Performance Evaluation of Networks: New Problems and Challenges

Alain Jean-Marie
LIRMM, University of Montpellier
161 Rue Ada, 34392 Montpellier Cedex 5, France
ajm@lirmm.fr

ASIAN'00
Penang, 26 November 2000

Performance Evaluation

Introduction Motivation and approach

- Networks, contention, delays and losses
- Methodology
- Models

The Theory of Queues Classical results

- Markov Chains
- Performance measures

Current Challenges Some models with or without solution

- Internet and its evolution
- Web modeling
- Models of Multimedia applications
- From passive to active networking

Introduction

In a communication network using routing/switching (Internet, ATM, Frame Relay...), *queues* form along the communication path (statistical multiplexing, contention, ...)

These queues create *delay* and *losses*.

The problem is to know how to quantify these.

The approach is (usually) *stochastic*, given the uncertain nature of traffic.

Research for results permitting to define, calculate and guarantee the celebrated *quality of service* (QoS).

Methodology

How to obtain performance measures?

Real System: Define objectives

Instrument the system: place control points, place measurement points

Perform measurements

Change parameters

Do it again

Simulated System: Define objectives

Program a sufficient representation of the system, elements and behavior

Perform measurements

Change parameters

Do it again

Mathematical analysis: Define objectives

Establish a sufficient mathematical representation of the system, elements and behavior

Calculate measures: QoS $= f(x_1, \ldots, x_n)$.

For both Simulation and Analysis, one needs Models.

Models

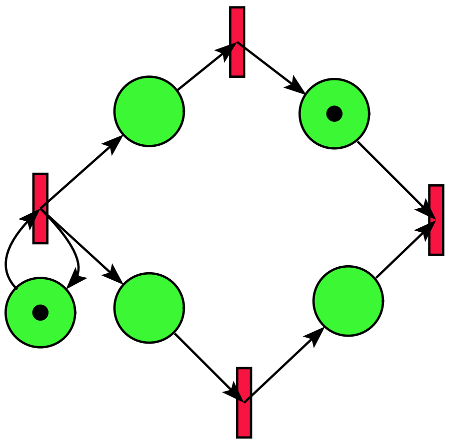

Petri Nets

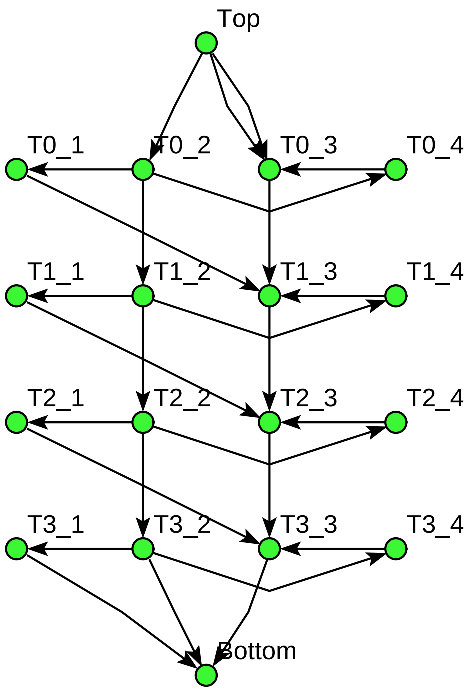

Task Graphs
(PERT networks, etc.)

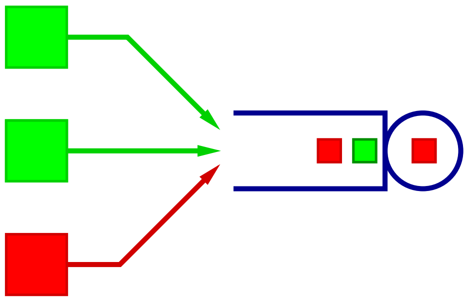

Queuing Networks

Uncertainty and randomness

Unknown quantities: arrival times of “events”, amount of resources claimed on the system.

- Stochastic models: unknown quantities are random variables.

  Random in, Random out $\Rightarrow$ performance measures are random in nature

  $\Rightarrow$ compute or measure their statistics (mean, variance, distribution...).

- Deterministic models: unknown quantities have bounds.

  Analysis reveals the worst case scenarios $\Rightarrow$ guaranteed performance.

Erlang's Model (1917)

A telephone cable with with N lines:

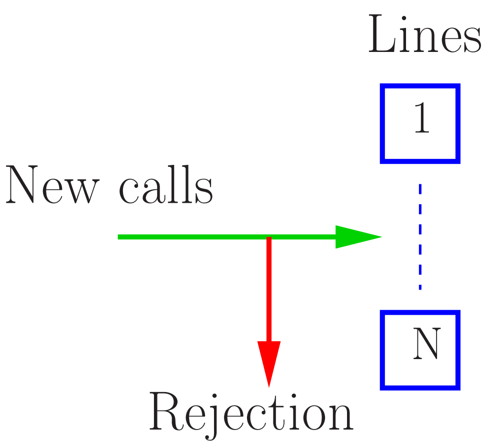

Call rate $= \lambda/s$

Average call duration $= m\ s$

$$\mathbb{P}\{\text{loss}\} = \frac{\rho^N/N!}{1+\rho+\ldots+\rho^N/N!} \qquad \rho = \lambda \times m$$

Penang's Model (2000)

A Web server with persistent HTTP connections and finite memory:

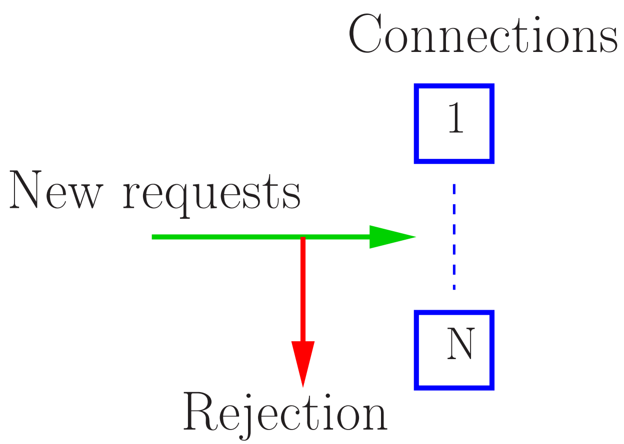

Hit rate $= \lambda/\mathrm{mn}$

Average connection length $= m$ mn

$$\mathbb{P}\{\text{no service}\} = \frac{\rho^N/N!}{1+\rho+\ldots+\rho^N/N!} \qquad \rho = \lambda \times m$$

Discrete Time Markov Chains

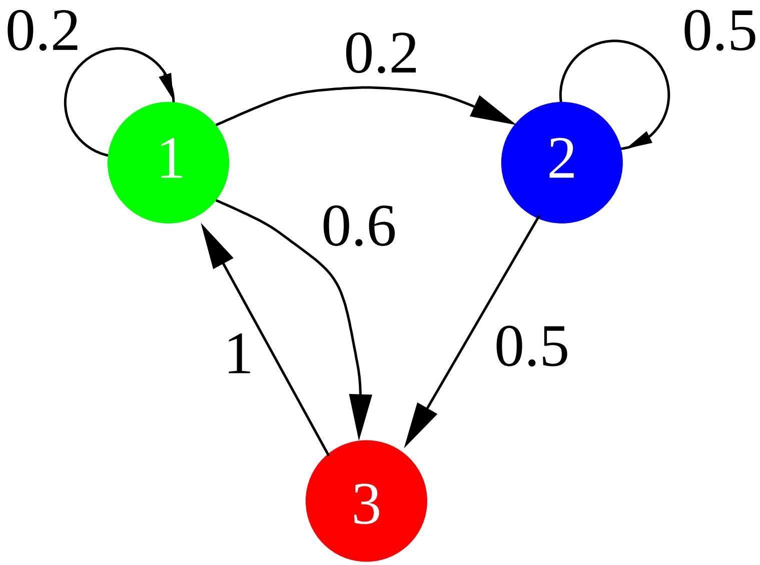

$$P = \begin{pmatrix} 0.2 & 0.2 & 0.6 \\ 0 & 0.5 & 0.5 \\ 1 & 0 & 0 \end{pmatrix}$$

Transition diagram

Transition Matrix

Probability vectors:

$$
\begin{array}{lll}
\pi_0 & = & (1, 0, 0) \\
\pi_1 & = & (0.2, \quad 0.2, \quad 0.6) \\
\pi_2 & = & (0.64, \quad 0.14, \quad 0.22) \\
\pi_3 & = & (0.348, \quad 0.198, \quad 0.454) \\
\pi_4 & = & (0.5236, \quad 0.1686, \quad 0.3078) \\
\vdots & \vdots & \vdots \\
\pi_\infty & = & (5/11, 2/11, 4/11)
\end{array}
$$

$$\boxed{\boldsymbol{\pi} \quad = \quad \boldsymbol{\pi}\, \mathsf{P}}$$

$\Rightarrow$ the *stationary probability*.

Its computation is reduced to the solution of a linear system!

Continuous Time Markov Chains

Let $\{X(t), t \in \mathbb{R}^+\}$, having the following properties. When X enters state i:

- X stays in state i a random time, with an exponential distribution with average $1/\tau_i$, independent of the past; then

- X jumps instantly in state j with probability p_{ij}.

Let

$$Q_{ij} \;=\; \tau_i p_{ij} \qquad Q_{ii} \;=\; -\tau_i \;.$$

$\mathbf{Q}$ is *infinitesimal generator*.

Example

$$\tau = \begin{pmatrix} 6 \\ 2 \\ 3 \end{pmatrix} \qquad \mathsf{P} = \begin{pmatrix} 0 & 1 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{3} & \frac{2}{3} & 0 \end{pmatrix} \qquad \mathsf{Q} = \begin{pmatrix} -6 & 6 & 0 \\ 1 & -2 & 1 \\ 1 & 2 & -3 \end{pmatrix}$$

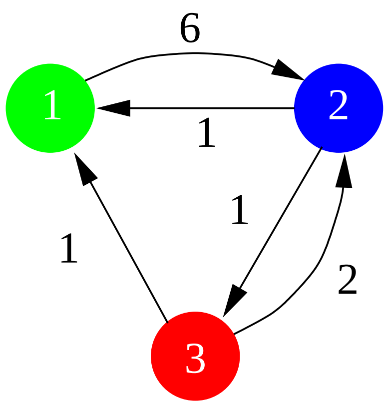

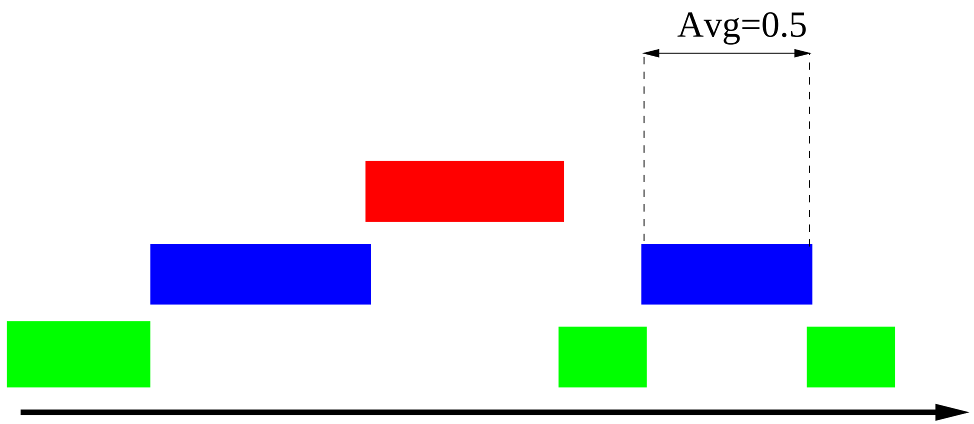

Equilibrium equations

If $\lim_t \boldsymbol{\pi}_t = \boldsymbol{\pi}$ exists, then:

$$\boxed{\mathbf{0} \;=\; \boldsymbol{\pi}\,\mathrm{Q}\,.}$$

These equilibrium equations can be written: $\forall i \in \mathcal{E}$,

$$(\sum_{j\neq i} q_{i,j})\pi(i) \;=\; \sum_{j\neq i} \pi(j) q_{j,i}\;.$$

Interpretation: entering flow = outgoing flow.

Queues

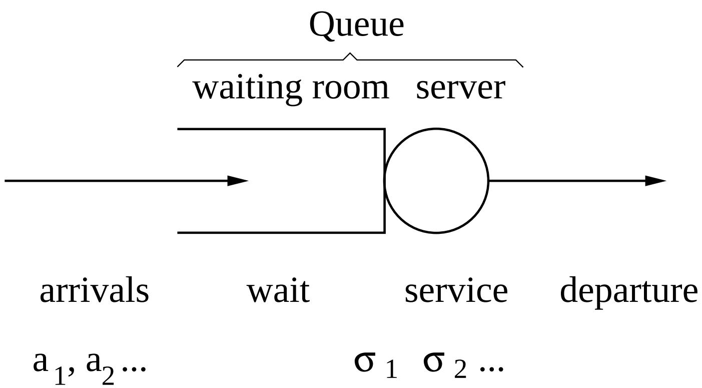

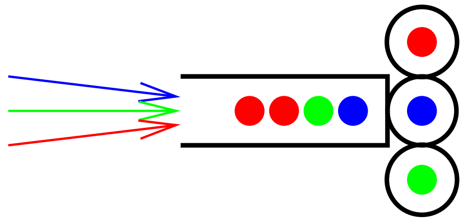

Usual representation of a queue

The elements that compose a queue are:

- one or several servers
- a waiting room
- (possibly) classes of customers
- an arrival process per class
- a process of service durations
- a service discipline

The M/M/1 queue

Characteristics: Infinite waiting room, 1 server, FIFO.

arrivals: a Poisson process with throughput λ

services: exponential distribution with average $1/\mu$:

$\{N(t)\}$ is a Markov Chain: a *birth and death process*

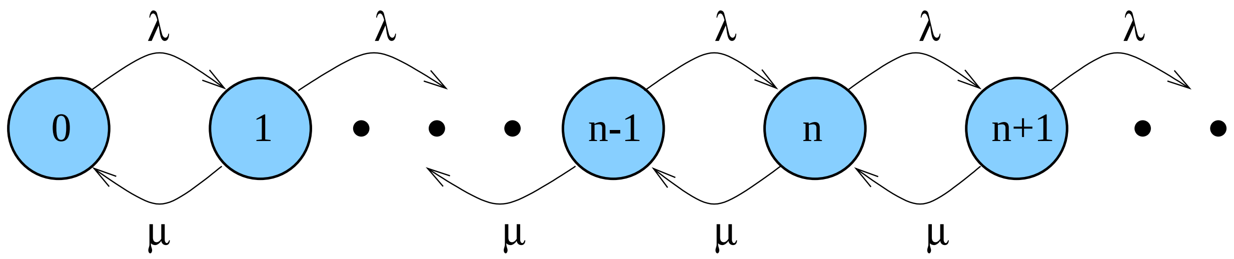

Performance Measures:

$$
\begin{aligned}
\text{Stability} \quad &\Longleftrightarrow \quad \lambda < \mu \\
\mathbb{P}\{N \geq n\} \quad &= \quad \left(\frac{\lambda}{\mu}\right)^{n} \\
\mathbb{P}\{N = n\} \quad &= \quad \left(1 - \frac{\lambda}{\mu}\right)\left(\frac{\lambda}{\mu}\right)^{n} \\
\mathbb{E}N \quad &= \quad \frac{\lambda}{\mu - \lambda} \\
\mathbb{E}R \quad &= \quad \frac{1}{\mu - \lambda}
\end{aligned}
$$

Product form solutions: Jackson and Kelly Networks

N queues (stations) with services $\sim$ Exp.

External arrivals of customers (Poisson processes), and a routing mechanism.

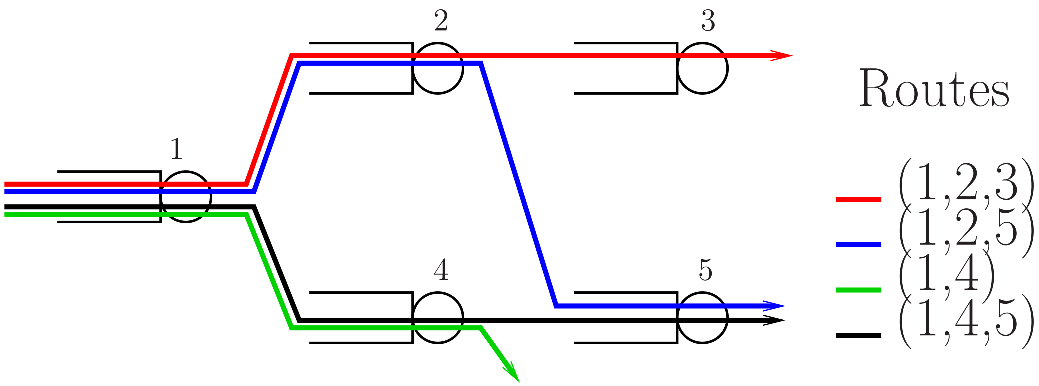

Global entering throughput in station i: λ_i

When stability, the stationary probability distribution is:

$$\boxed{p(n_1, \ldots, n_N) \;=\; \prod_{i=1}^{N} \left(1 - \frac{\lambda_i}{\mu_i}\right) \left(\frac{\lambda_i}{\mu_i}\right)^{n_i},}$$

$\Rightarrow$ justification of the end-to-end response time formula:

$$\boxed{T \;=\; \sum_{i=1}^{M} \frac{1}{C_i - L_i}}$$

C_i: capacity of the link/switch, L_i: entering traffic (in packets/s).

Application: Capacity Planning

Assuming known: traffic rates and routes.

Problem: allocate link/node capacities so as to minimize collective average.

$$\boxed{\min_{(\mu_1,\ldots,\mu_N)\in\mathcal{M}} \qquad \overline{T}(\mu_1,\ldots,\mu_N)}$$

$\mathcal{M} = \{$feasible allocations$\}$.

Application: Route Planning

Assuming known: node and link capacities, Origin/Destination traffics.

Problem: allocate routes so as to minimize collective average.

$$\boxed{\min_{\mathbf{x}\in\mathcal{R}} \qquad \overline{T}(\mathbf{x})}$$

where $\mathcal{R} = \{\text{feasible route allocations}\}$.

Routing

Consider a network with distributed routing based on distance vector tables.

According to the response time formula at nodes for Kelly networks plus the propagation delay, a reasonable metric for link $n = (i \to j)$ is:

$$\boxed{D_n \;=\; \frac{1}{\mu_n - \lambda_n} + d_n\ .}$$

Analytical approach to flow control

Consider the closed Jackson network with W customers, each node having capacity μ

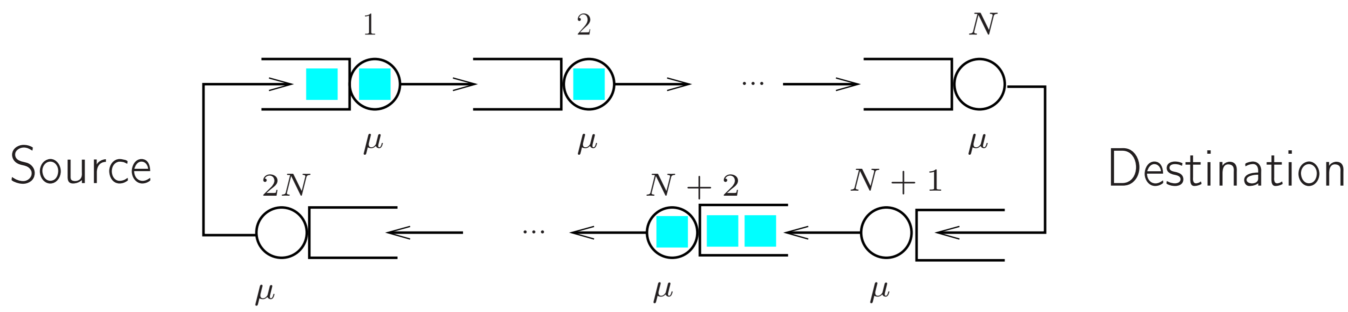

The throughput and RTT (Round Trip Time) of packets are

$$\boxed{\theta = \frac{W\mu}{W+2N-1}}$$

$$\boxed{RTT = \frac{W+2N-1}{\mu}}$$

Theoretical Challenges

Examples of open questions:

- less restrictive assumptions on traffic models:
  non-Poisson arrival processes
  non-exponential services
  long range dependence of processes

- finite capacities, losses, feedback

- service disciplines and stability

- distributions of end-to-end response times

Modeling Internet and its Evolutions

Study the protocols for:

- controlling the flow of connections: TCP
- controlling the flow of non-connection streams
- preventing the congestion
- introducing some level of "differentiated service"

Models of TCP

Principles of the TCP congestion control (simplified)

- TCP sources use a window W (bytes) of unacknowledged packets
- the window decreases when packet losses are detected
- the window grows
  - exponentially fast in the slow start or fast recovery modes
  - linearly in the congestion avoidance mode

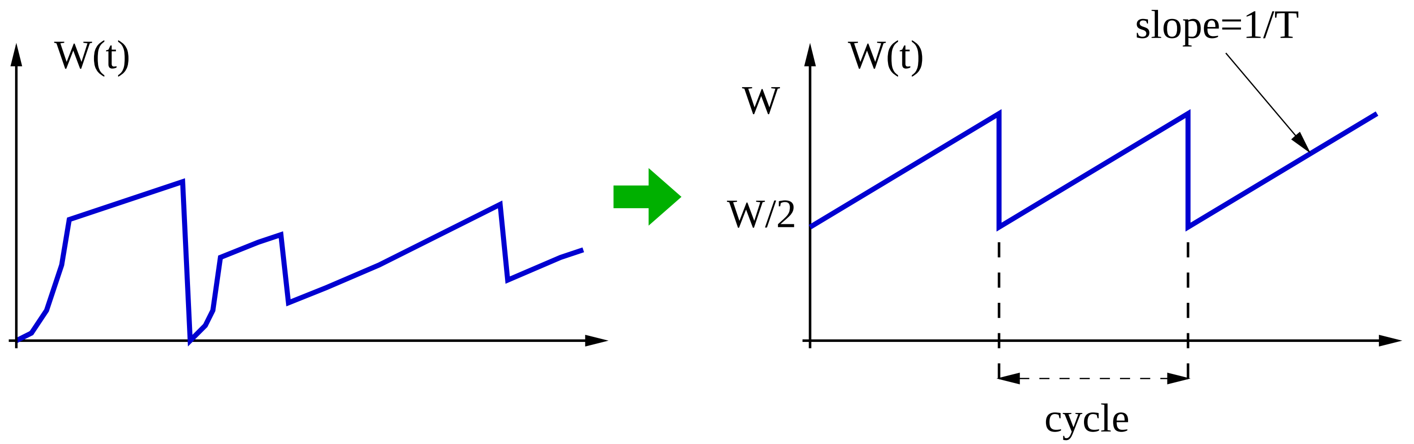

Deterministic cycle analysis: the effective throughput is:

$$\boxed{\theta \quad \simeq \quad \frac{1}{T}\sqrt{\frac{3}{2p}}}$$

This is the "square root formula": relates the loss probability p and the throughput θ.

Stochastic cycle analysis:

Times S_n between losses: i.i.d. random variables with coefficient of variation c^2.

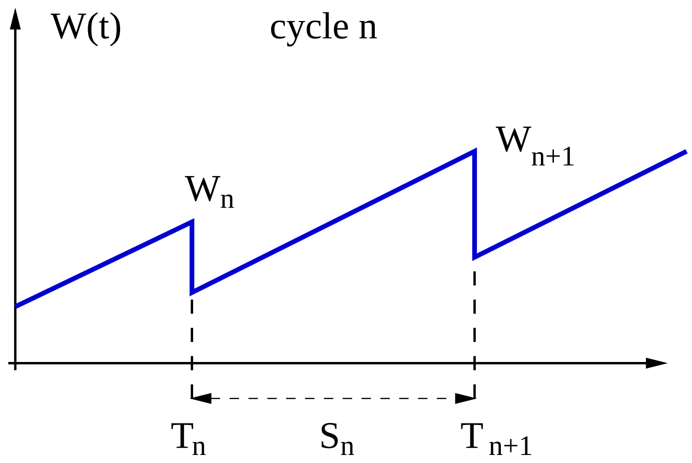

$$\boxed{\theta = \frac{1}{T}\sqrt{\frac{3+c^2}{2p}}}$$

Differentiated Services

Objective Improve the "quality of service" of the Internet by adding some kind of service definition and guarantees.

Move away from "best effort" and its lack of response time/throughput guarantees.

Means Using 1 bit, define 2 classes with one of them having:

- better throughput ("Assured Forwarding")
- better delay characteristics (Expedited Forwarding, aka: Premium Service)

Model for "Assured Forwarding"

Model:

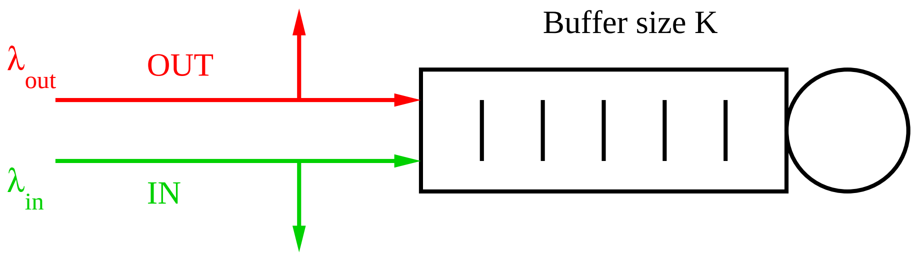

Router = single server queue

Input traffic = two classes, IN (high priority, tagged) and OUT, low priority.

Buffer management: RED = Random Early Detection / Discard

RIO = RED on IN and OUT

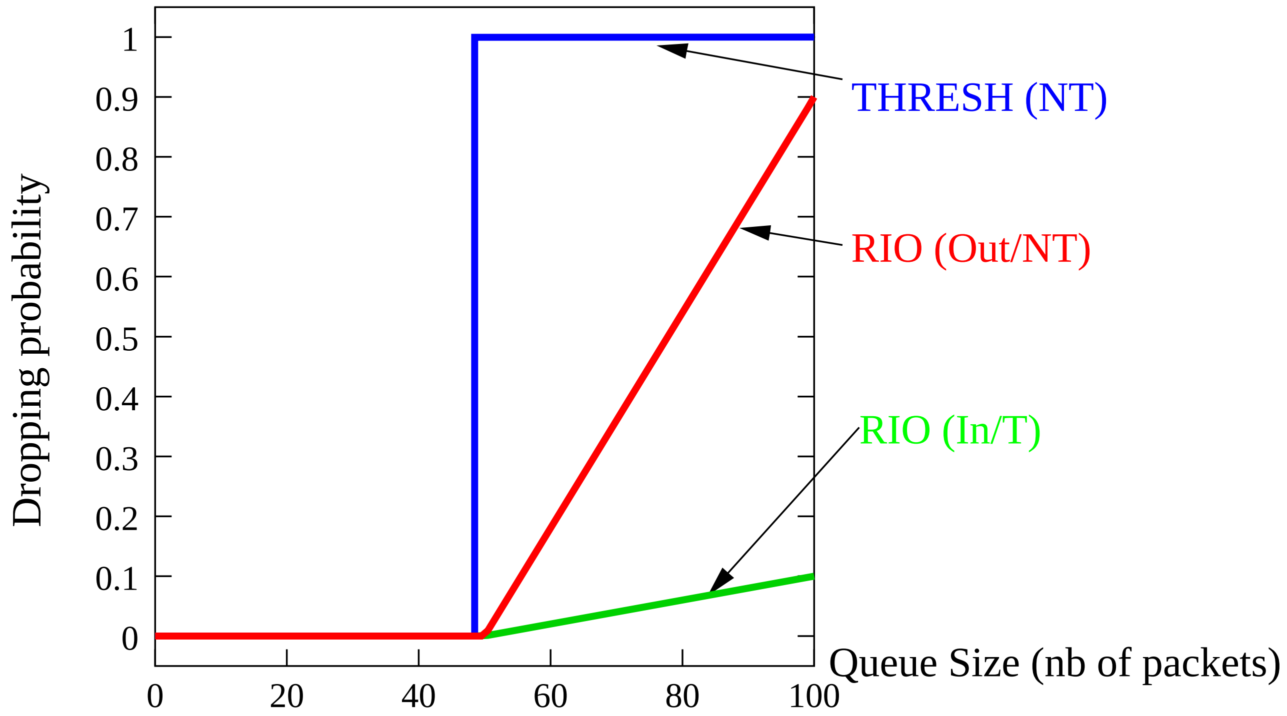

Probability of accepting a packet:

$$\boxed{\alpha(n) \quad = \quad \frac{\lambda_{\text{in}}}{\lambda}\,\alpha_{\text{in}}(n) \;+\; \frac{\lambda_{\text{out}}}{\lambda}\,\alpha_{\text{out}}(n)}$$

with $\lambda = \lambda_{\text{in}} + \lambda_{\text{out}}$.

Evolution of the number of customers $N(t)$:

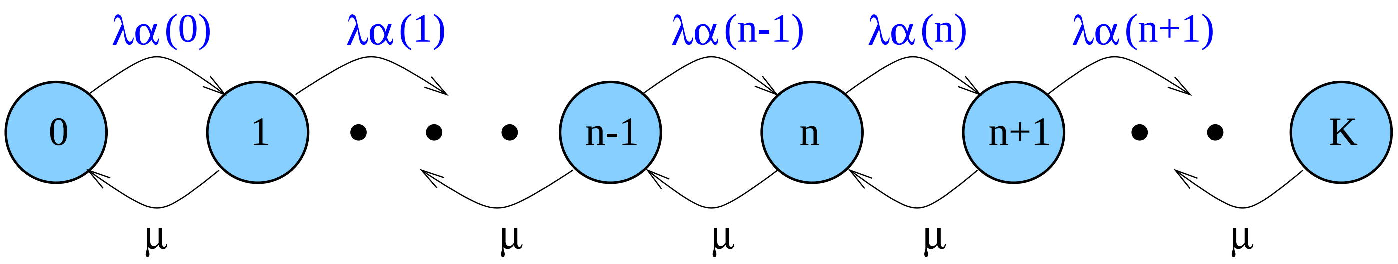

Solution:

$$
\begin{aligned}
\pi(n) &= \pi(0) \left(\frac{\lambda}{\mu}\right)^n \prod_{i=0}^{n-1} \alpha(i) \\
\pi(0) &= \left[\sum_{n=0}^{K} \left(\frac{\lambda}{\mu}\right)^n \prod_{i=0}^{n-1} \alpha(i)\right]^{-1}
\end{aligned}
$$

Computation of performance measures, including throughputs:

$$
\begin{aligned}
{\lambda_{\text{in}}}^{\text{eff}} &= \lambda_{\text{in}} \sum_{n=0}^{K-1} \alpha_{\text{in}}(n)\pi(n) \\
{\lambda_{\text{out}}}^{\text{eff}} &= \lambda_{\text{out}} \sum_{n=0}^{K-1} \alpha_{\text{out}}(n)\pi(n)
\end{aligned}
$$

Results

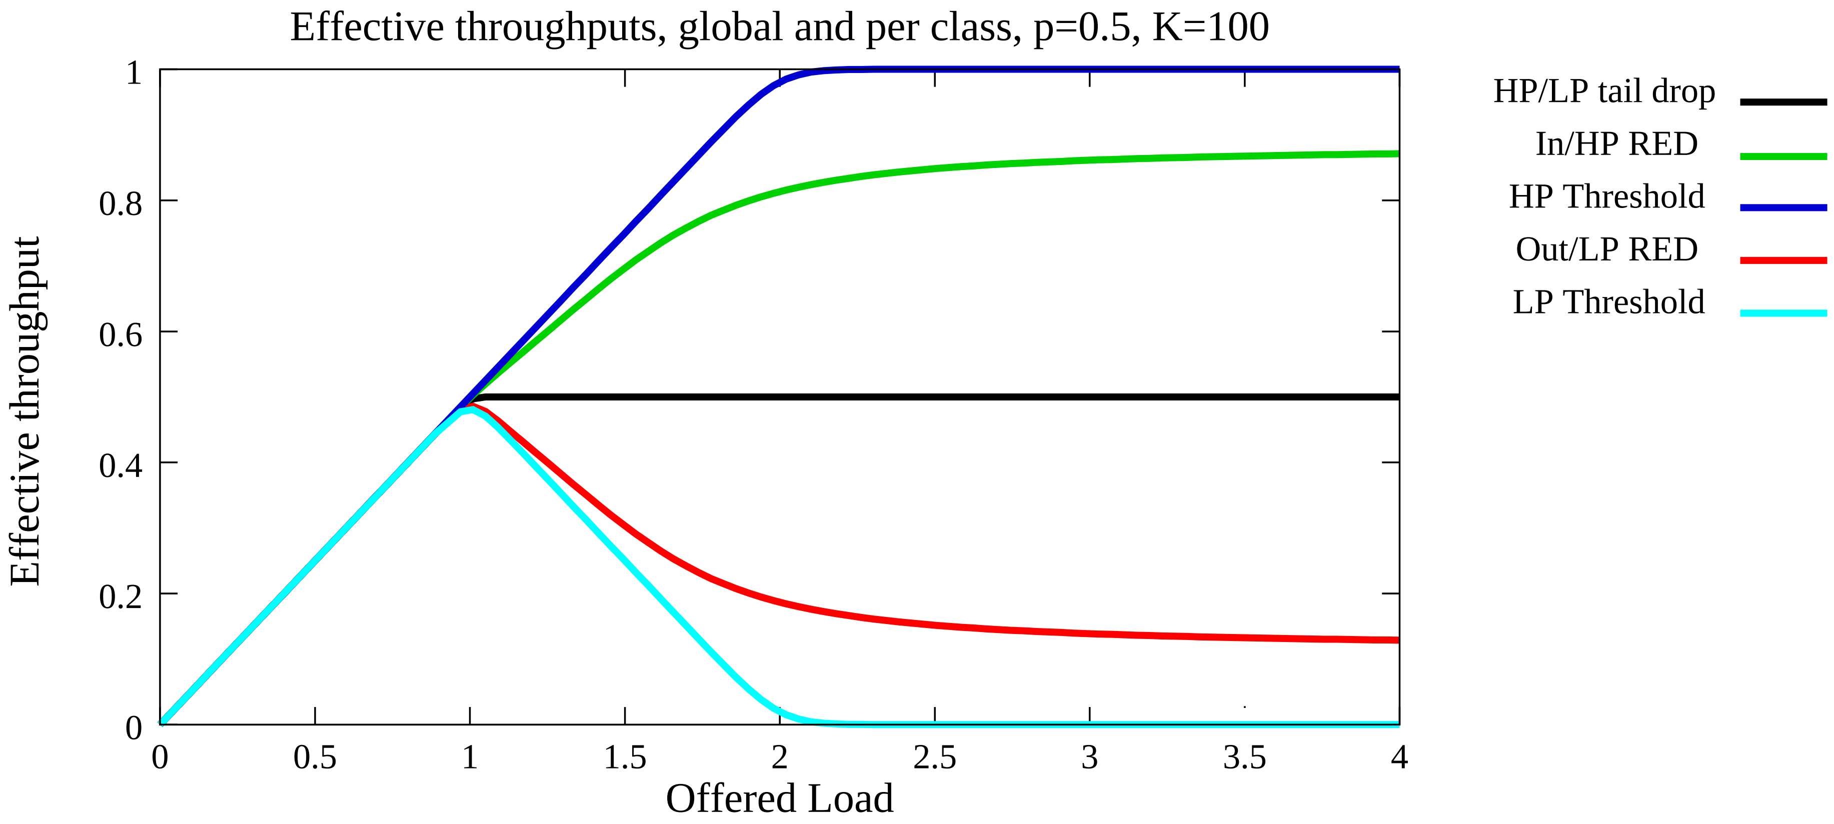

Effective throughputs, global and per class, p=0.5, K=100
Effective throughput
Offered Load
0
0.2
0.4
0.6
0.8
1
0
0.5
1
1.5
2
2.5
3
3.5
4
HP/LP tail drop
In/HP RED
HP Threshold
Out/LP RED
LP Threshold

Research Issues

- Investigate average queue length measurements:

$$\hat{q}_{n+1} \;=\; \alpha q_n \;+\; (1-\alpha)\hat{q}_n$$

- Find the proper value of α. Depends on the traffic?
- Investigate RIO based on the queue length of tagged packets instead of total queue length
- Investigate the interactions between RIO and TCP

Web Server Optimization

Caching: keep documents "closer" to the user.

Modeling: Sequence of documents $D_1, D_2, \ldots$.

Document d has the probability p_d of appearing.

Problem: maximize the hit ratio $= \mathbb{P}\{D_n$ is in the cache$\}$, or minimize the "cost" of retrieving a document.

Difficulty: the probabilities p_d are unknown...

Traditional solution for memory pages: LRU (Least Recently Used), tries to rank the probabilities.

Shown to be optimal if arrivals are i.i.d. Commonly used in computer systems.

Problems:

- does not take into account the size of the documents
  $\rightarrow$ document hit rate $\neq$ byte hit rate
- may be slow to converge
- may not work well if arrivals are not i.i.d.

Some solutions:

- *climbing* algorithms
- not use (empirical) hit frequency p_d but *density*

$$\delta_d = \frac{p_d}{s_d}, \qquad s_d \text{ size of } d.$$

- LFU (Least Frequently Used): estimate document frequency p_d

Other Web problems:

- use of the SRPT (shortest remaining processing time) policy in web servers
- replication/distribution of data

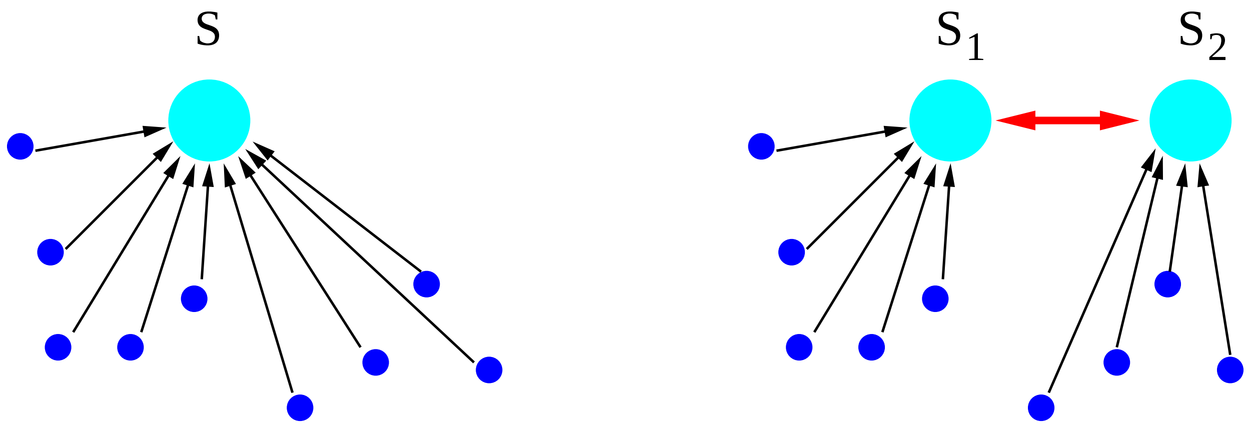

- realistic workload models for benchmarking

Multimedia over IP

Voice over IP

- influence of coding
- influence of the locality of losses
- usefulness of the FEC (Forward Error Correction) mechanisms

Video over IP: real time / video on demand

- compute buffer/bandwidth requirements
- optimal smoothing of traffic

For both:

- multi-level coding with adapted FEC
- adaptive transmission rate
- adaptive redundancy level

Example: video/voice playout

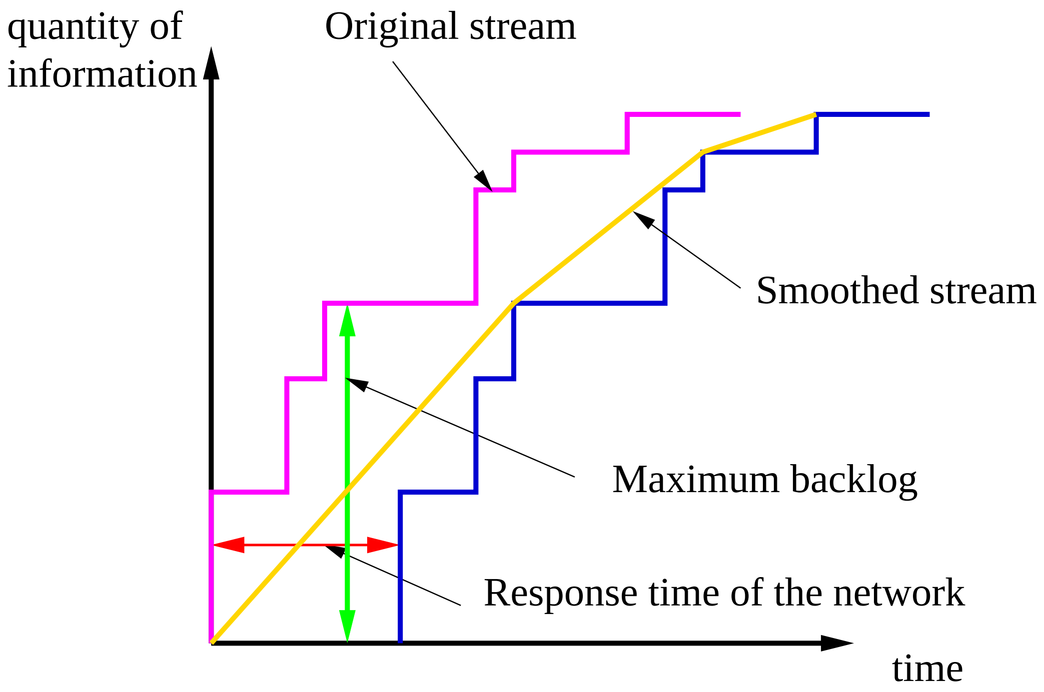

FEC

Forward Error Correction consists in adding redundancy to data so that it can cope with loss.

Assume a stream of packets of the same size, grouped in blocks of size n.

It is possible to add k packets to each block so that any k losses in the super-block of $n + k$ packets can be recovered.

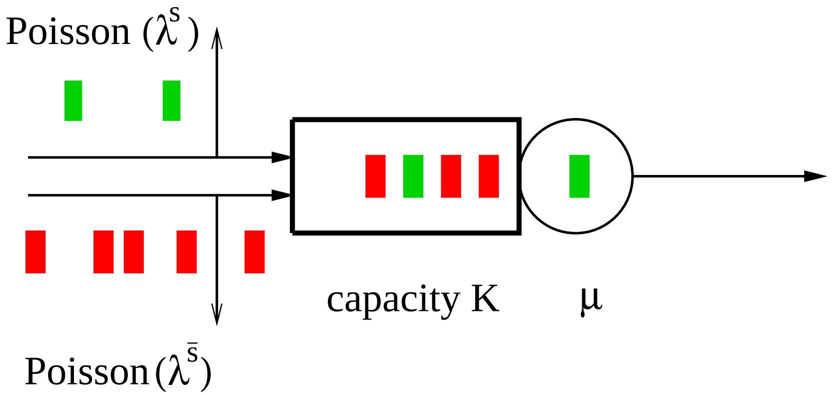

- $\nearrow$ probability loss recovery for the group

- $\nearrow$ load, $\nearrow$ probability of losing an individual packet

Does the compromise bring a global benefit? What is the optimal value of k?

Active Networking

The control is escaping the network:

- users pay for some service
- users may want to choose the route for their traffic
- proposals for "active packets" learning their way

This poses new types of problems relevant to Game Theory.

Example 1: definition of *fairness*?

Minimax criterion $\max_{\text{user } u} \; \min_{\text{link } \ell} \; \theta(\text{user } u \text{ on link } \ell)$

Average criterion $\max_{\phi} \left| \text{all user } u \text{ receives proportion } \phi \text{ of his demand} \right.$

Example 2: Braess' Paradox

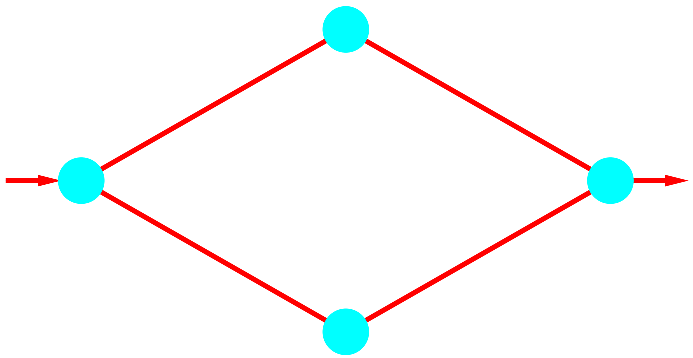

Utility U_1

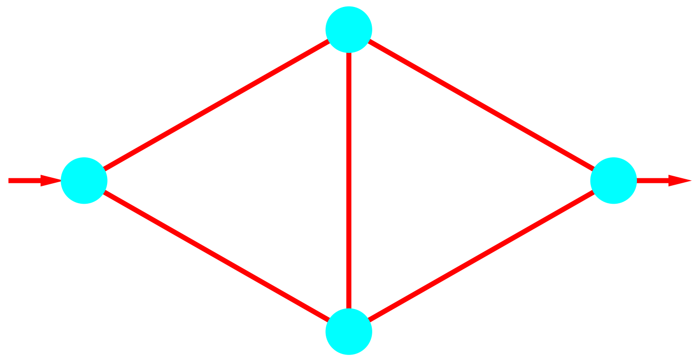

Utility $U_2 < U_1$!!!

Game Theory concepts:

- Nash Equilibrium $(x_1^*, \ldots, x_n^*)$

$$U_i(x_1^*, \ldots, x_i^*, \ldots, x_n^*) \;\geq\; U_i(x_1^*, \ldots, x_i, \ldots, x_n^*) \; .$$

- Pareto Efficient solution $(y_1^*, \ldots, y_n^*)$:

$$\nexists (x_1, \ldots, x_n), \qquad U_i(x_1, \ldots, x_n) \;\geq\; (y_1^*, \ldots, y_n^*) \quad \forall i \; .$$

In general, Nash equilibria are very inefficient...

Conclusion

Many challenges, and more:

- group communications: control of multicast
- ATM networks (rate control algorithms, call admission control, pricing, policing)
- mobile communications (capacity planning, prevention of interruptions)
- satellite communications (noisy, asymmetric transmission, low orbit constellations)
- optical communications (wavelength allocations, hot potato routing)

Many opportunities for Performance Evaluation.