}^{I} \frac{\lambda^{(i)}}{\lambda} \mathbf{E}\left[(S^{(i)})^k\right],$$

$$\mathbf{E}\left[D^k\right] = \sum_{i=1}^{I} \frac{\lambda^{(i)}}{\lambda} \mathbf{E}\left[(D^{(i)})^k\right],$$

$$\mathbf{E}^0\left[X^k(t-T)\right] = \frac{1}{\mathbf{E}[D]} \mathbf{E}\left[\int_0^D X^k(u)\mathrm{d}u\right]$$

$$= \frac{1}{\mathbf{E}[D]} \sum_{i=1}^{I} \frac{\lambda^{(i)}}{\lambda} \mathbf{E}\left[\int_0^{D^{(i)}} (X^{(i)}(u))^k \mathrm{d}u\right].$$

We can also start from the characterization of the total throughput for each class (call this process $R^{(i)}(t)$), and then characterize the process $R(t)$, which is no other than the sum of these throughputs. The characterization of the total throughput for a class can be easily done by using our previous results with only the parameters of that class. The moments of the process $R(t)$ can be immediately deduced from those of the processes $R^{(i)}(t)$. In particular,

$$\mathbf{E}[R(t)] = \sum_{i=1}^{I} \mathbf{E}\left[R^{(i)}(t)\right], \qquad V_R = \sum_{i=1}^{I} V_{R^{(i)}}.$$

H. Poisson shot-noise and spectral density

As we highlighted at the beginning, the process $R(t)$ is a Poisson shot-noise process [3], [7], where $X(t-T)$ is the response of a linear system to a Dirac pulse that arrives at time T. $R(t)$ is then the convolution of the responses of the system to all the Dirac pulses that arrive before t. Denote by $\tilde{X}(f)$ the Fourier transform of the function $X(t)$ (called the Transfer function of the system) and by $S_R(f)$ the spectral density of the centred process $R(t) - \mathbf{E}[R(t)]$,

$$\tilde{X}(f) = \int_{-\infty}^{+\infty} X(t) e^{-j2\pi f t} \mathrm{d}t$$

$$S_R(f) = \int_{-\infty}^{+\infty} C_R(\tau) e^{-j2\pi f \tau} \mathrm{d}\tau,$$

with $C_R(\tau) = \mathbf{E}[R(t-\tau)R(t)] - (\mathbf{E}[R(t)])^2$ being the auto-covariance function of process $R(t)$. Using Theorem 1 in [4], we find

$$S_R(f) = \lambda \mathbf{E}\left[\left|\tilde{X}(f)\right|^2\right].$$

We can also invert the expression of the spectral density to get the auto-covariance function of $R(t)$. With some transformations, we obtain the following result for $\tau \geq 0$,

$$C_R(\tau) = \lambda \mathbf{E}\left[1\{D > \tau\} \int_0^{D-\tau} X(u)X(u+\tau)\mathrm{d}u\right].$$

For the particular case $\tau = 0$, we clearly see (as one should expect) how the expression of $C_R(\tau)$ coincides with that of the variance of $R(t)$ given by (9).

IV. Simulations

In this section, we present a validation of the proposed model by simulation. We use the ns simulator [12] to study two scenarios. In the first scenario, all flows have the same size but different durations. In the second scenario, both the size and duration of flows change during each simulation. In both scenarios, a flow corresponds to a TCP connection traversing the modeled link in one direction.

A. Simulation scenario

A set of TCP Newreno sources transmit files of a fixed size for a duration of 1000 seconds. Each flow crosses a 10Mbps link (the modeled link) on their path to destination. The delayed acknowledgement option of TCP is used and packets have a size of 500 bytes. Before arriving at the 10Mbps link, all connections experience a packet loss probability of 3% (to introduce some randomness in the TCP flows). TCP flows are generated according to a Poisson process. The rate of the Poisson process and the size of the files are chosen in such a way that the 10Mbps link always remains under-utilized. The round-trip time of all TCP connections is set to 80ms. Finally, we compute the rate with which data packets cross the 10Mbps link and we store the variation of this rate as a function of time. We also measure the size and the duration of the different TCP transfers, which produces samples for the processes $\{D_n\}$ and $\{S_n\}$. The instantaneous total throughput is measured by averaging the number of packets that cross the 10Mbps link over intervals of 200ms ($\simeq$ twice the round-trip time). The moments of the total throughput are first computed by our model using λ and the samples of $\{S_n\}$ and $\{D_n\}$, and second by measurement using directly the samples of the process $R(t)$.

B. Constant-size flows

We set the arrival rate of TCP transfers to 1 transfer per second and we give the file size different values between 10 Kbytes and 1 Mbytes. We run a simulation for each value of the file size. Other rates for the Poisson process were also considered; similar results were found.

We compare first the measured average total throughput to that given by our model for the different file sizes. Figure 3 shows a good match of simulation and modeling results. The x-axis in the figure shows the log of the file size. The 95% confidence intervals are so narrow that we decided not to plot them. We verify in Figure 3 how the utilization of the 10Mbps link increases with the size of the files.

In Figure 4 we plot the coefficient of variation of the total throughput ($\sqrt{V_R}/\mathbf{E}\left[R\right]$) given by the simulation for the different file sizes. We compare the coefficient of variation of the total throughput given with the two particular cases of our model: the constant throughput case (Equation (10)) and the linear throughput case (Equation (11)). The expectation of the ratio S^2/D required by our model is computed from the samples $\{S_n\}$ and $\{D_n\}$. We make two major observations:

- The linear throughput case gives a better approximation of the coefficient of variation than the constant-throughput case, and both cases underestimate the variation of the total throughput. This is because neither models captures accurately the variation of the flow throughput caused by the saw-tooth nature of TCP congestion control [1], [14]. By increasing the averaging interval of 200ms, we can absorb TCP throughput variations and, therefore, improve the performance of our two approximations.
- We observe that the coefficient of variation of the total throughput decreases with the size of files which indicates that the total throughput becomes smoother and smoother. This is because the coefficient of variation of $N(t)$, the number of active flows, decreases. Indeed, the coefficient of variation of $N(t)$ is equal to $1/\sqrt{\lambda \mathbf{E}\left[D\right]}$, which decreases with any increase in λ or $\mathbf{E}\left[D\right]$. In our case here, we have an increase in $\mathbf{E}\left[D\right]$ due to an increase in S.

Figure 5 shows the histogram of the total throughput for a simulation with file sizes of 500 Kbytes. We also plot in the same figure the distribution of a normal random variable having the same average and variance as $R(t)$. We notice a good match of the two distributions which is due to the large number of concurrent TCP flows.

An ISP could use our model (together with the distribution of the normal random variable) to dimension the links of its backbone network. For example, we know that in 70% of time, the total throughput is located somewhere

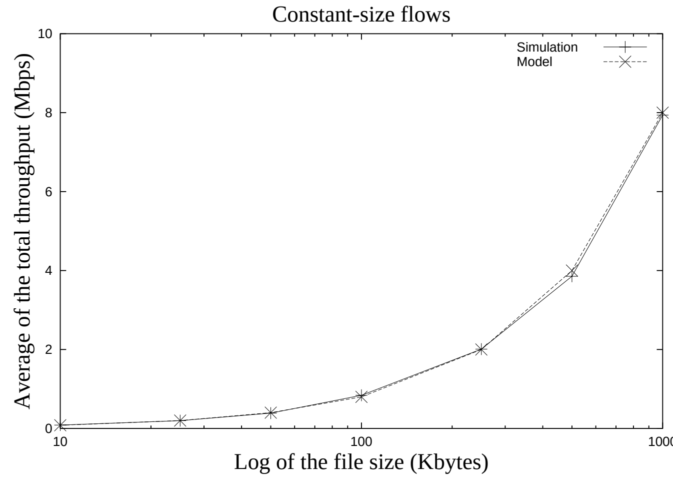

Fig. 3. Average of the total throughput for constant flow size simulations

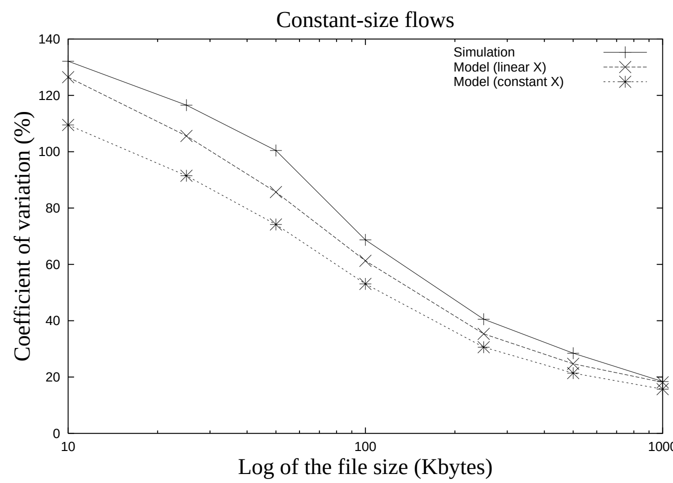

Fig. 4. Coefficient of variation of the total throughput for constant flow size simulations

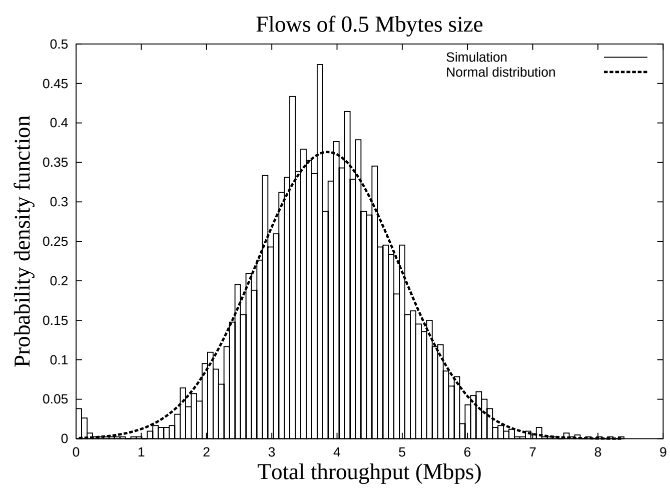

Fig. 5. Histogram of the total throughput for a simulation with flows of 0.5Mbytes size

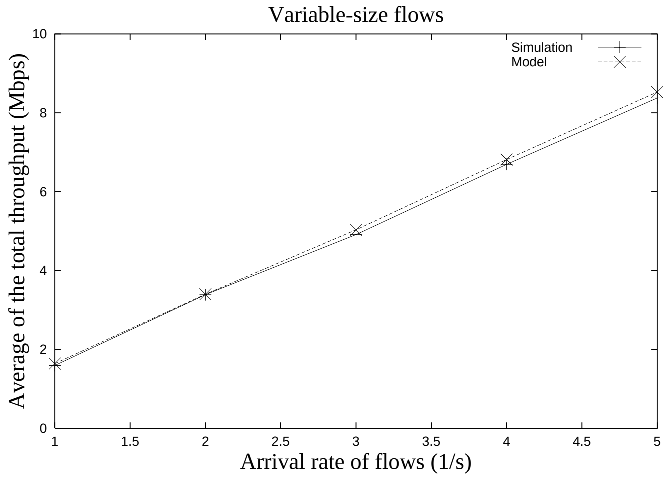

Fig. 6. Average of the total throughput for variable flow size simulations

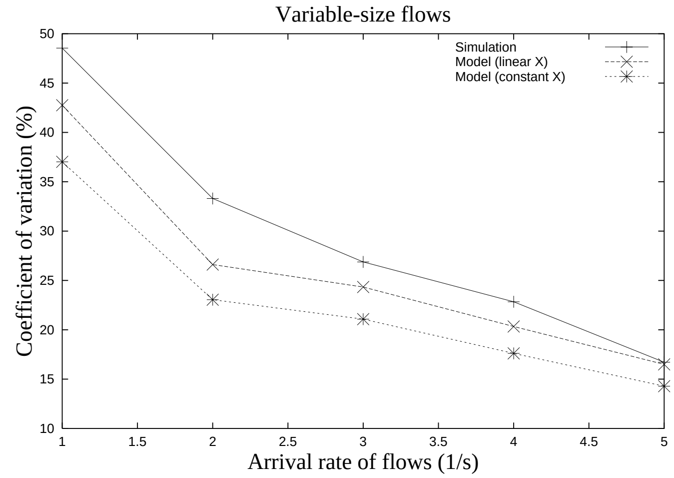

Fig. 7. Coefficient of variation of the total throughput for variable flow size simulations

between $\mathbf{E}[R] - \sqrt{V_R}$ and $\mathbf{E}[R] + \sqrt{V_R}$ [9]. One can also use Large Deviations techniques [17] to find an approximate value for the probability that the total throughput exceeds a certain level beyond $\mathbf{E}[R] + \sqrt{V_R}$. The bandwidth of a link can then be chosen as a function of the demand and the distribution of flow sizes and durations, in such a way that congestion is avoided or that some level of QoS is met.

C. *Variable-size flows*

We repeat the previous simulation with variable file sizes. For each transfer, we pick a real number randomly between 1 and 3 with a uniform distribution. The size of the file in Kbytes is then 10 power the selected number: it results in an average file size equal to 214 Kbytes. We give the Poisson arrival process a rate between 1 and 5 transfers per second. For each rate, we run a set of simulations for a duration of 1000s.

We measure the average of the total throughput as well as its coefficient of variation and we plot the results in Figures 6 and 7, respectively. The x-axis in both figures corresponds to the different rates chosen for the Poisson process. We also plot in both figures the approximation of the two measures given by our model. In the second figure, we consider again the two approximations of flow throughput dynamics: constant and linear. As in the previous section, both figures show a good match of simulation and modeling results, even though the size of the files changes during the simulation. We verify in the figures that the utilization of the 10Mbps link increases and that the total throughput smoothes as the arrival rate of the flows increases.

Trace	Length	Average Link Utilization
1	24h	44 Mbps
2	8h 30m	28 Mbps
3	11h 20m	77 Mbps
4	3h 40m	101 Mbps
5	9h 30m	67 Mbps
6	24h	17 Mbps
7	10h	42 Mbps
8	18h 40m	54 Mbps
9	13h 20m	76 Mbps
10	8h	33 Mbps

TABLE I
SUMMARY OF TRACES

V. EXPERIMENTAL VALIDATION

In this section we compare the prediction realized with our model to measurements carried out on the Sprint IP backbone.

The data we use were collected from OC-3 backbone links (155 Mbps) using the passive monitoring infrastructure described in [8] and deployed in some of the Point-of-Presence (POPs) of the network. In short, the infrastructure consists of passive monitoring systems that tap the optical link to collect packet level traces between access routers and backbone routers. Every packet is timestamped using a GPS clock signal which provides accurate timing information.

In this study we use data from 10 different internal POP links collected on August 9th, 2000 starting at 17:00 UTC. Table I provides a summary of traces.

The traces have different link utilizations (ranging from around 20Mbps to around 100Mbps), resulting in different trace lengths. The 10 monitored links are provisioned so that congestion never happens.

A. Methodology

We compare the average and the variance of the total throughput of the collected traces with the results obtained from our model when the input data (i.e., flow arrival rate and moments of S and S^2/D) are directly derived from the traces.

We divide each trace into 30 minutes intervals (to keep the arrival rate stationary, and to get many points for comparison). For each interval, we compute the average of the total throughput and its coefficient of variation, the number of flows and, for each flow, its size and duration. Samples of the instantaneous total throughput are computed using averaging intervals of 200ms (Section III-E). This is comparable with the average value of the round-trip time we measure on these links. We also present results obtained with an averaging interval equal to 500ms to show that we can smooth the total throughput by increasing the averaging interval, and thus, improve the match between the experimental observations and our model.

In the backbone trace, we define a flow as a set of packets having the same protocol number, source and destination addresses and port numbers. The size of a flow is measured in bytes, while the duration is equal to the time difference between the first and the last packet of the flow. A flow made of only one packet is discarded (the duration would be zero), although that packet is counted for the purpose of the average and variance of the total throughput.

We do not distinguish among TCP and UDP flows. Long-lived flows (i.e., flows that last longer than the 30 minutes interval) are divided into several pieces, and are fed as input to the model, with a smaller size and duration over many intervals.

B. Results

Figure 8 shows the average total throughput computed on each 30 minutes interval from the packet-level trace and the average total throughput given by our model. The measured total throughput is plotted on the x-axis, while the y-axis shows the corresponding average total throughput evaluated according to the model. A point on the diagonal crossing the graphs represents a perfect match between the model and the measurements. As we can see from the graph, the model gives a very accurate estimate of the average of the total throughput over the entire range of link utilizations. Indeed, our simple expression for the average of the total throughput holds for any process of flow arrivals and for any dynamics of flow throughput. Clearly,

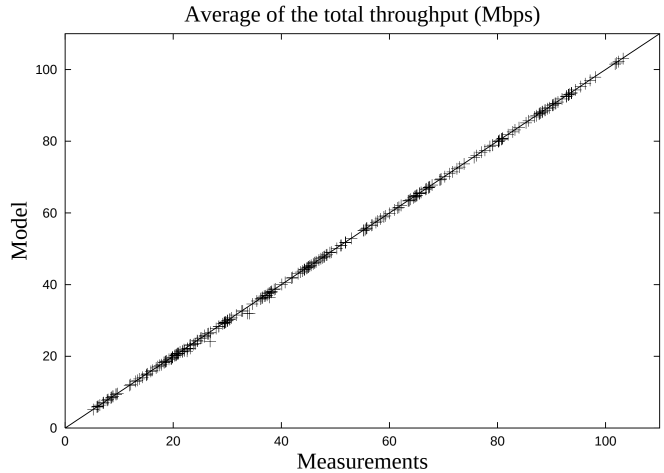

Fig. 8. Average of the total throughput

this is not the case for our expressions of the higher moments of the total throughput which are sensitive to the Poisson assumption and to the dynamics of flow throughput. Nevertheless, we will see in the next two figures that our model still gives good approximations of these higher moments. Note also that our expression of the average of the total throughput is not sensitive to the choice of the averaging interval.

In Figures 9 and 10, we compare the coefficient of variation given by measurements and that given by our model with rectangular (constant throughput) and triangular (linearly increasing throughput) shots. Both show a good match with better results for the linear throughput case; they however under-estimate the variation of the real total throughput since they do not capture all the dynamics of flow throughput. Most flows being TCP [8], the flow throughput varies more frequently than our model can predict, which results in the underestimation. Nevertheless, the result is still quite good if we look at the small number of parameters required by the model. To capture exactly the dynamics of the flow throughput, we would need more information about the flows. We are currently investigating the worthiness of introducing more parameters into the model to refine the approximations for the moments of the total throughput.

As explained in Section III-E, we can smooth the samples of the total throughput by increasing the averaging interval. When smoothing, we will eliminate some of the oscillations of the total throughput that are not captured by our approximations for flow dynamics, and this should improve the results. Figure 11 provides the same results as Figure 10 with an averaging interval equal to 500ms (instead of 200ms). Figure 11 confirms our expectations and our model gives a better approximation in this case.

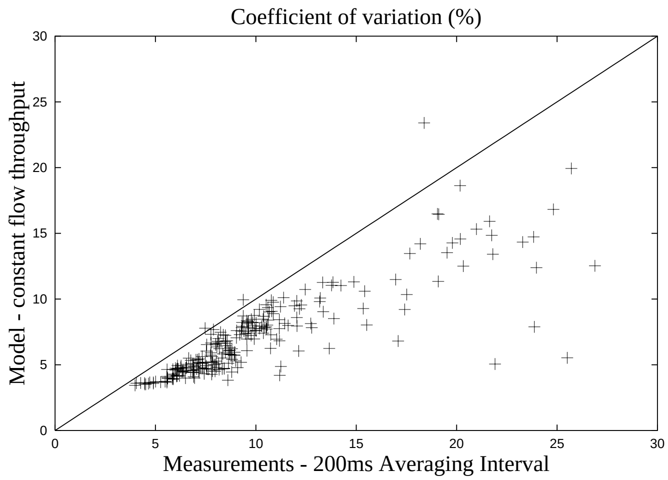

Fig. 9. Coefficient of variation of the total throughput for the constant-rate case

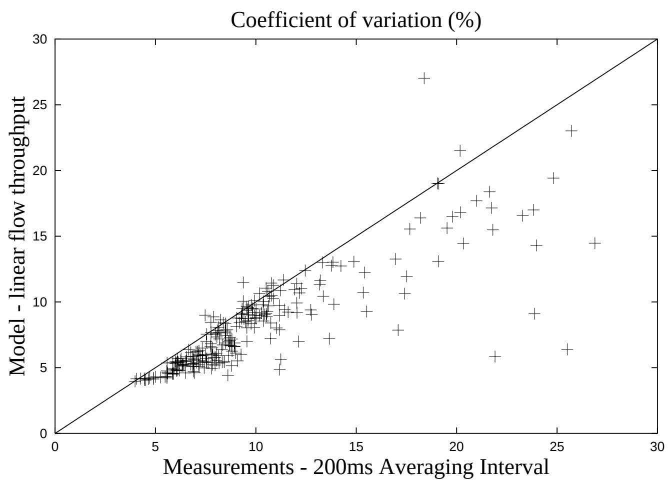

Fig. 10. Coefficient of variation of the total throughput for the linear-rate case

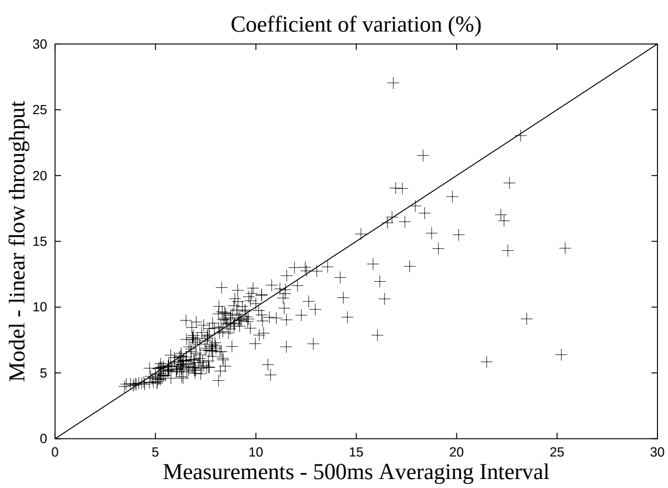

Fig. 11. Coefficient of variation of the total throughput for the linear-rate case and a 500ms averaging interval

VI. CONCLUSIONS

In this paper, we proposed a simple model for the traffic on an IP backbone network. The model is inspired from Poisson shot-noise and M/G/∞ queue theories. With only 3 parameters (λ, rate of flows, $\mathbf{E}[S]$, average size of a flow, and $\mathbf{E}[S^2/D]$, average value of the ratio of the square of a flow size and its duration), the model is able to find good approximations for the total throughput on a backbone link and for its variations. We believe that this will be very useful for dimensioning and provisioning backbone links so as to absorb the demand of Internet users and ISPs.

We are working on various extensions of the present work. We stated in the paper a result for the auto-covariance function of the total throughput. Using this result, we are investigating the correlation of Internet traffic and its relation with the flow arrival rate and distributions of flow size and flow duration. We are also designing flow models that are specific to TCP. Using information on the packet loss rate and the round-trip time of flows, we are computing better approximations for the moments of flow throughput. Finally, we are studying the worthiness of introducing classes of flows and of considering more complex flow arrival processes than Poisson. The results are promising but the challenge is to improve our evaluation of the throughput without increasing much the complexity of the model. We want the model to be usable with current management tools used by ISPs.

REFERENCES

[1] E. Altman, K. Avratchenkov, and C. Barakat, ”A stochastic model for TCP/IP with stationary random losses”, *ACM SIGCOMM*, September 2000.

[2] S. Ben Fredj, T. Bonald, A. Proutiere, G. Regnie and J. Roberts, ”Statistical Bandwidth Sharing: A Study of Congestion at Flow Level”, *ACM SIGCOMM*, August 2001.

[3] P. Brémaud, “An insensitivity property for light-tail shot noise traffic overflow asymptotics”, *Technical Report DSC-EPFL*, Feb. 2000.

[4] P. Brémaud and L. Massoulié, “Power spectra of general shot noises and Hawkes point processes with a random excitation”, preprint, private communication.

[5] J. Cao, W. Cleveland, D. Lin and D. Sun, “Internet traffic tends to Poisson and independent as the load increases”, *Bell Labs Tech Report*, 2001.

[6] N. Cardwell, S. Savage and T. Anderson, ”Modeling TCP Latency”, *IEEE INFOCOM*, March 2000.

[7] D. Daley and D. Vere-Jones, “An introduction to the theory of point processes”, *Springer-Verlag*, 1988.

[8] C. Fraleigh, C. Diot, B. Lyles, S. Moon, P. Owezarski, D. Papagiannaki and F. Tobagi, “Design and Deployment of a Passive Monitoring Infrastructure”, *Passive and Active Measurement Workshop*, April 2001.

[9] R. Jain, ”The art of computer systems performance analysis”, *Wiley*, 1991.

[10] L. Kleinrock, ”Queueing Systems”, *Wiley*, 1975.

[11] W. Leland, M. Taqq, W. Willinger and D. Wilson, ”On the self-similar nature of Ethernet traffic,” *ACM SIGCOMM*, September 1993.

[12] The LBNL Network Simulator, ns, http://www.isi.edu/nsnam/ns/

[13] C. Nuzman, I. Saniee, W. Sweldens and A. Weiss, “A Compound Model for TCP Connection Arrivals”, *ITC workshop*, Sep. 2000.

[14] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, “Modeling TCP Throughput: a Simple Model and its Empirical Validation”, *ACM SIGCOMM*, September 1998.

[15] M. Parulekar and A. M. Makowski, ”M/G/Infinity Input Processes: A Versatile Class of Models for Network Traffic”, *IEEE INFOCOM*, Apr. 1997.

[16] V. Paxson and S. Floyd, ”Wide-Area Traffic: The Failure of Poisson Modeling”, *IEEE/ACM Transactions on Networking*, vol. 3, no. 3, pp. 226-244, June 1995.

[17] A. Weiss, ”An introduction to large deviations for communication networks”, *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 6, pp. 938-952, October 1995.

[18] R. Wolf, ”Stochastic modeling and the theory of queues”, *Printice Hall*, 1989.