A flow-based model for Internet backbone traffic

Chadi Barakat, Patrick Thiran

Gianluca Iannaccone, Christophe Diot, Philippe Owezarski

ACM SIGCOMM Internet Measurement Workshop
Marseille, 6th November 2002

Introduction

- Objective: use information on flows to characterize the traffic in an IP backbone network.
  - Average and variance of the traffic.
  - Correlation.
  - If possible, the distribution.
- Backbone networks: links are over-provisioned so that the utilization does not exceed 50%.
- In the literature, the focus has been mainly on the modeling of congested links.

Characteristics of the model

- Simple: few parameters easy to compute by a router.

- Protocol and application independent: the model works for any definition of flow.

- What is a flow? A stream of packets, having an arrival time, a size (a volume), and a duration.

Without loss of generality, we use two definitions of flow:

- Packets having the same 5-tuple (e.g., TCP connections).
- Packets having the same /24 destination address prefix (e.g., packets sent to machines located on the same LAN).

Why to model?

- Impact of a change in the characteristics of flows on backbone traffic, and hence on the dimensioning of the links of the backbone:
  - Arrival rate of flows.
  - The distributions of flow sizes and flow durations:
    - ➘ New application, new network, increase in access bandwidth.
  - The correlation of flow sizes and flow durations:
    - ➘ Increase in the multiplexing level.
  - The dynamics of flows' transmission rates:
    - ➘ New transport protocol.
- Other applications:
  - Prediction of backbone traffic traffic, generation of the traffic.
  - Compute traffic in the backbone using measurements at the edges. Use this computation to optimize routing tables.

Outline

- ❑ Our fluid model for the traffic (Poisson shot-noise).
- ❑ Computation of the moments of the traffic.
  - Our focus here is on the first two moments: mean and variance.
- ❑ Validation of the model using traces collected on the Sprint IP backbone (OC-12 links $\Leftrightarrow$ 622 Mbps).
- ❑ Some results on the case of TCP flows.

The model

- Flows arrive at times $\{T_n\}$, have sizes $\{S_n\}$, and last for $\{D_n\}$
- X(t): **shot**, models flow transmission rate

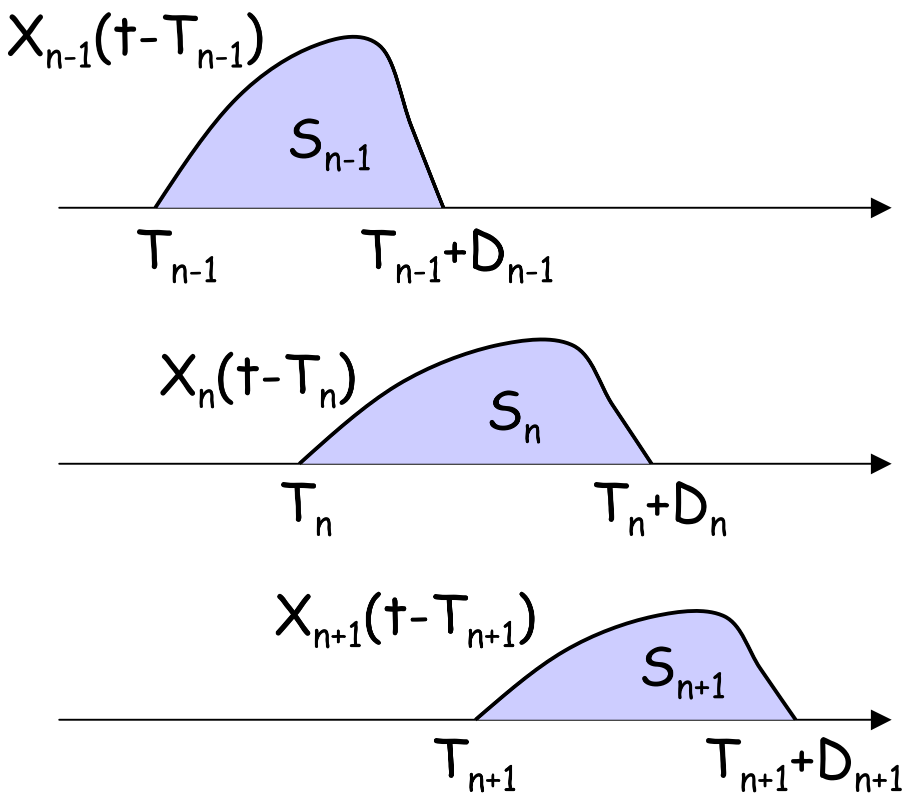

After superposition

Total rate to characterize

$$R(t) = \sum_{n=0}^{\infty} X_n(t - T_n)$$

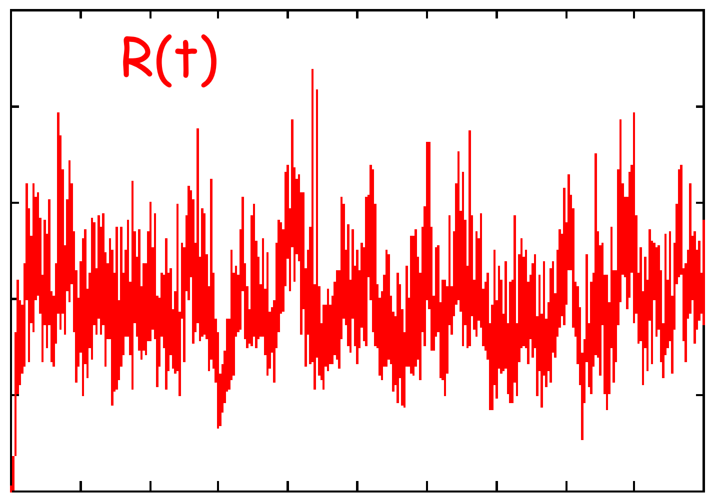

First moment of R(t)

- Easy to compute without any particular assumption.
- Let λ be the arrival rate of flows (λ finite):

$$E[R(t)] = \lim_{t \to \infty} \frac{1}{t} \int_0^t R(u)du = \lambda E[S_n]$$

Intuitive result:

Average traffic = Average arrival rate $\times$ Average flow size

- Result independent of the shot X(t).
- This is not the case for the higher moments of R(t) ...

Higher moments of R(t)

- Further assumptions are needed, justified by the high degree of multiplexing in the backbone:
  - Flow arrivals form a homogeneous Poisson process (λ).
  - Shots are iid ($\Rightarrow$ $\{S_n\}$ and $\{D_n\}$ are iid sequences).

- Relaxing the assumptions:
  - Using the theory of reversible systems, it is very easy to prove that our results on the moments of the traffic hold for much more general processes than Poisson.
  - The case where shots have different distributions can be solved by using classes of flows (for simplicity, we consider here only one class).

Laplace Stieltjes Transform

- Objective: Compute the Laplace Stieltjes Transform of the total rate R(t), i.e. $E\left[e^{-sR(t)}\right]$, Re(s) > 0.

- In the particular case $X_n(t) = 1_{\{0 < t < Dn\}}$, R(t) is the number of clients in an M/G/∞, of which we know the distribution and the LST:

$$P\{N(t) = k\} = \frac{(\lambda E[D_n])^k}{k!} e^{-\lambda E[D_n]}$$

$$N^*(z) = E\left[z^{N(t)}\right] = e^{\lambda E[D_n](z-1)}$$

Main analytical results

- The LST of R(t) in the stationary regime

$$E\left[e^{-sR(t)}\right] = \exp\left(\lambda E\left[\int_0^{D_n} e^{-sX_n(u)}du\right] - \lambda E[D_n]\right)$$

- The average of R(t) is clearly $\lambda E[S_n]$

- The variance of the total rate

$$V_R = \lambda E\left[\int_0^{D_n} X_n^2(u)du\right]$$

Variance and higher moments require models for the shot

Modeling shots

- **Measurement-based approach:**
  - Choose some parametric model for the shot shape, for example the power b shot: $X(t) = a\,t^{b}$, $b > 0$.
  - Compute the total rate using this shape, e.g. the variance.
  - Use measurements to find the optimal parameters of the shot, e.g. parameters that preserve the variance of the traffic.
  - Clearly, the shot shape depends on the measure of the traffic we want to capture, e.g. variance, higher moments, correlation.

- **Protocol-based approach:**
  - Use protocol information to determine the shape of the shot, e.g., case of TCP flows.

Measurement-based approach

- We look for shots of the form **$X(t)= at^b$**.
- **Main result:** **$V_R = k(b)\, \lambda\, E[S_n^2/D_n]$**, $k(b)=(b+1)^2/(2b+1)$.

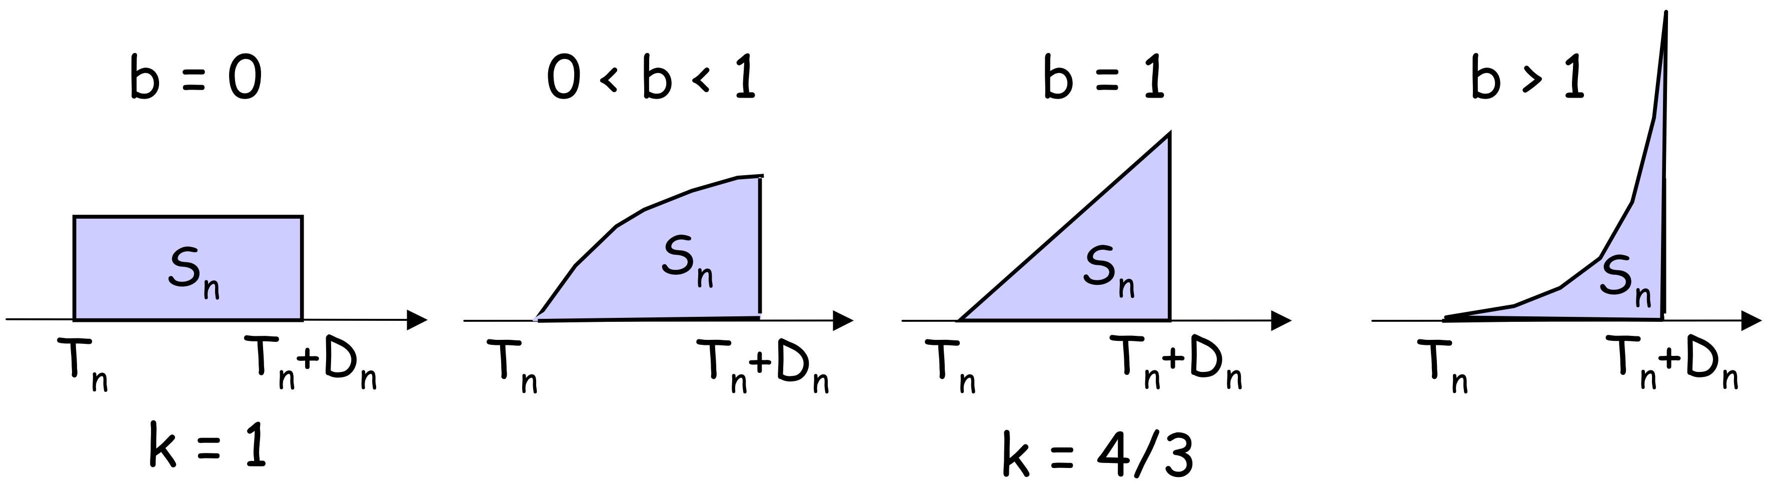

- Find the power b that preserves the variance:

$$\textbf{Measured variance} = k(b)\, \lambda\, E[S_n^2/D_n]$$

Measurement results

- Samples of R(t) obtained by averaging traffic over 200 ms.
- Each point in the figures represents 30 minutes of data.

Coefficient of variation = $\sqrt{V_R}/E[R]$

Flows defined by 5-tuple

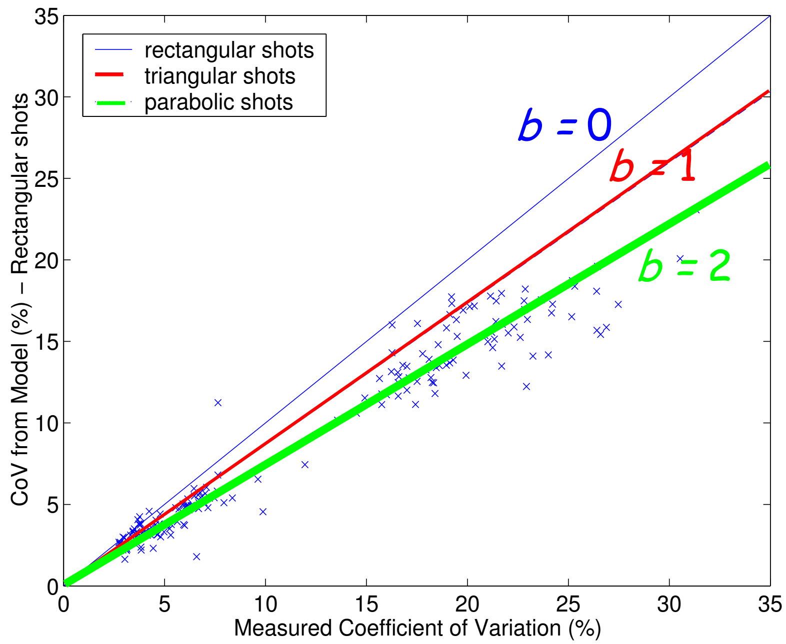

Flows defined by prefix

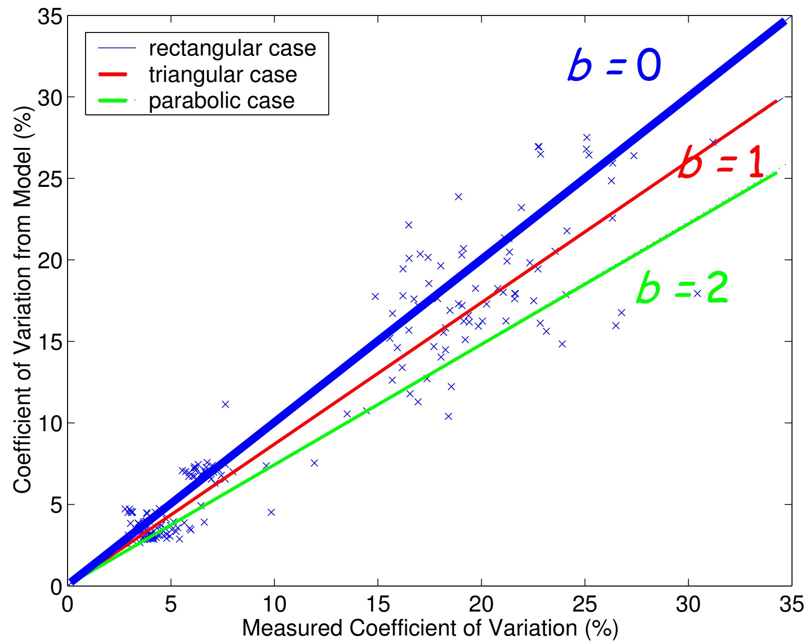

Histogram of power b

❑ For 5-tuple flows, we plot the histogram of the optimal power b over all the 30-minutes traces.

Optimal b: the shot power that gives the same variance for R(t) as the measured one.

Parabolic shots are more suited to this kind of flows.

Rectangular shots are more suited to /24 flows.

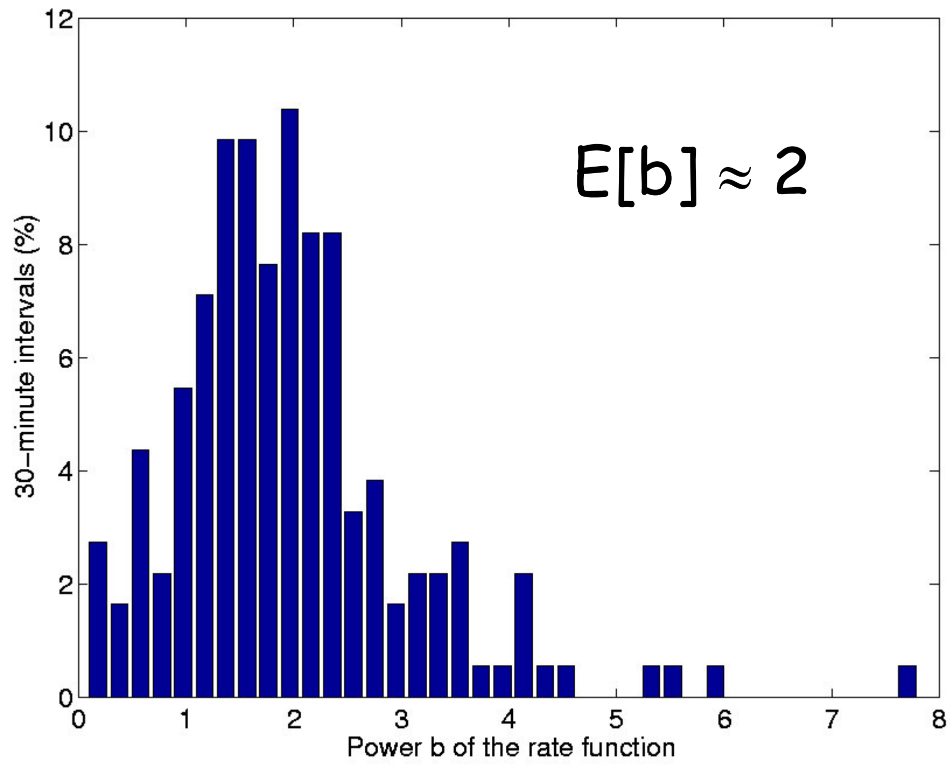

TCP-dependent shot

- We look for a shot shape that preserves the variance of TCP traffic (TCP flows carry most of Internet traffic).

TCP shot

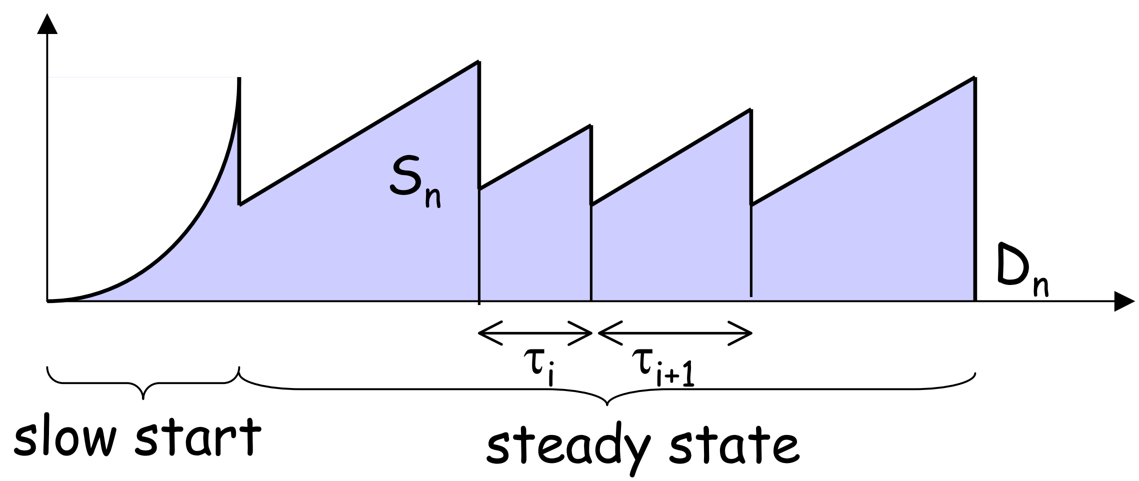

- Times between loss events τ_i are assumed to be i.i.d.
- Let $L^{(k)} = E[\tau_i^k]/E^k[\tau_i]$ be the k-th normalized moment of τ_i

Selected results

- For long-lived TCP flows (using a fluid AIMD model):

$$V_R = k(L)\, \lambda\, E[S_n^2 / D_n]$$

where $k(L) = (2 + 4L^{(2)} + L^{(3)}) / (3 + 1.5L^{(2)})$.

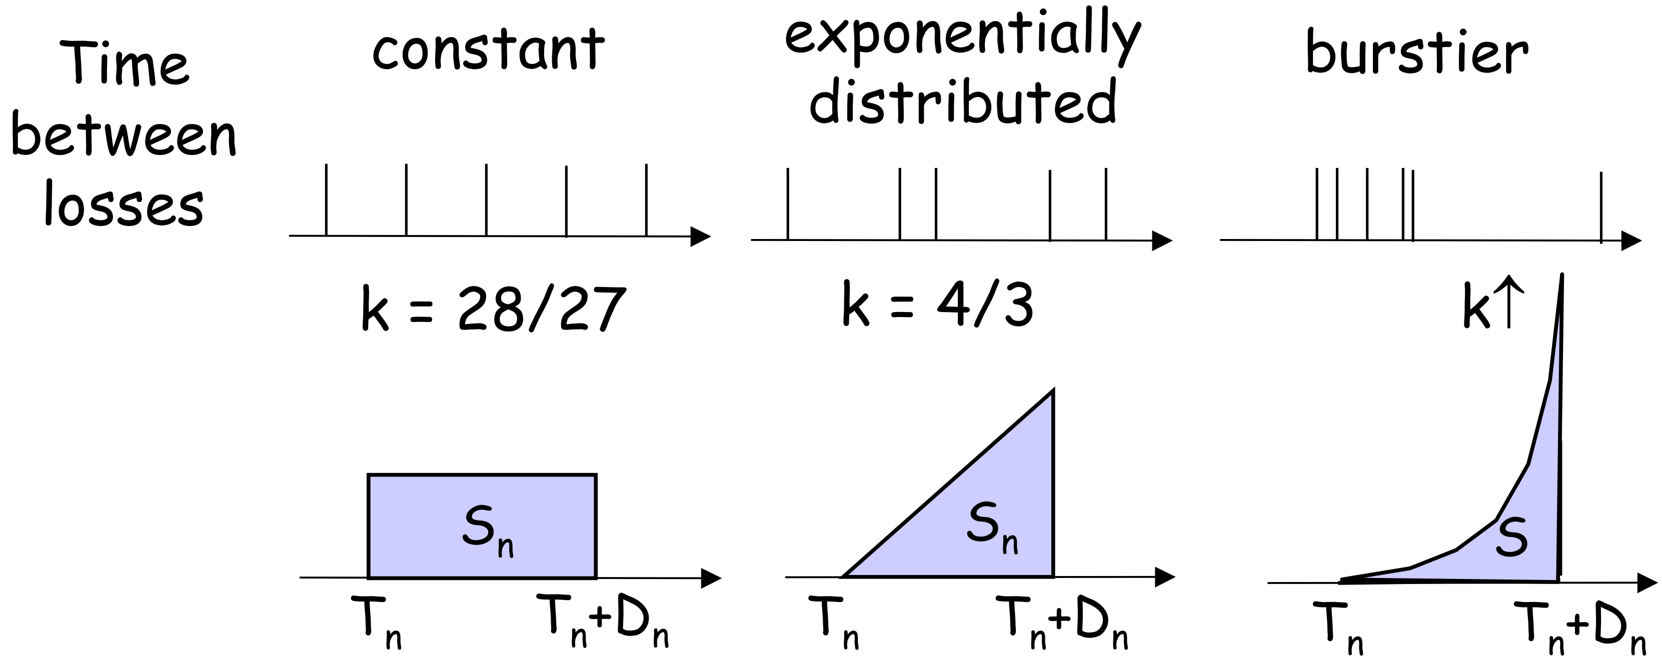

Impact of flows' characteristics on backbone dimensioning

Suppose that the links of the backbone are dimensioned to:
$E[R] + A(\varepsilon)\sqrt{V_R}$, with $A(\varepsilon)$ the congestion probability:

- Among the set of shots, rectangular shots give the lowest variance (require the least bandwidth).
- The traffic in the backbone smoothes when the arrival rate of flows increases:
  - $\sqrt{V_R}$ increases as $\sqrt{\lambda}$ instead of λ, whereas E[R] increases as λ.
- The smaller the correlation of the flow sizes and durations, the larger the variability of traffic in the backbone (the more the bandwidth required to absorb the oscillations of R(t)).

General conclusions

- A parsimonious model for backbone traffic. Three parameters are enough to compute average and variance:
  - Arrival rate of flows λ, $E[S_n]$, and $E[S_n^2 / D_n]$.

- A general model independent of the definition of flow. Flow can be defined based on 5-tuple, /24 prefix, etc.

- A new component called "shot", that allows to specify the model to different applications and transport protocols.