

On Internet backbone traffic modeling *

Chadi Barakat[†]
INRIA
Chadi.Barakat@sophia.inria.fr

Patrick Thiran
LCA - EPFL
Patrick.Thiran@epfl.ch

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

1. INTRODUCTION

The motivation of this work is to design a traffic model that can be used in routers or by network administrators to assist in network design and management. Currently, network operators have very basic information about the traffic. They mostly use SNMP, which provides average throughput information over 5 minutes intervals. An analytical model can provide more accurate information on the traffic such as its variation and its auto-correlation at short timescales.

In contrast to other works (see [2] and the references therein), we choose to model the traffic on a link that is *not* congested (congestion possibly appears elsewhere in the Internet). This assumption is valid (and in fact is the rule) for backbone links that are generally over-provisioned (i.e., the network is designed so that a backbone link does not reach 50% utilization in the absence of link failure [4]). This choice is driven by our main objective, which is to provide a link dimensioning tool usable in backbone network management.

We opt for a model of the traffic at the flow level. Modeling the traffic at the packet level is very difficult, since traffic on a link is the result of a high level of multiplexing of numerous flows whose behavior is strongly influenced by the transport protocol and by the application. A flow in our model is a very generic notion. It can be a TCP connection or a UDP stream (described by source and destination IP addresses, source and destination port numbers and the protocol number), or it can be a destination address prefix (e.g., destination IP address in the form a.b.0.0/16). The definition of a flow is deliberately kept general, which allows our model to be applied to different applications and to different transport mechanisms. The model can however be specified to some particular traffic types such as FTP and HTTP. By specifying the model to a certain traffic type, one must expect to obtain better results.

Data flows arrive to a backbone link at random times, trans-




Figure 1: Backbone traffic modeled as a multiplexing of flows represented by "shots"

port a random volume of data, and stay active for random periods. Given information on flows, our model aims to compute the total (aggregate) rate of data observed on the backbone link. We are interested in capturing the dynamics of the total data rate at short timescales (i.e., of the order of hundreds of milliseconds). This dynamics can be completely characterized using simple mathematical tools, namely the shot-noise process [3]. Our main contribution is the computation of simple expressions for important measures of backbone traffic such as its average, its variance, and its auto-correlation function. These expressions are functions of a few number of parameters that can be easily computed by a router (e.g., using a tool such as NetFlow, which provides flow information in Cisco routers¹).

Our model can be helpful for managing and dimensioning IP backbone networks. Knowing the average and the variance of the traffic allows an ISP to provision the links of its backbone so as to avoid congestion. Congestion can be avoided at short timescales of the order of hundreds of milliseconds. The auto-correlation function of the traffic can be used to propose predictors for its future values. The prediction of the traffic has diverse applications in managing the resources of the backbone. One interesting application is the use of a short-term prediction to optimize packet routing and load balancing. Our model can also be used to assess the impact on backbone traffic of changes made in the rest of the Internet such as the addition of a new customer, a new application, or a new transport mechanism. The ISP can plan the provisioning of its backbone so as to absorb the resulting change of traffic before this change takes place.

2. THE GENERAL MODEL

Let $T_n, n \in \mathbb{Z}$, denote the arrival time of the n -th flow to the backbone link under consideration. Let S_n and D_n denote the size and the duration of the n -th flow. The size of a flow is the volume of data it transports during its lifetime. We

*More details on this work can be obtained upon request from the authors.

[†]The work has been done while the author was working at EPFL, with a grant from Sprint Labs.

¹<http://www.cisco.com/warp/public/732/Tech/netflow>

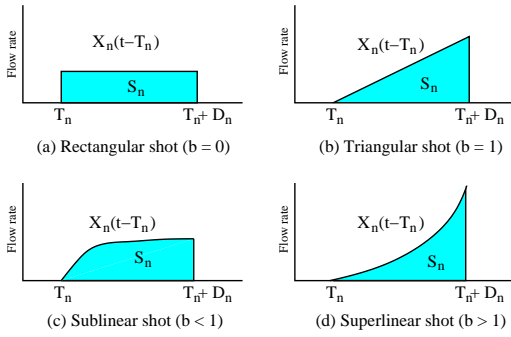


Figure 2: Simple models for shots

introduce the set of functions $\{X_n(\cdot)\}$, which we call *flow rate functions* or *shots*. We define $X_n(u)$ as the data rate of flow n at time $T_n + u$. We take $X_n(u) = 0$ for $u < 0$ and $u > D_n$. The total rate of data through the backbone link, which we denote by $R(t)$, is the result of multiplexing of all shots: $R(t) = \sum_{n \in \mathbb{Z}} X(t - T_n)$. This multiplexing is illustrated in Figure 1. Functions $\{X_n(\cdot)\}$ are assumed to be independent and identically distributed. The total data rate $R(t)$ can be seen as a shot-noise process [3]. A shot-noise is the response of a linear system to a train of Dirac pulses arriving at times T_n , with the shot $X_n(\cdot)$ being the impulse response of the system. We use results on shot-noise to compute the most important measures of backbone traffic, namely its average, its variance, and its auto-correlation function. The expressions we obtain are function of a small number of parameters: (i) the average arrival rate of flows $\lambda = 1/\mathbb{E}[T_n]$, (ii) the distributions of S_n and D_n , and (iii) the shot shape $X_n(\cdot)$.

3. SHOT MODELING

The major contribution of our model is the introduction of shots to model flow rates. This is a new component we add to existing traffic models. The results of our model strongly depends on the shot shape we consider. The shot shape depends on the dynamics of flow rate, which in turn depends on many factors such as the definition of flows, the transport mechanism, the application nature, etc. We identify two methods for modeling shots, that we summarize in the next two sections. Note that our model allows the simultaneous use of different shot shapes in order to solve the problem of flows having different dynamics.

3.1 Protocol-agnostic shots

The first method to model shots is to take some *a priori* shapes and to study the performance of the model for each one of them. The shape giving the best performance is considered as the appropriate one. For example, one can consider the family of shots of the form $X_n(u) = au^b$, and choose the coefficients a and b that minimize the deviation between the measured total rate and the modeled total rate. This family of shots is illustrated in Figure 2. Note that a and b are linked by $\int_0^{D_n} X_n(u) du = S_n$. The form of the shot given by this method depends on the performance measure of the backbone traffic we want to capture. The shot shape capturing the variance of the traffic might be different than the shot shape capturing its marginal distribution or its auto-correlation function.

We compare in Figure 3 the coefficient of variation of the total rate ($\sqrt{\mathbb{E}[R^2(t)] - \mathbb{E}[R(t)]^2} / \mathbb{E}[R(t)]$) given by our model




Figure 3: Coefficient of variation of the total rate

to that given by real measurements carried out on the backbone of Sprint. Each point in the figure represents a 30 minutes trace, and flows are defined based on the well known 5-tuple. The results of the model are obtained under the assumptions that flow arrivals form a Poisson process, and that shots are triangles. In Figure 3, the closer the point to the diagonal, the better the agreement between the model and the actual trace.

3.2 Protocol-based shots

In some important cases, we can make use of protocol information to derive the shape of shots, that is not based on measurements as in the previous method. The most typical example is TCP, whose dynamics shapes the flows and can be captured by analytical models (see [1] for an example of a model for long-lived TCP flows). The first advantage of this method is that it allows the simultaneous use of different shot shapes for flows having different dynamics. The second advantage is that it provides shots able to capture all moments of the traffic, not only its variance or its auto-correlation function. The only problem with this method is in the difficulty to model flows that do not have a well defined dynamics (e.g., uncontrolled UDP flows, flows defined based on address prefixes).

Using a fluid model inspired from [1], we model the shot of a TCP flow of size S_n and of duration D_n . We find that the shot shape is only a function of (i) the moment of the total rate to capture, and (ii) the distribution law of times between packet losses. We also find that for long-lived TCP flows, and for packets lost according to a Poisson process, the average and the variance of the backbone traffic are the same as those obtained when using triangles ($b = 1$) to model shots. More general, we find that the higher the variability of times between packet losses, the larger the power coefficient b , as one expects.

4. 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] D. Daley and D. Vere-Jones, "An introduction to the theory of point processes", *Springer*, 1988.
- [4] C. Fraleigh, S. Moon, C. Diot, B. Lyles and F. Tobagi, "Packet-Level Traffic Measurements from a Tier-1 IP Backbone", *Sprint ATL Technical Report TR01-ATL-110101*, November 2001.