Simulation Analysis and Fixed Point Approach for Multiplexed TCP flows

Urtzi Ayesta, Konstantin Avrachenkov,
Eitan Altman, Chadi Barakat, Parijat Dube

INRIA-Sophia Antipolis, France

Motivation

- Propose a simple model for an Internet router crossed by TCP traffic:
  - TCP carries the majority of Internet traffic.
  - The router implements the DropTail policy.

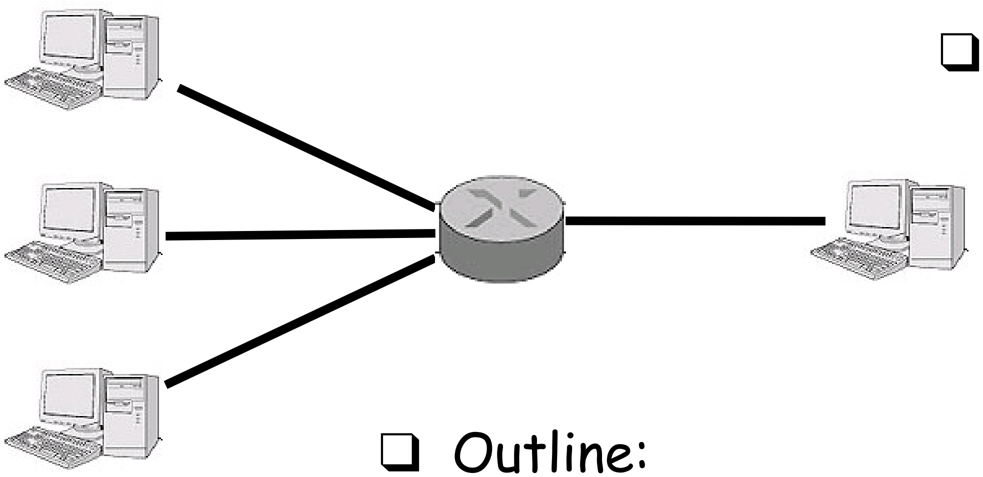

- Performance measures to compute:
  - Loss probability in the router.
  - Utilization of the bandwidth.
  - Useful for network dimensioning !
- Outline:
  - A Fixed Point Approach based on the M/M/1/K model.
  - We validate the model with NS and we study the cases when it applies.

Model for TCP traffic

- Short-lived TCP transfers that arrive at the router at a rate λ and have an average size E[S]:

$$\text{Nominal load} \quad \rho_0 = \frac{\lambda.E[S]}{C_{Bottleneck}} = \frac{\sum_{i=1}^{N} \lambda_i.E[S]}{C_{Bottleneck}}$$

- "M" long-lived persistent TCP connections.
  - Suppose that the bottleneck drops packets with probability p.
  - Let T(p) be the average sending rate of a persistent TCP connection.

$$\text{Total load} \quad \rho = \frac{\rho_0}{1-p} + \frac{M.T(p)}{C_{Bottleneck}}$$

Fixed Point Approach

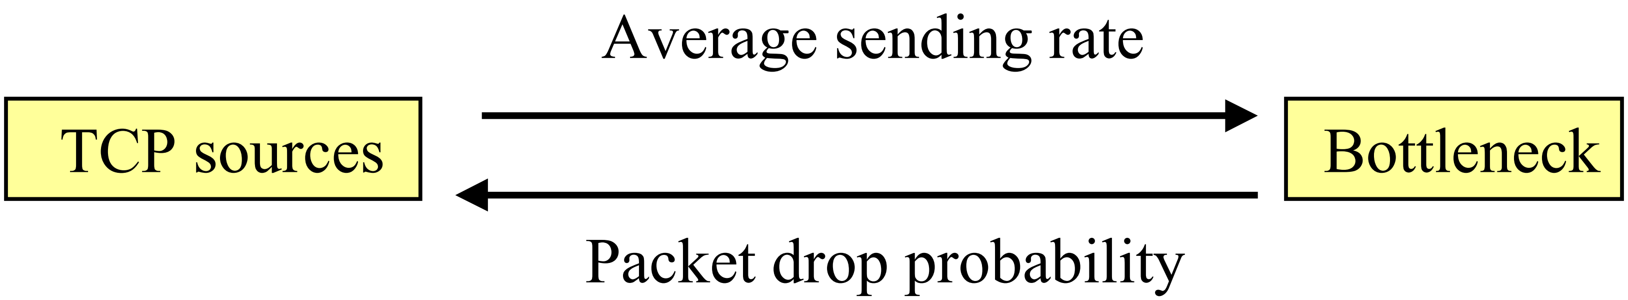

- We have an expression of the load ρ as a function of p.
- If we find a model for the network that gives p as a function of the load ρ, the problem can be solved for:
  - The packet loss probability, the load, the average queue size.
  - These measures can then be used to compute performance measures at the application layer e.g. HTTP latency.
- A reference model for the network is the M/M/1/K …

Fixed Point Approach based on the reference M/M/1/K model

- Packets arrive as Poisson.
- Packet transmission times are exponential.
- K is the size of the buffer in the bottleneck router.

$F(\rho) = \rho_0$

$$p = \frac{\rho^K (1-\rho)}{1-\rho^{K+1}} \quad \text{and} \quad \rho = \frac{\rho_0}{1-p} + \frac{M.T(p)}{C_{Bottleneck}}$$

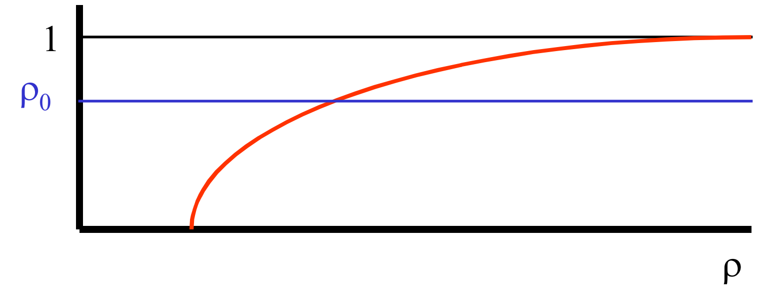

- Proposition: For $\rho_0 < 1$ and T(p) non increasing with p, the solution of the system exists and is unique (**proof given in the paper**).
- The service time distribution can be relaxed by taking an M/G/1/K model.
- But, what about the packet arrival process? and does this model scale with the capacity of the network (case of high speed links) … ?

Packet arrival process: Simulation analysis

- The packet arrival process tends to Poisson as the number of multiplexed flows increases:
  - Distribution of times between packet arrivals tends to exponential.
  - Times between packet arrivals tend to be independent.
- But what about the topology of the network, and the speeds of links ?
- We simulate with NS:
  - Prop delays between 20 ms and 60 ms.
  - Files arrive as Poisson.
  - File size exponential and Pareto.
  - One persistent TCP connection per access link.
  - Scenarios: Only short transfers, only long transfers, mixture of short and long transfers.

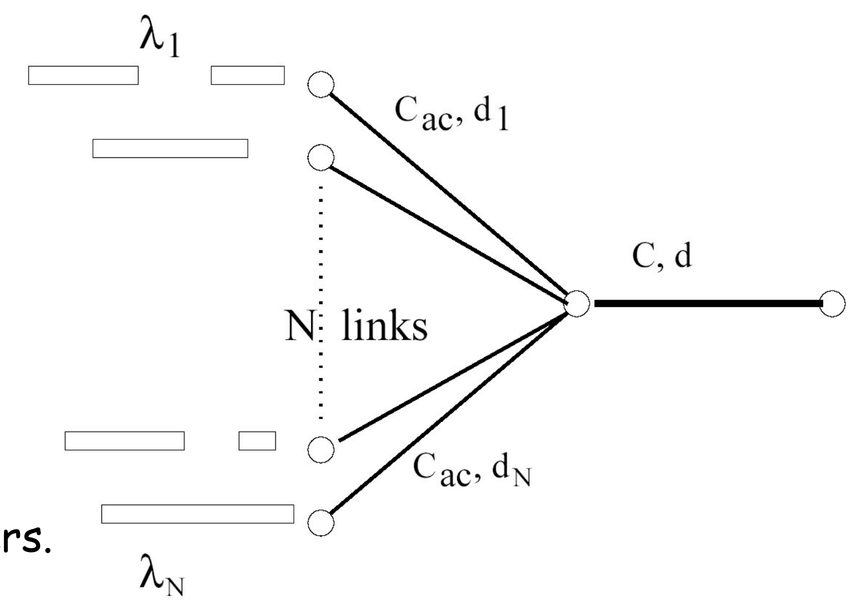

Summary of the findings

- Distribution of packet inter-arrival times:
  - Shows steps at the transmission times of access and bottleneck links, for all values of N and ρ (all access links are assumed to have the same speed).
    - If access links are faster than the bottleneck, we have two steps at $1/C_{ac}$ and $1/C$.
    - If access links are slower than the bottleneck, we have a step at $1/C_{ac}$.
  - The presence of long-lived TCP connections reduces the steps, and make the distribution closer to exponential.
  - At high speed, the step is very close to zero !
    - Packets arrive in batches, a batch model should give better results than Poisson.
  - If access links are much slower than the bottleneck, inter-arrival time distribution is very close to exponential, $C_{ac} \approx \frac{2.C}{N}$

INRIA SOPHIA ANTIPOLIS

Some results about the distribution

Only short TCP transfers, Pareto file size of average 10 Kbytes

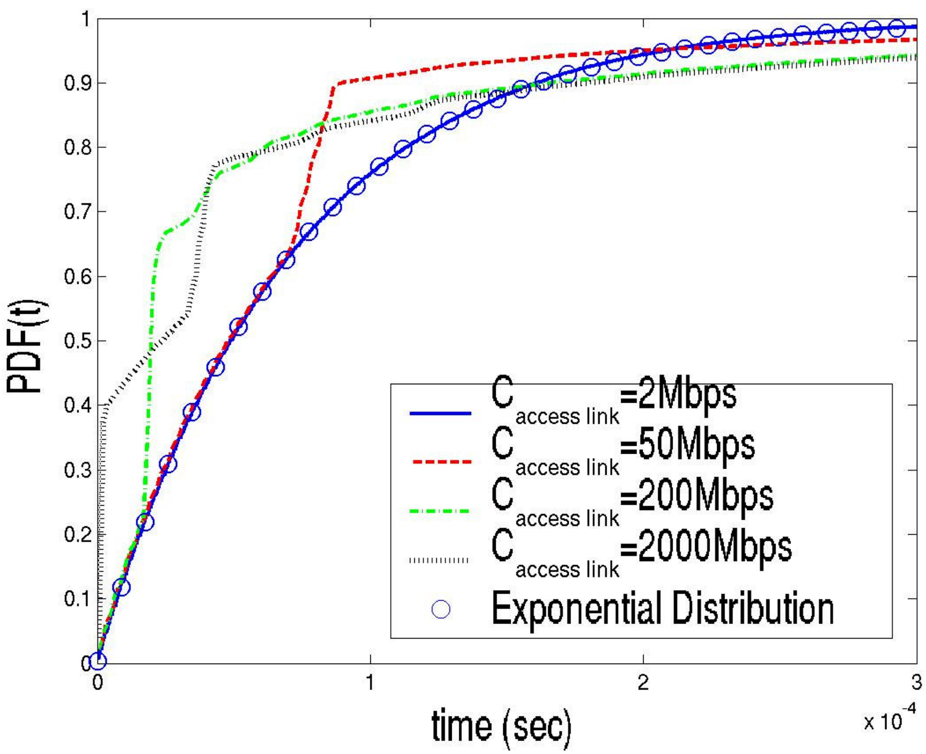

Bottleneck 100 Mbps,
100 access links, load 0.6

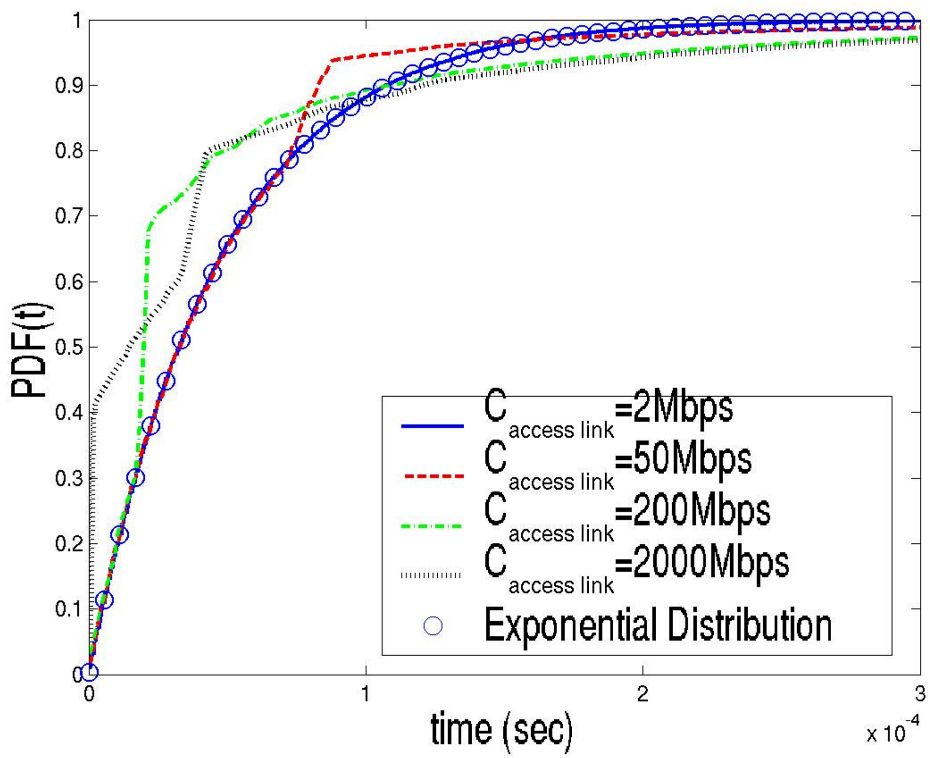

Bottleneck 100 Mbps,
100 access links, load 0.9

Some results about the distribution

Persistent and short connections, 80% of load due to short connections, pareto file size of average 10 Kbytes for short connections

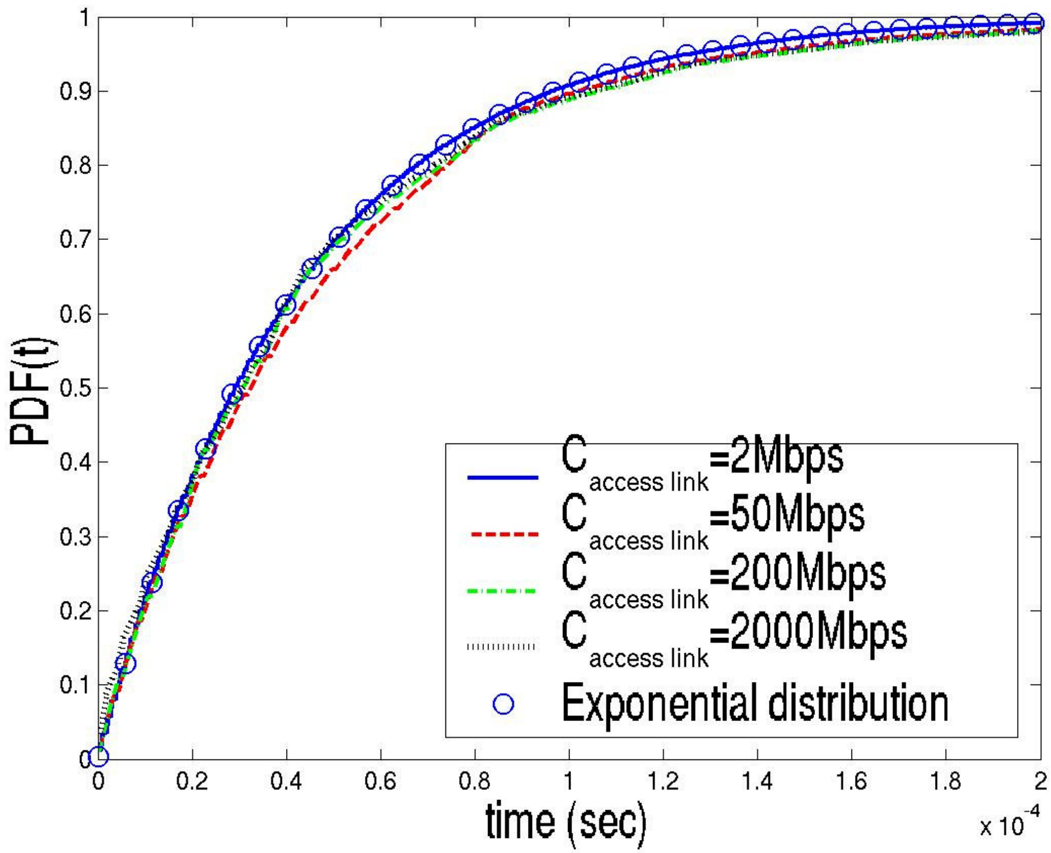

Bottleneck 100 Mbps, 100 access links

Auto-correlation of packet inter-arrival times

- The number of lags that corresponds to a maximum auto-correlation is approximately:

$$\frac{\text{transmission time of a packet over an access link}}{\text{average time between packet arrivals at the bottleneck}} \approx \frac{C_{bottleneck}}{C_{access}}$$

  - Packets arrive in bursts over the access link.
- For slow access links, the correlation of the traffic is less important:
  - *And we have seen that for slow access links, the distribution is close to exponential.*
  - So, the packet arrival process tends to Poisson as the speed of the bottleneck increases with respect to the access
    - Packets are spaced at the access before arriving at the backbone !

Auto-correlation of packet inter-arrival times

Only short TCP transfers of average 10 Kbytes

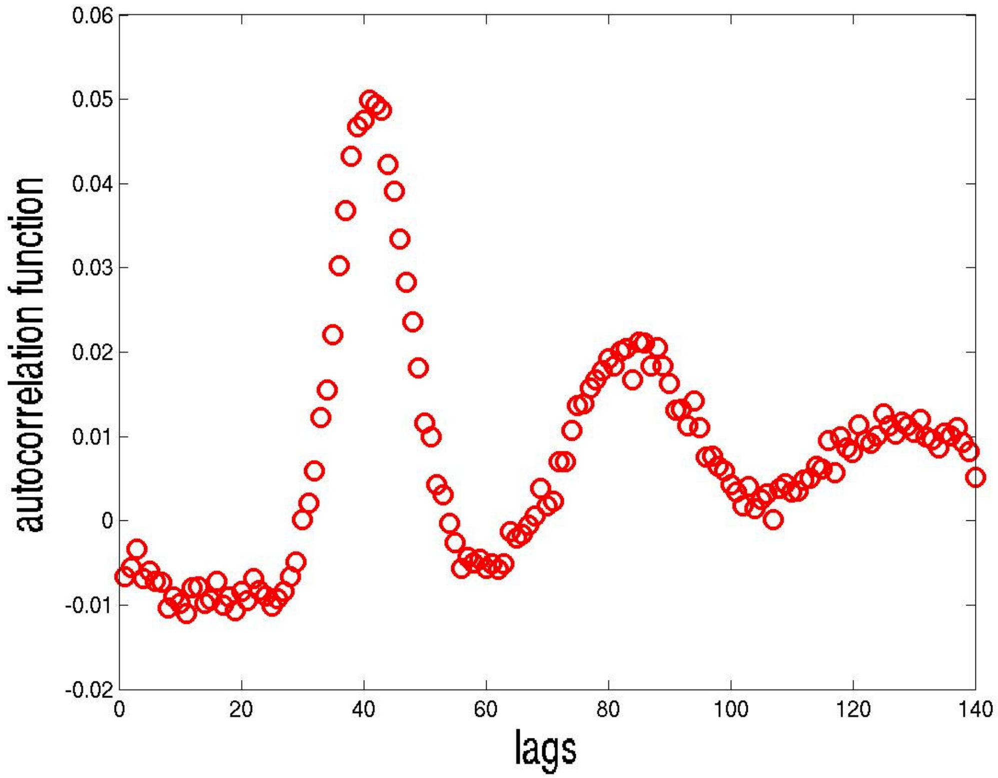

Correlation function of the packet arrival process:
Bottleneck 100 Mbps, 100 access links and C_{ac} = 2Mbps

Auto-correlation of packet inter-arrival times

Only persistent connections

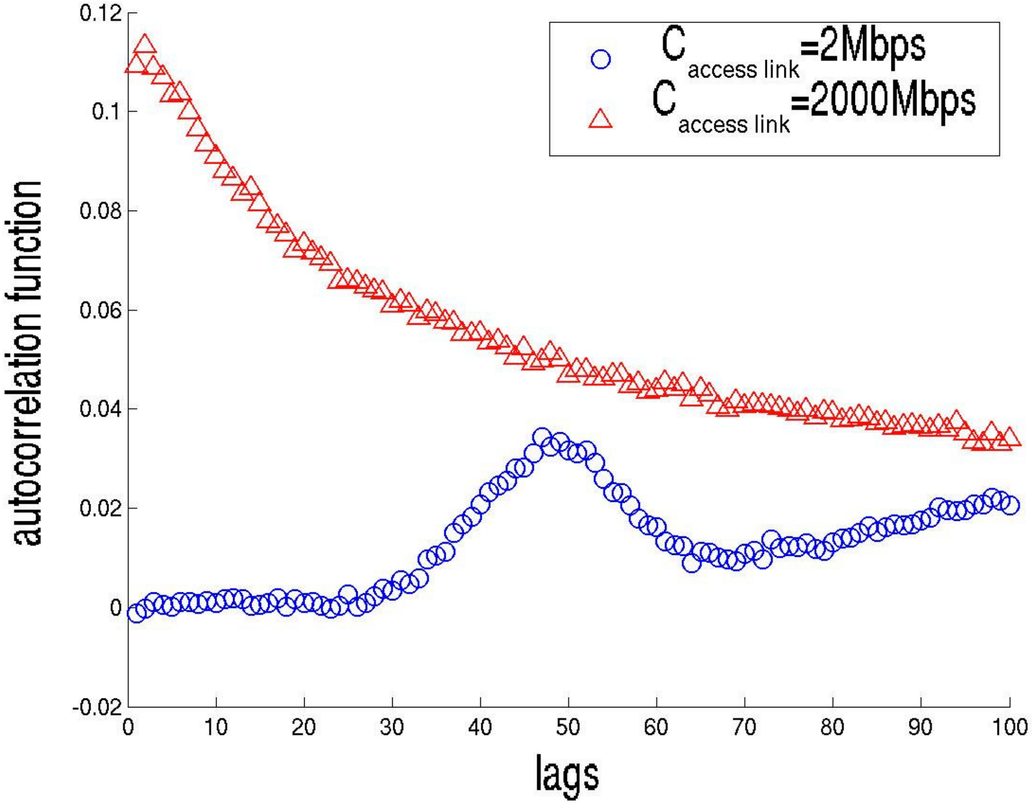

Correlation function of the packet arrival process:
Bottleneck 100 Mbps, 100 access links

A batch model for high speed links

- Our approximation for the most bursty case (high speed links):
  - A batch size equal to the window size W.
  - The TCP connection does not suffer from any loss (stays in slow start).
  - Batches arrive as Poisson at the bottleneck.
  - The probability of having a window size $W = w_i \in \{1,2,4,8 \ldots\}$ is

$$P(W = w_i) = \sum_{j=\lfloor \log(w_i)+1 \rfloor}^{\infty} \frac{1}{j} P(2^{j-1} - 1 < FileSize \le 2^j - 1)$$

  - The queue length is a Markov process, π_i its stationary distribution. The packet loss probability is given by:

$$p = \frac{1}{E[W]} \sum_{i=0}^{K} \pi_i \sum_{w_j = K-i+1}^{\infty} (w_j - K + i) P(W = w_j)$$

Analytical vs. Simulation results

Only short TCP connections, exponential file size of average 10 Kbytes

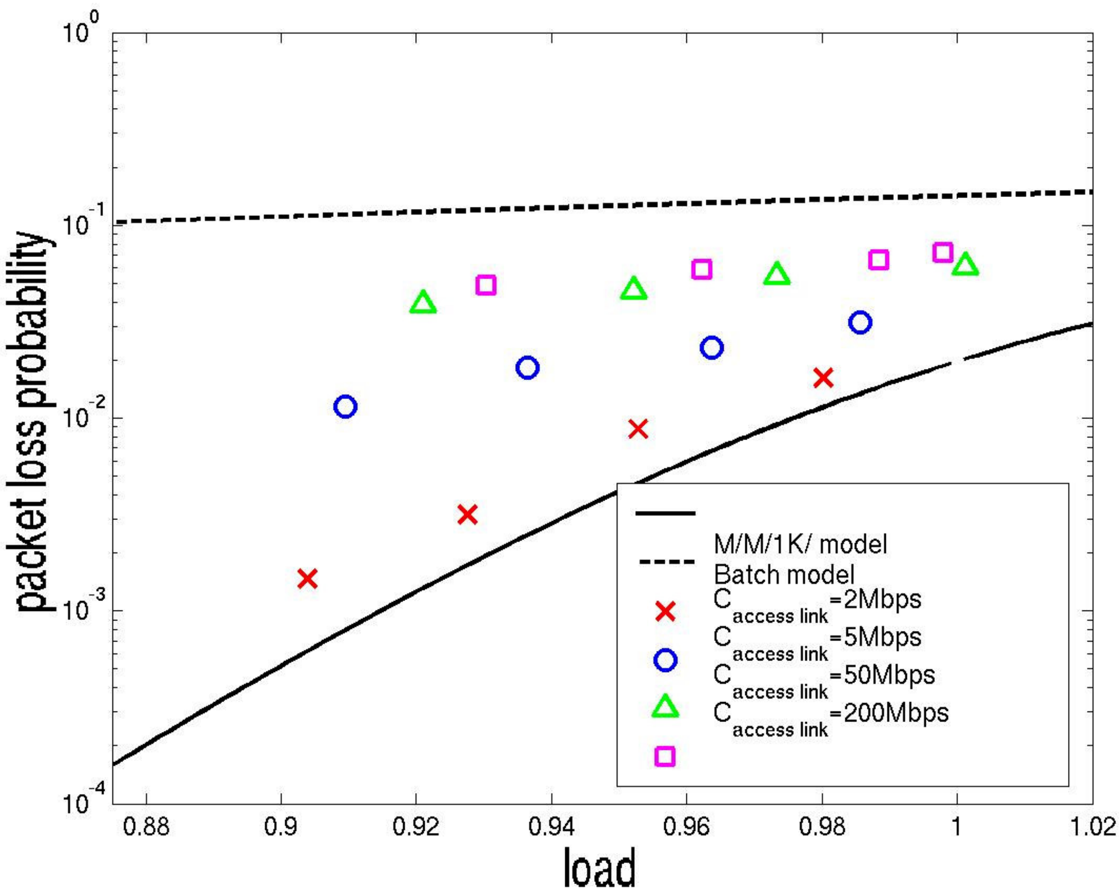

Packet loss probability: Bottleneck 100 Mbps, 100 access links

INRIA SOPHIA ANTIPOLIS

Conclusions

- A simple model for Internet routers.
  - Can be used to compute performance at transport and application layer.
- Simulation analysis of the arrivals of TCP packets to a router:
  - Bursty, with the times between packets in a burst function of links' speed.
  - The maximum correlation appears at a number of lags function of the ratio of link speeds (bottleneck to access).
  - Approaches Poisson when the access links are slow compared to the bottleneck (traffic spread at the access before arriving at the bottleneck).
  - A Poisson batch model to be used at high link speed.