CS 5229: Advanced Compute Networks

Traffic Model and Engineering

Dr. Chan Mun Choon
School of Computing, National University of Singapore

Aug 26, 2010

Readings

- Readings
  - HK Choi and John O. Limb, “A Behavior Model for Web Traffic,” ICNP 1999.
  - Vern Paxon and Sally Floyd, “Wide-Area Traffic: The Failure of Poisson Modeling,” IEEE Transaction on Network, pp. 226-244, June, 1995.
- Reference
  - S. Keshav, “An Engineering Approach to Computer Networking”, Chapter 14: Traffic Management

Motivation for Traffic Models

- In order to predict the performance of a network system, we need to be able to "describe" the "behavior" of the input traffic
  - Often, in order to reduce the complexity, we classify the user behavior into classes, depending on the applications
  - Sometimes, we may be even able to "restrict" or shape the users' behavior so that they conform to some specifications
- Only when there is a traffic model is traffic engineering possible

An example

- Executive participating in a worldwide videoconference
- Proceedings are videotaped and stored in an archive
  - Edited and placed on a Web site
  - Accessed later by others
- During conference
  - Sends email to an assistant
  - Breaks off to answer a voice call

What this requires

- For video
  - *sustained bandwidth of a few hundred Kbps (depends on quality and screen size)*
  - *low loss rate*
- For voice
  - *sustained bandwidth of at least 8 kbps*
  - *low loss rate*
- For interactive communication
  - *low delay (< 100 ms one-way)*
- For playback
  - *low delay jitter*
  - *or .....*
- For email and archiving
  - *reliable bulk transport*

Traffic management

- Set of policies and mechanisms that allow a network to *efficiently* satisfy a *diverse* range of service requests
- Tension is between **diversity** and **efficiency**
- Traffic management is necessary for providing *Quality of Service (QoS)*

Time Scale of Traffic Management

- Less than one round-trip-time (cell-level)
  - Perform by the end-points and switching nodes
  - Scheduling and buffer management
  - Regulation and policing
  - Policy routing (datagram networks)
- One or more round-trip-times (burst-level)
  - Perform by the end-points
  - Feedback flow control
  - Retransmission
  - Renegotiation

Time Scale (cont.)

- Session (call-level)
  - End-points interact with network elements
  - Signaling
  - Admission control
  - Service pricing
  - Routing (connection-oriented networks)
- Day
  - Human intervention
  - Peak load pricing
- Weeks or months
  - Human intervention
  - Capacity planning

The two camps

- Can increase welfare either by
  1. matching services to user requirements
  2. increasing capacity blindly
- Which is cheaper? Utilization vs. complexity
  - depends on technology advancement
  - User behavior/expectation/tolerance
  - small and smart or big and dumb
- Smarter ought to be better?
  - otherwise, to get low delays for some traffic, we need to give *all traffic* low delay, even if it doesn't need it
  - But, if bandwidth is cheap and control is complex, may be cheaper to increase capacity

Telephone traffic models (Call)

- How are calls placed?
  - call arrival model
  - studies show that time between calls is drawn from an exponential distribution
  - call arrival process is therefore *Poisson*
  - memoryless: the fact that a certain amount of time has passed since the last call gives no information of time to next call
- How long are calls held?
  - usually modeled as exponential
  - however, measurement studies (in the mid-90s) show that it is *heavy tailed*
    - A small number of calls last a very long time
    - Why?

Packet Traffic Model for Voice

- A single voice source is well represented by a two state process: an alternating sequence of active or talk spurt, follow by silence period
  - Talk spurts typically average 0.4 – 1.2s
  - Silence periods average 0.6 – 1.8s
  - Talk spurt intervals are well approximated by exponential distribution, but not true for silence period
  - Silence periods allow voice packets to be multiplexed

Ref: Chapter 3 of "Broadband Integrated Networks", by Mischa Schwartz, 1996.

Internet traffic modeling

- A few apps account for most of the traffic
  - WWW, FTP, E-mail
  - P2P
- A common approach is to model apps
  - time between app invocations
  - connection duration
  - # bytes transferred
  - packet inter-arrival distribution

Hyoung-Kee Choi and John Limb,
"A Behavioral Model of Web Traffic,"
ICNP 1999.

Web Download Model

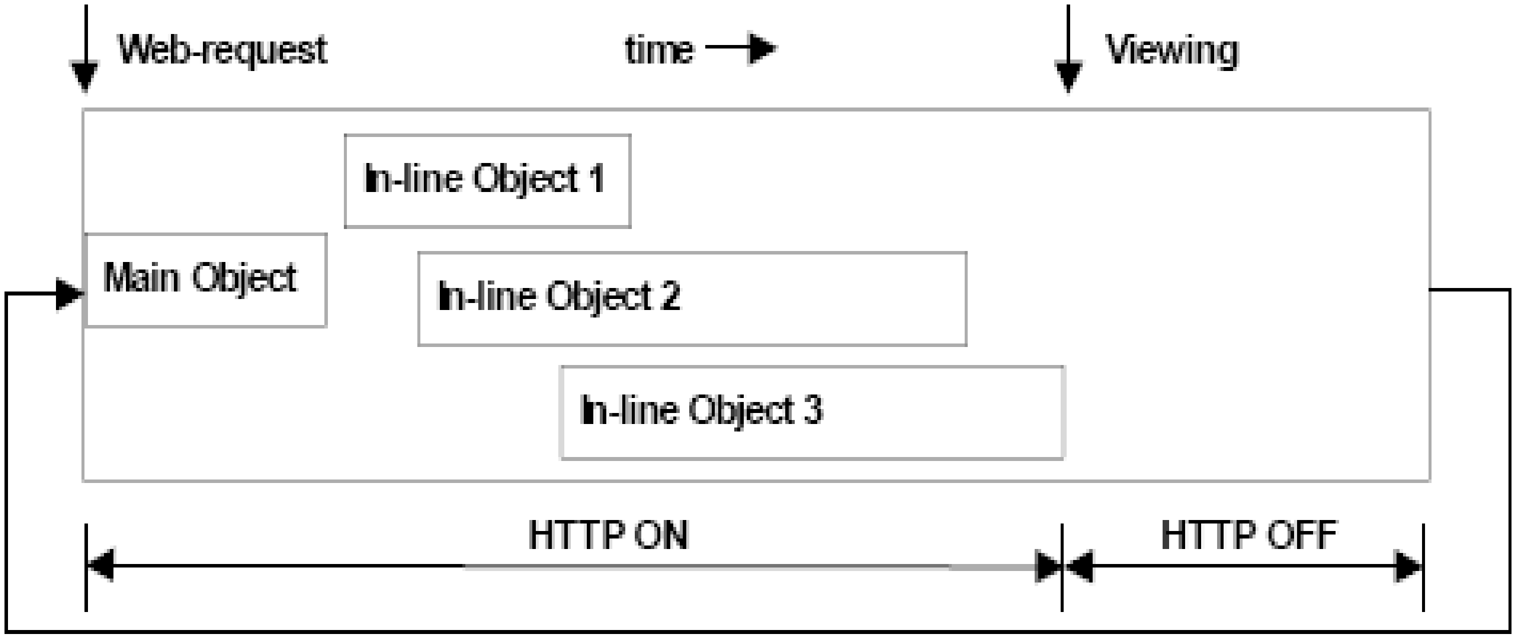

Figure 1: Overview of the basic model

Measurement

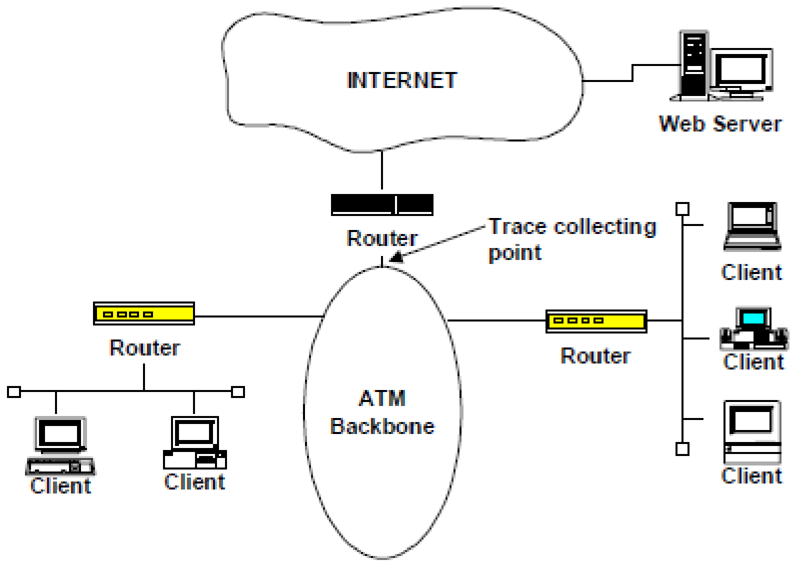

Figure 2: Perspective of the Georgia Tech campus network.

- 1 hour trace (done in 1998)
- > 1900 clients
- ~ 24,000 Web-requests

Q-Q Plot

- a **Q-Q plot** is a graphical method for comparing two probability distributions by plotting their quantiles against each other.
- If the two distributions being compared are similar, the points in the Q-Q plot will approximately lie on the line *y* = *x*.
- Refence: http://en.wikipedia.org/wiki/Q-Q_plot

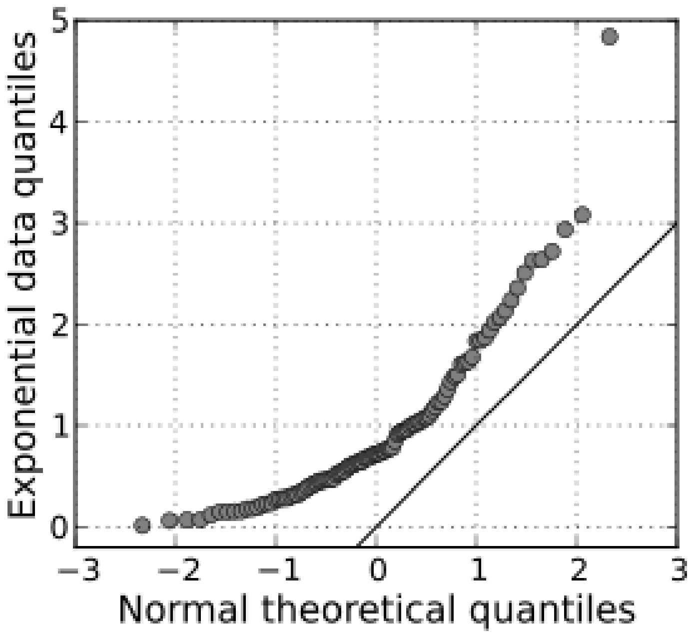

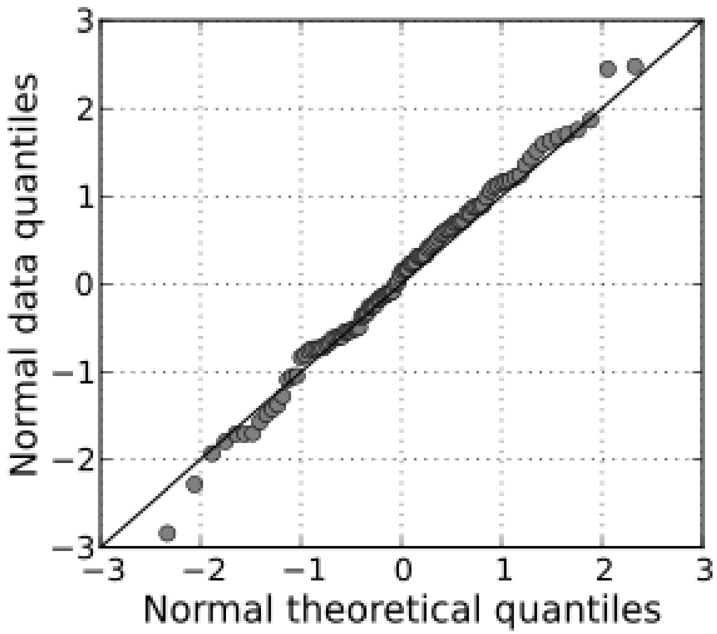

Results

- Distribution models: Weibull, Lognormal, Gamma, Chi-square, Pareto and Exponential (Geometric) distributions.

Parameters		Mean	Median	S.D.	Best-fit
Request size		360.4	344	106.5	LN
Object size	Main	10710	6094	25032	LN
	In-line	7758	1931	126168	LN
Parsing time		0.13	0.06	0.187	G
Number of In-line objects		5.55	2	11.4	G
In-line Inter-Arrival time		0.86	0.17	2.15	G
Viewing (OFF) time		39.5	11.7	92.6	W
Number of Web-requests	Non-cached	12.6	5	21.6	LN
	Cached	1.7	1	1.7	GM

Table 3: Summary statistics for HTTP parameters
(LN=Lognormal, G=Gamma, W=Weibull and GM=Geometric)

Traditional Assumptions

- Packet and connection arrivals are often assumed to be Poisson
- A number of studies have shown that the inter-arrivals are clearly not exponential
- Use of Poisson models under-estimate the “burstiness” of traffic

Why is the result important?

- Congestion Modeling
  - Congestion can be longer than expected, with losses concentrated over a small period
  - Slight increase in traffic can result in large increase in loss rate

M/G/1 Queue

$\overline{X}$: Average service time

$\overline{X^2}$: Second moment of service time

(b) Mean residual service time

$$R = \frac{\lambda \overline{X^2}}{2}$$

(c) Pollaczek-Khinchin formula

$$W = \frac{R}{1-\rho} = \frac{\lambda \overline{X^2}}{2(1-\rho)}$$

Check for Poisson Traffic

- For a exponential distribution with mean μ,
  - $E(\mu^2) = 2/\mu^2$
- Mean Waiting Time (W)

  $= \lambda\ (2/\mu^2) / 2(1-\rho)$

  $= (\lambda/\mu) / \mu(1-\rho)$

  $= \rho / (\mu-\lambda)$

"Uniform" Traffic

- Link speed is 32,000 bps
- λ = 3 packet per second (Poisson arrivals)
- Average packet size is 400 bytes, μ = 10 pkts/s
- Uniform scenario:
  - Size of packet is constant, 400 bytes or 3,200 bits
- ρ = 0.3, E(X) = 0.1s
- $E(X^2) = 0.01$
- W = (3 * 0.01) / 2(1 – 0.3)

  = 0.02143s or <u>21.43ms</u>

"Bursty" Traffic (1)

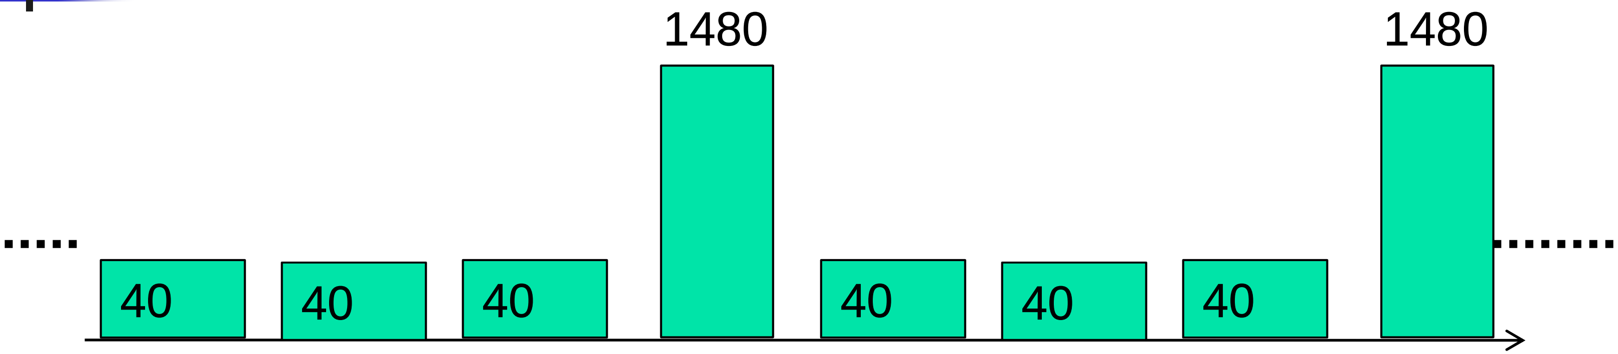

- ρ = 0.3, E(X) = 0.1s
- P(size=40bytes) = 0.75, P(size=1480bytes) = 0.25
- $E(X^2) = 0.75*(40*8/32000)^2 + 0.25*(1480*8/32000)^2$
  $= 0.000075 + 0.034225 = 0.0343$
- W = 3 (0.0343) / 2(1 - 0.3) = 0.0735s or 73.5ms

Bursty Traffic (2)

- $\rho = 0.3$, E(X) = 0.1s
- P(size=10bytes) = 0.9, P(size=3910bytes) = 0.1
- $E(X^2) = 0.9(10*8/32000)2 + 0.1(3910*8/32000)2$

  ~ 0.09556
- W = 3 (0.09556) / 2(1 – 0.3)

  = 0.2048s or 204.8ms

- More “bursty” traffic leads to longer waiting time (or more loss)
- What is bursty traffic?

Heavy Tail

- **heavy-tailed distributions** are probability distributions whose tails are not exponentially bounded

$$\lim_{x \to \infty} e^{\lambda x} \Pr[X > x] = \infty \quad \text{for all } \lambda > 0.$$

Intuitively, there is a small, but non negligible, chance that x can be very large

Example

- Exponential:
  - $P(X>x) = e^{-\lambda x}$

- Weibull distribution ($0<k<1$):
  - $P(X>x) = e^{-(x/\lambda)^k}$

Long Tail

$$\lim_{x \to \infty} \Pr[X > x + t | X > x] = 1,$$

- If the long-tailed quantity exceeds some high level, the probability approaches 1 that it will exceed any other higher level.
- All long-tailed distributions are heavy-tailed, but the converse is false, and it is possible to construct heavy-tailed distributions that are not long-tailed.

Example

- Weibull distribution (0<k<1):
  - P(X>x) = $e^{-(x/\lambda)^k}$

Pareto Distribution

- Pareto originally used this distribution to describe the allocation of wealth among individuals
  - A larger portion of the wealth of any society is owned by a smaller percentage of the people
  - Sometimes expressed more simply as the 80-20 rule
- Other examples:
  - The sizes of human settlements (few cities, many hamlets/villages)
  - File size distribution of Internet traffic which uses the TCP protocol (many smaller files, few larger ones)
- Zipf's Law or zeta distribution: discrete counterpart of Pareto

Pareto Distribution

The classical Pareto distribution with shape parameter β and location parameter a has the cumulative distribution function [HK80]:

$$F(x) = P[X \leq x] = 1 - (a/x)^{\beta}, \quad a, \beta \geq 0, \quad x \geq a,$$

with the corresponding probability density function:

$$f(x) = \beta a^{\beta} x^{-\beta-1}.$$

If $\beta \leq 2$, then the distribution has infinite variance, and if $\beta \leq 1$, then it has infinite mean.

Example

- Pareto: $P(X>x) = x^{-1.5}$ $(a=1, \beta=1.5)$

Exponential/Pareto Distribution

- Exponential Distribution: $P(X>x) = e^{-x/3}$
- Pareto Distribution: $P(X>x) = x^{-1.5}$ ($a=1, \beta=1.5$)
- Means of both distributions are 3

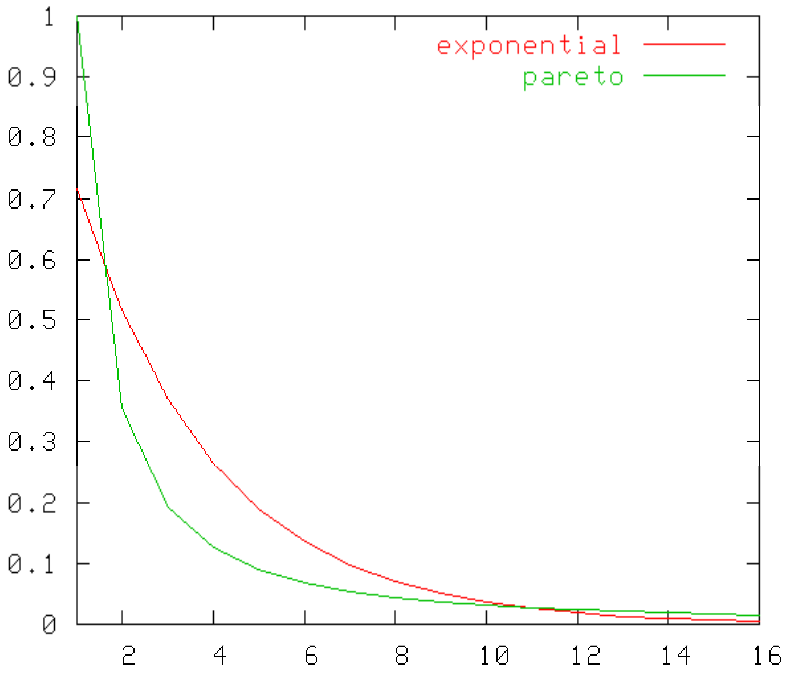

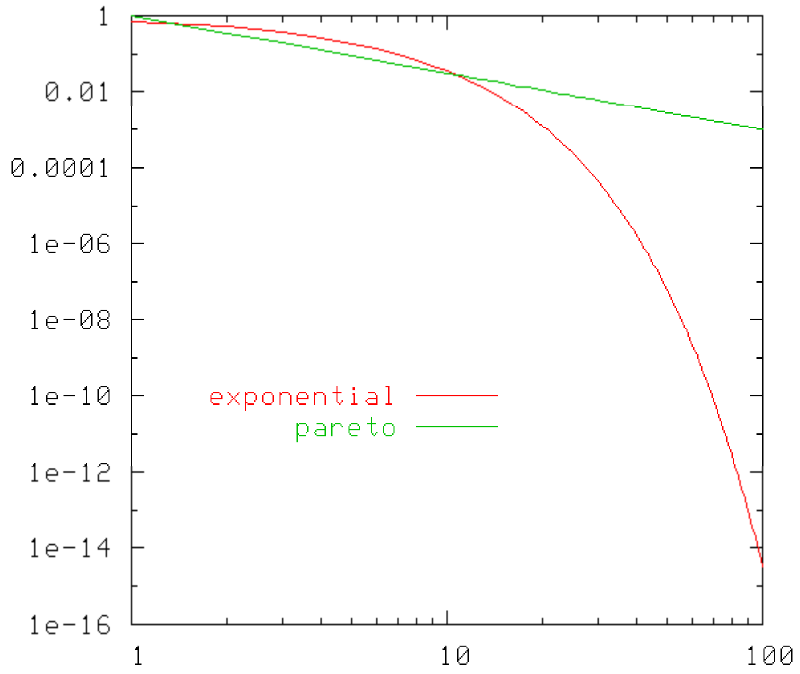

Conditional Mean Exceedance (CME)

$$\mathrm{CME_x} = E[X - x | X \geq x].$$

- Uniform distribution
  - CME is a decreasing function of x
  - The longer you wait, the sooner you will be done
- Exponential distribution
  - CME is independent of x (memoryless)
  - Additional waiting time is independent of time already spent waiting
- Long Tail
  - CME is increasing with x
  - The longer you have waited, the more likely to wait for a longer period of time

- W. Leland, M. Taqqu, W. Willinger, and D. Wilson, “On the Self-Similar Nature of Ethernet Traffic (Extended Version),” *IEEE/ACM Transactions on Networking, 2(1),* pp. 1-15, February 1994.

Vern Paxson and Sally Floyd,

"Wide-Area Traffic: The Failure of Poisson Modeling,"

Transaction of Networking, June 1995.

TCP Trace

Dataset	Date	Duration	What
Bellcore (BC)	10Oct89	13 days	17K TCP conn.
U.C.B. (UCB)	31Oct89	24 hours	38K TCP conn.
coNCert (NC)	04Dec91	24 hours	63K TCP conn.
UK-US (UK)	21Aug91	17 hours	26K TCP conn.
DEC 1-3	See refs.	24 hours $\times 3$	195K TCP conn.
LBL 1-8	See refs.	30 days $\times 8$	3.7M TCP conn.

Table 1: Summary of Wide-Area TCP Connection Traces

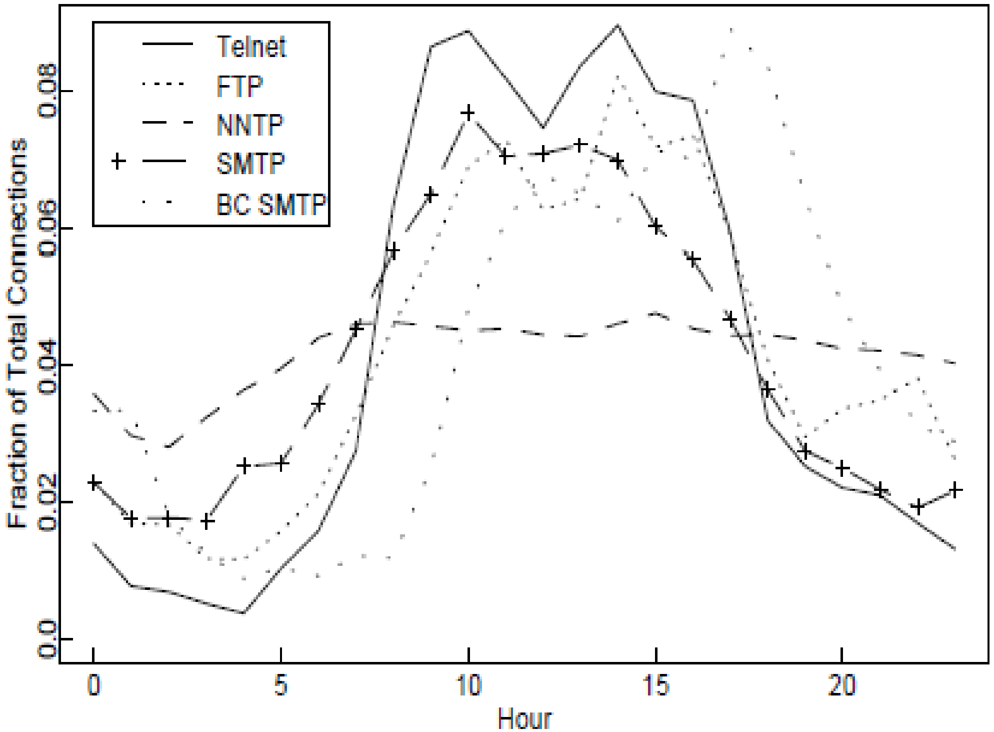

Figure 1: Mean, relative, hourly connection arrival rate for LBL-1 through LBL-4 datasets.

Telephone traffic is fairly well modelled during one-hour intervals using homogeneous Poisson arrival processes

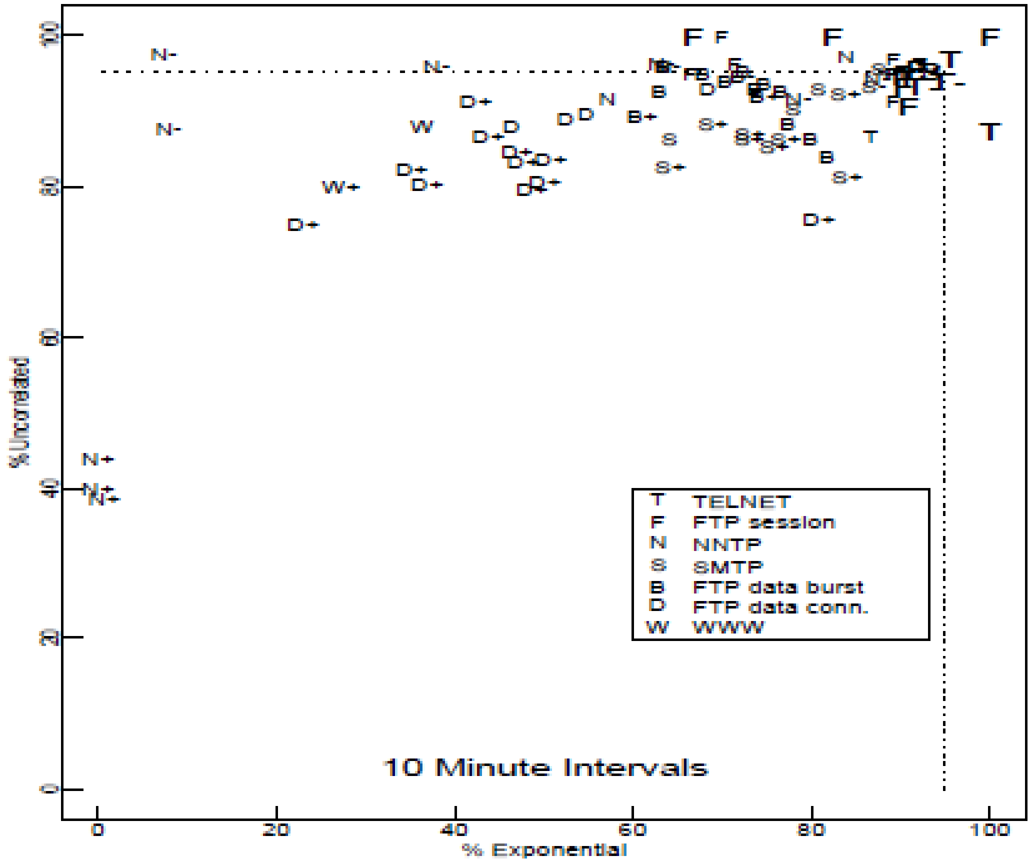

Figure 2: Results of testing for Poisson arrivals.

TELNET connection arrivals and FTP session arrivals are very well modelled as Poisson, both for 1-hour and 10-minute fixed rates.

Packet Trace

Dataset	Date	When	What
LBL PKT-1	Fri 17Dec93	2PM-4PM	1.7M TCP pkts.
LBL PKT-2	Wed 19Jan94	2PM-4PM	2.4M TCP pkts.
LBL PKT-3	Thu 20Jan94	2PM-4PM	1.8M TCP pkts.
LBL PKT-4	Fri 21Jan94	2PM-3PM	1.3M pkts.
LBL PKT-5	Fri 28Jan94	2PM-3PM	1.3M pkts.
DEC WRL-1	Wed 08Mar95	10PM-11PM	3.3M pkts.
DEC WRL-2	Thu 09Mar95	2AM-3AM	3.9M pkts.
DEC WRL-3	Thu 09Mar95	10AM-11AM	4.3M pkts.
DEC WRL-4	Thu 09Mar95	2PM-3PM	5.7M pkts.

Table 2: Summary of Wide-Area Packet Traces

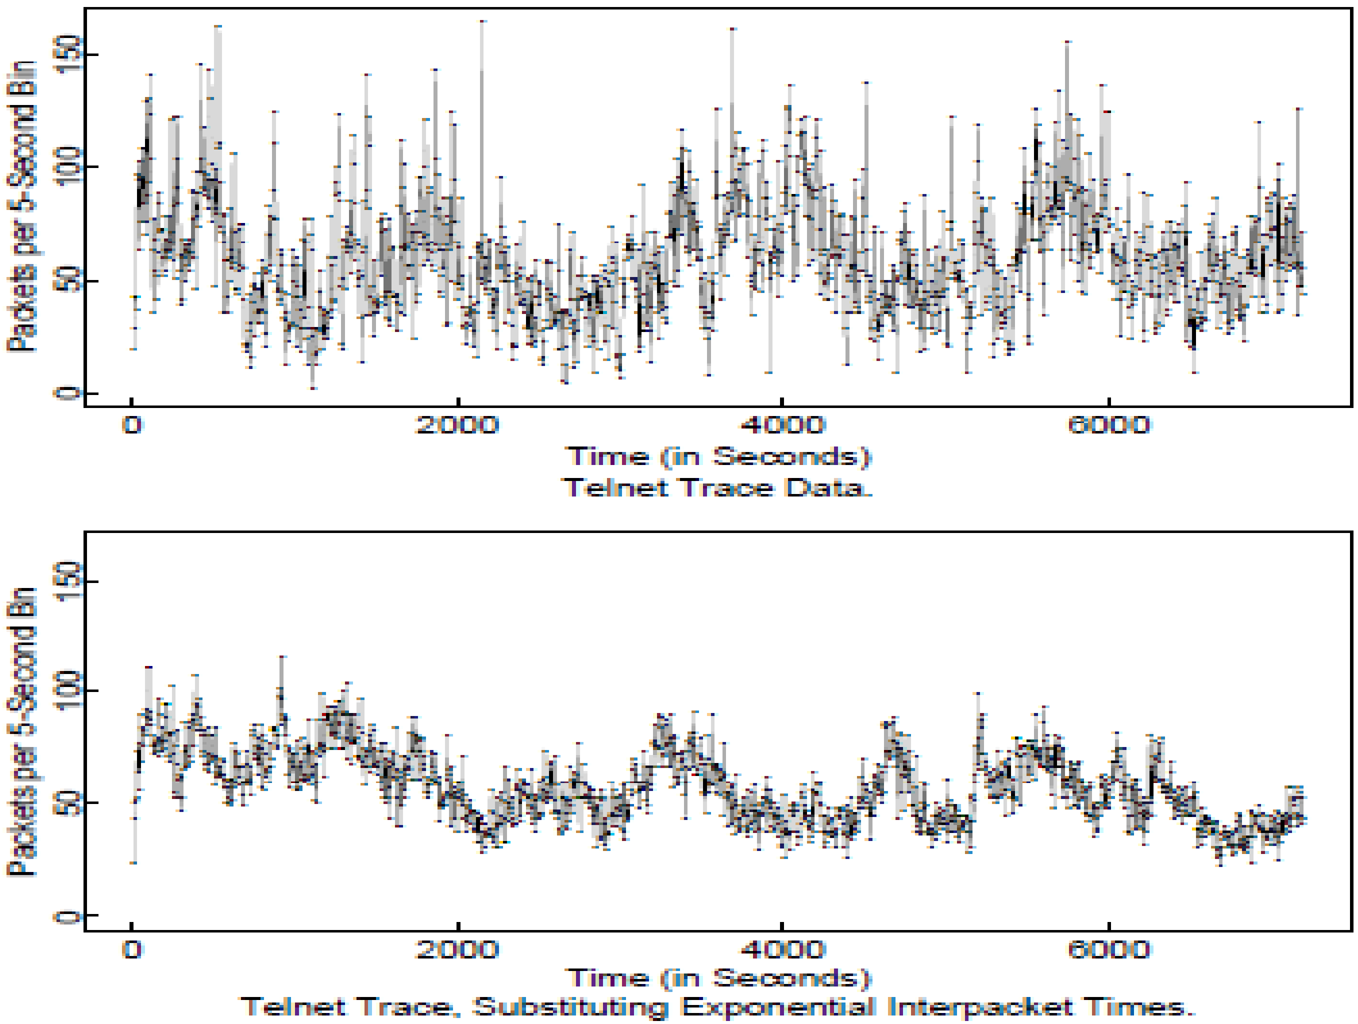

Figure 6: Comparisons of actual and exponential TELNET packet interarrival times.

Variance Time-Plot

- Variance-time plots are obtained by plotting log(var(X(m))) against log(m) ("time") and by fitting a simple least squares line through the resulting points in the plane
- For most processes, the result is a straight line with slope equals to -1
- For self similar process, the line is much much flatter, between -1 and 0

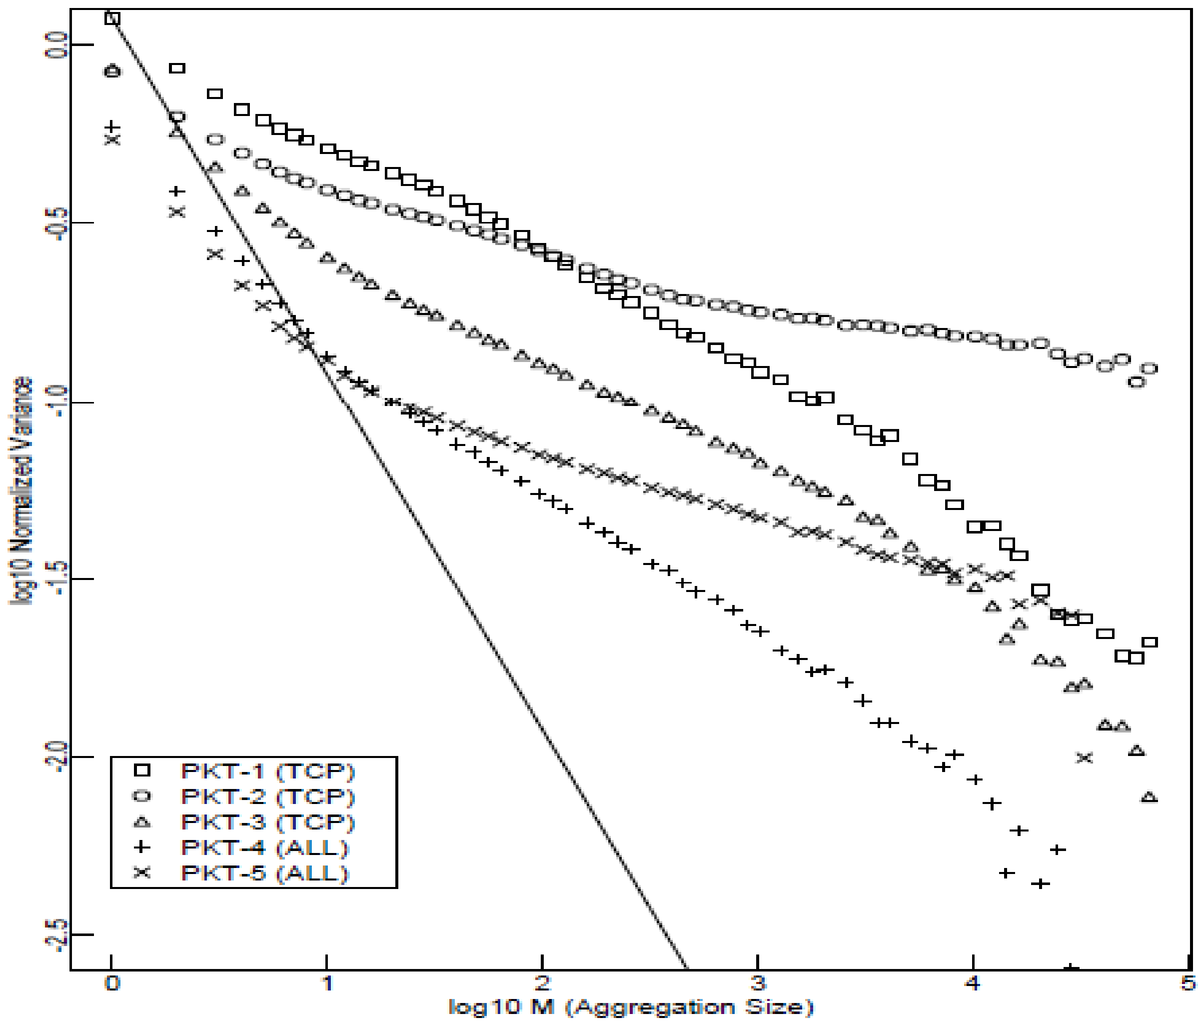

Figure 12: Variance-time plot for all TCP / all link-level packet arrivals in the LBL PKT datasets.

Summary of Observations

- For interactive TELNET traffic
  - connection arrivals are well-modeled as Poisson
  - However, packet arrivals are not Poisson
- Similarly, for FTP traffic
  - Session arrivals are Poisson
  - Data connections within a FTP session is not Poisson
  - Distribution of file size transfer for a data connection is heavy tail
- For SMTP/NNTP traffic, connection arrivals are not Poisson

T Karagiannis, et. al, “A Nonstationary Poisson View of Internet Traffic,” INFOCOM 2004

Poisson (still useful ...)

- Since the original data set was collected (1989, 1994), both link speeds and the number of Internet-connected hosts have increased by more than three orders of magnitude
  - In the 1994 packet trace, there are 26M packets in 12 hours
- Study the Poisson assumption's validity on several OC48 (2.5 Gbps) backbone traces taken from CAIDA (Cooperative Association for Internet Data Analysis)
- Traces: from Aug 2002 – Apr 2003,
  - In about 50 minutes, there are 434M packets

Cont’d

- Result can be interpreted at different time scales
  1. Packet arrivals appear Poisson at sub-second time scales.
  2. Internet traffic appears non-stationary at multi-second time scales.
  3. Internet traffic exhibits long-range dependence (LRD) at scales of seconds and above.
- As the Internet increases in size and the technologies connected to it change, the appropriate traffic models need to be reevaluated.

Traffic Shaping

- Traffic may be "shaped" or "smoothed" to reduce any adverse impact on the network
  - Usually, buffer the packets at the "access" routers and then send out packets at a smoothed, more regular rate
- The so called "leaky bucket" is a popular mechanism

Policing Mechanisms

Three common-used criteria:

- *(Long term) Average Rate:* how many packets/bits can be sent per unit time (in the long run)
  - crucial question: what is the interval length: 100 packets per sec or 6000 packets per min have same average!
- *Peak Rate:* e.g., 6000 pkts per min. (ppm) avg.; 15000 ppm peak rate
- *(Max.) Burst Size:* max. number of pkts/bits sent consecutively (with no intervening idle)

Example

- Policy
  - Average rate = 1000pps or 424Kbps
  - Peak rate = 2Kpps or 848Kbps
  - Burst Size = 1000 packets or 424Kb
- Policy 2
  - Average rate = 1000pps or 424Kbps
  - **Peak rate = 4Kpps or 1696Kbps**
  - Burst Size = 1000 packets or 424Kb
- Policy 3
  - Average rate = 1000pps or 424Kbps
  - Peak rate = 2Kpps or 848Kbps
  - **Burst Size = 2000 packets or 828Kb**

Example (Worst Case)

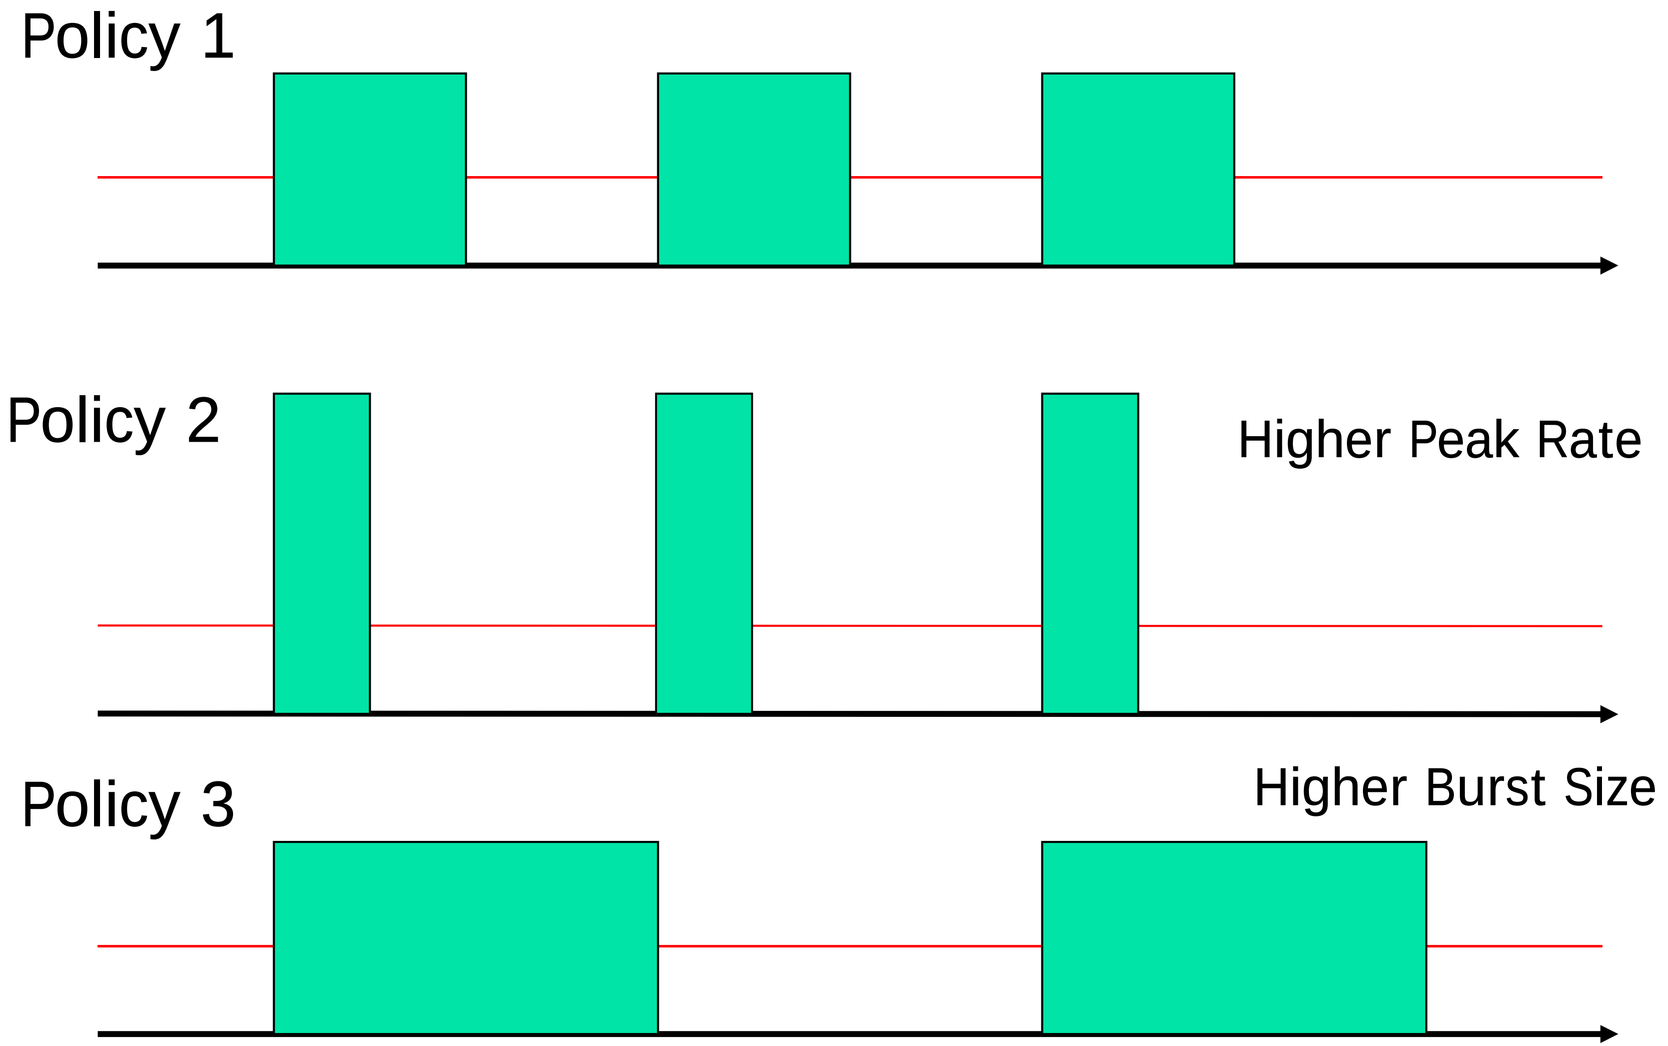

Policing Mechanisms

Token Bucket: limit input to specified Burst Size and Average Rate.

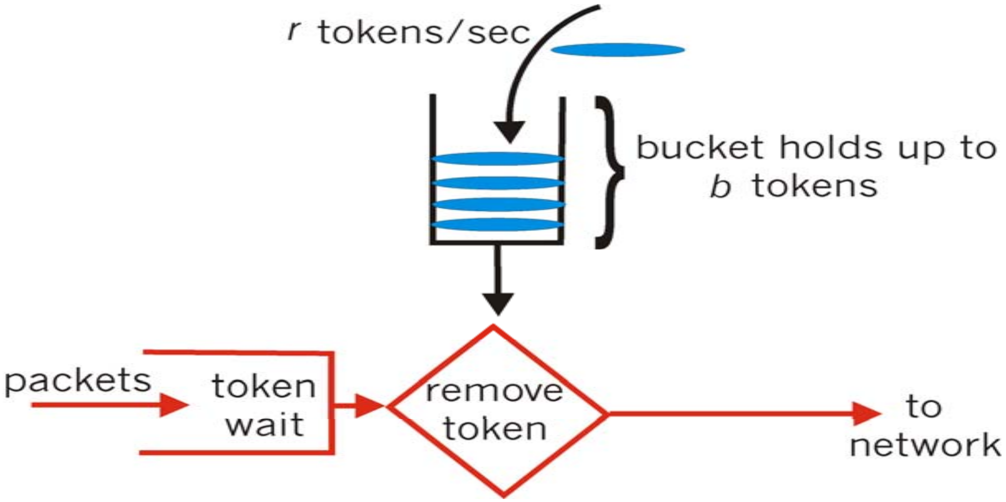

- bucket can hold σ tokens
- tokens generated at rate ρ *token/sec* unless bucket full
- *over interval of length t: number of packets admitted less than or equal to* $(\rho t + \sigma)$*.*

Policing Mechanism

- How useful is such a policing mechanism?
- What are the pros and cons?