Internet Traffic Modeling and Its Implications to Network Performance and Control

Kihong Park
Department of Computer Sciences
Purdue University
park@cs.purdue.edu

Outline

- Motivation
- Traffic modeling
- Performance evaluation
- Traffic control

Motivation

Traffic Measurement

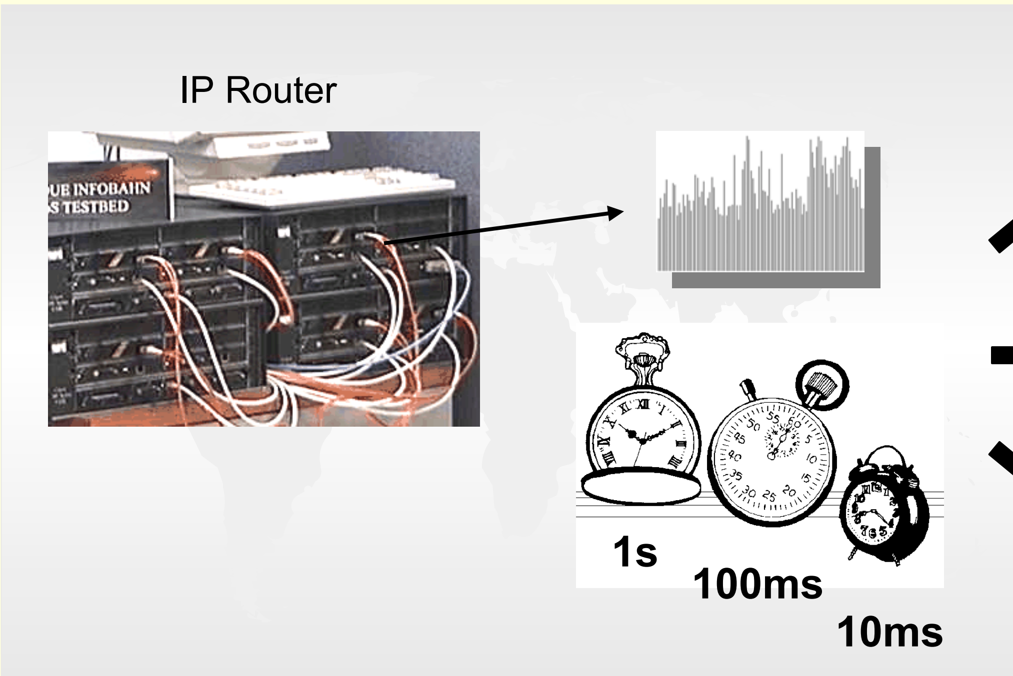

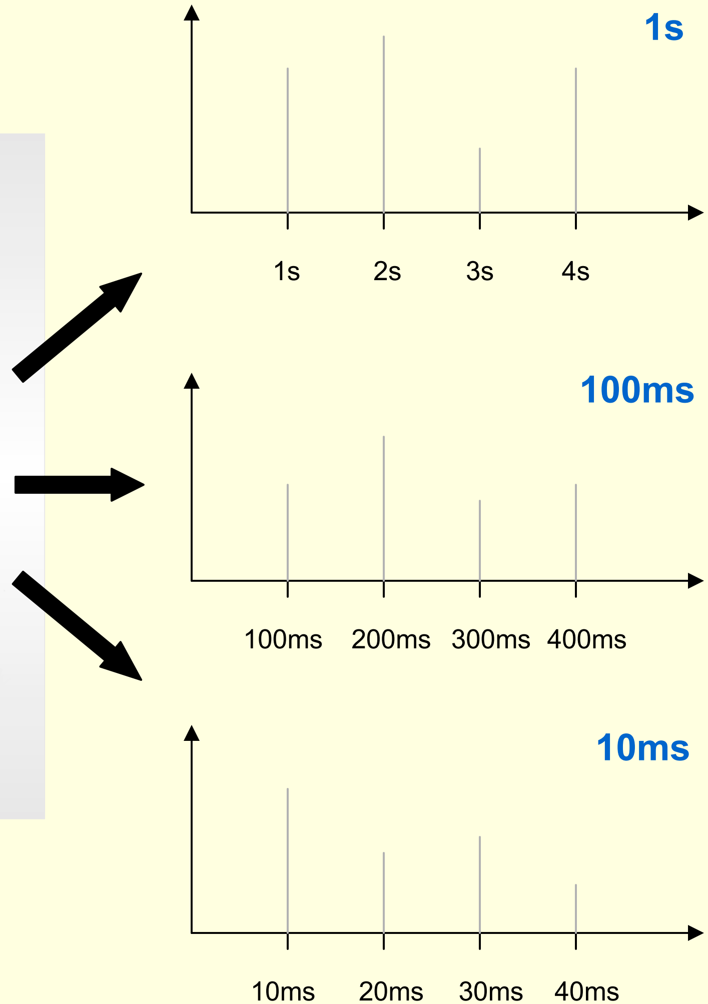

Traffic Burstiness

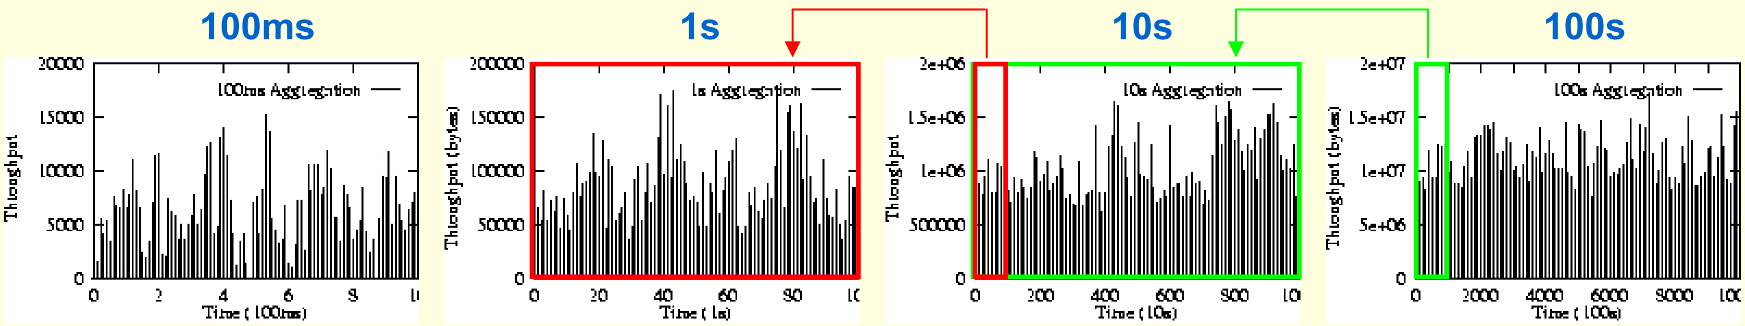

Network Traffic

- Bursty across multiple time scales: 100ms ~ 100s
- Fractal or self-similar: the whole resembles its parts

In Contrast…

100ms

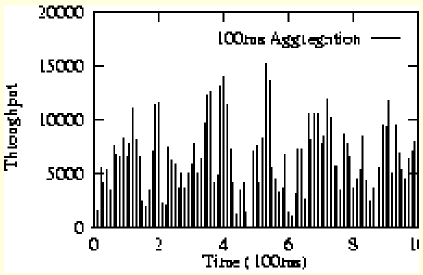

1s

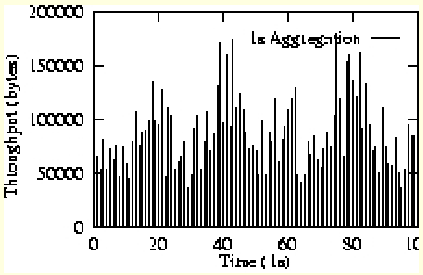

10s

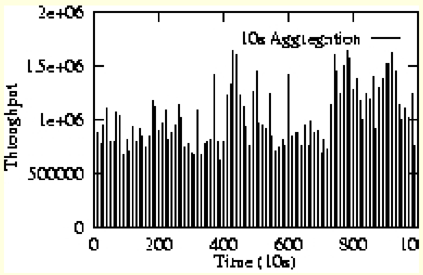

100s

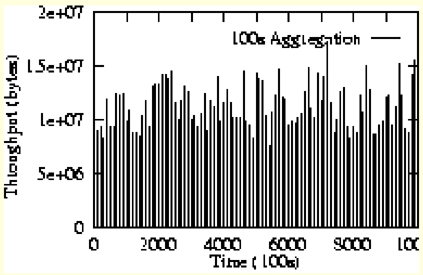

Network Traffic

100ms

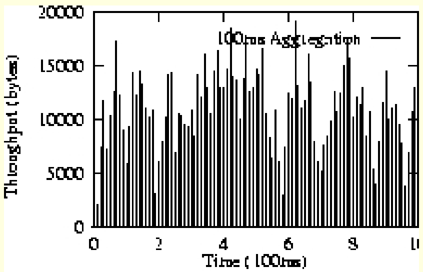

1s

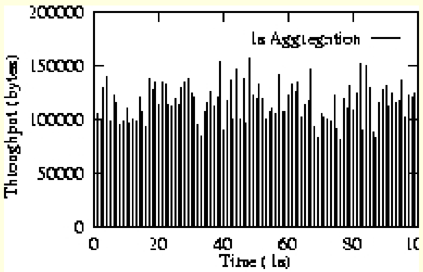

10s

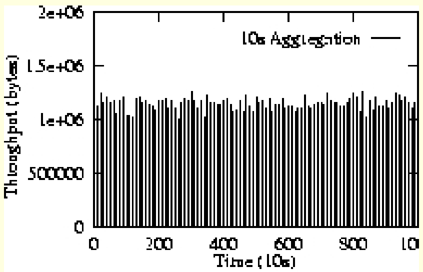

100s

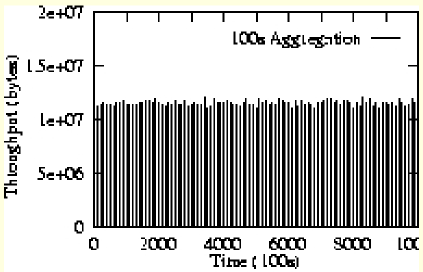

Poisson Traffic

Self-Similar Burstiness

- Burstiness preserved across multiple time scales
- Deterministic self-similarity

2nd order structure

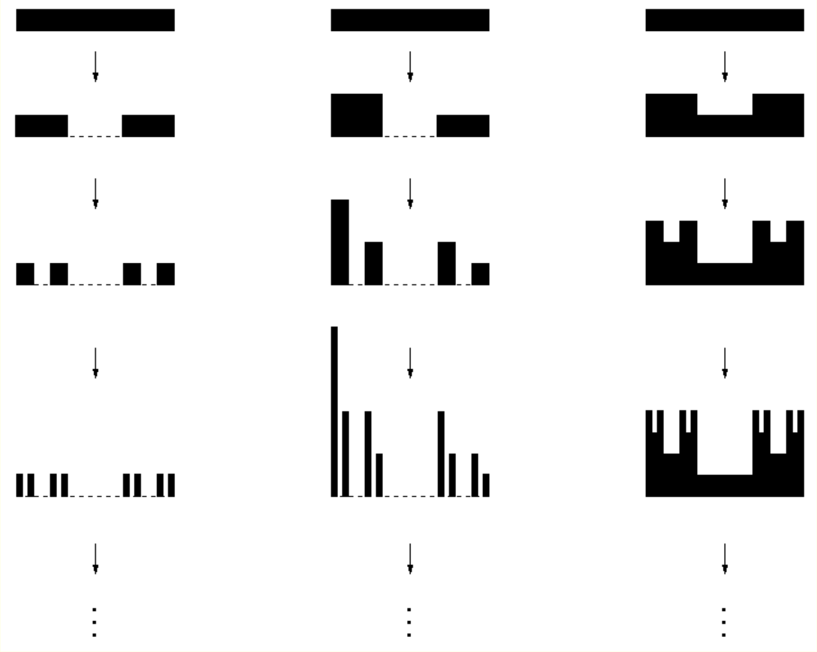

"Cantor Traffic"

Sustained Contention

100ms 1s 10s 100s

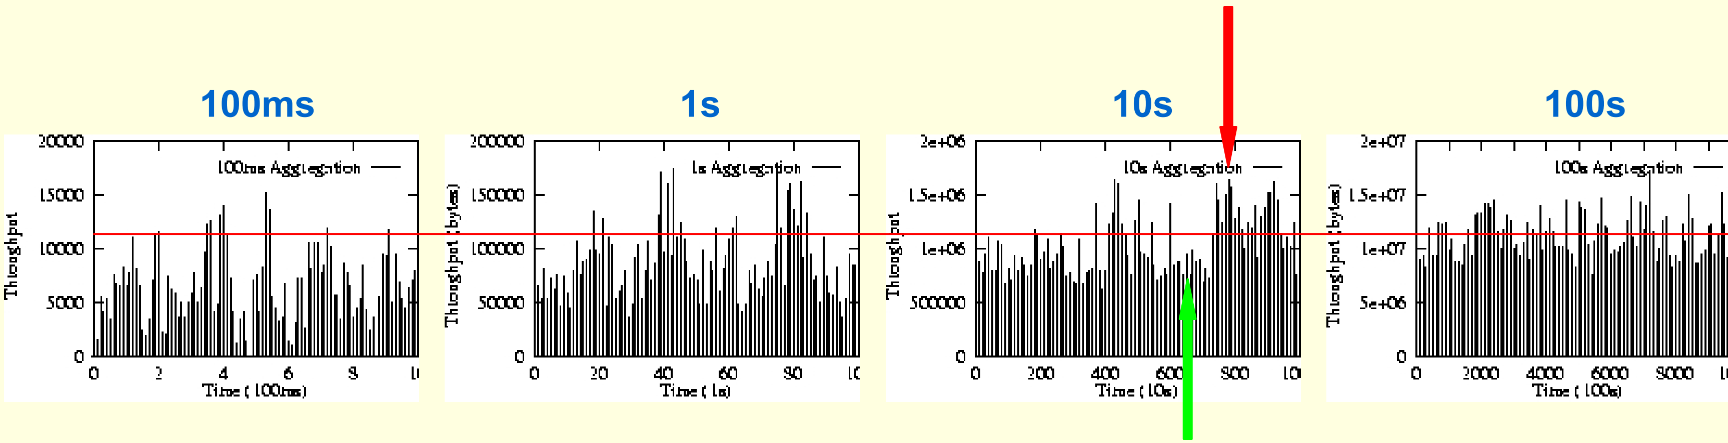

Network Traffic

100Kb/s 1Mb/s 10Mb/s 100Mb/s

100ms 1s 10s 100s

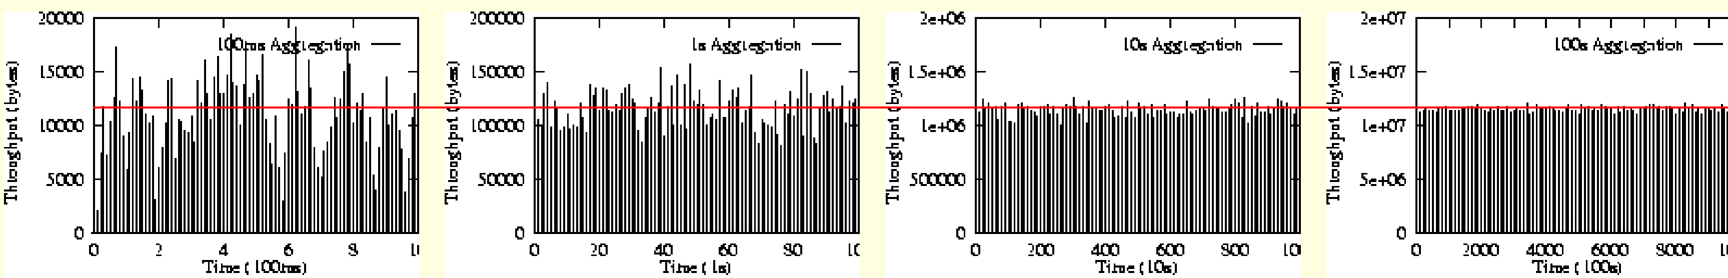

Poisson Traffic

Correlation at a Distance …

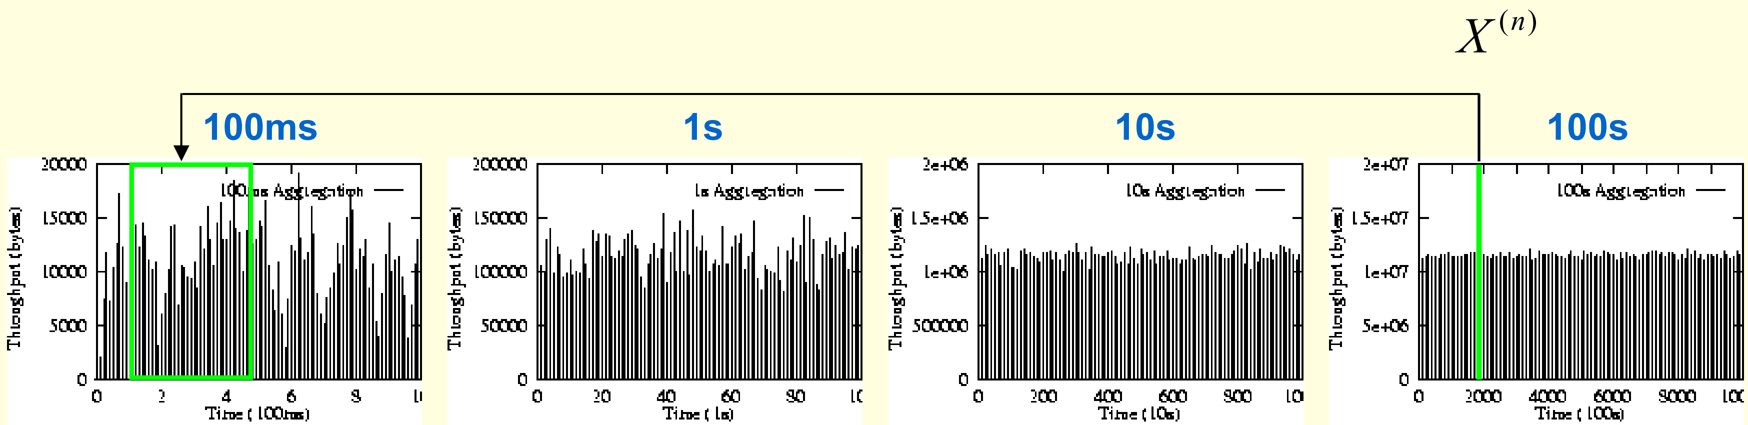

$$X^{(n)} = (X_1 + X_2 + \cdots + X_n)/n$$

- For example, if i.i.d. then
  - By LLN, concentrates around mean $E[X_1]$
  - Sample variance $V[X^{(n)}] = \sigma^2 n^{-1}$

Presence of Strong Correlation

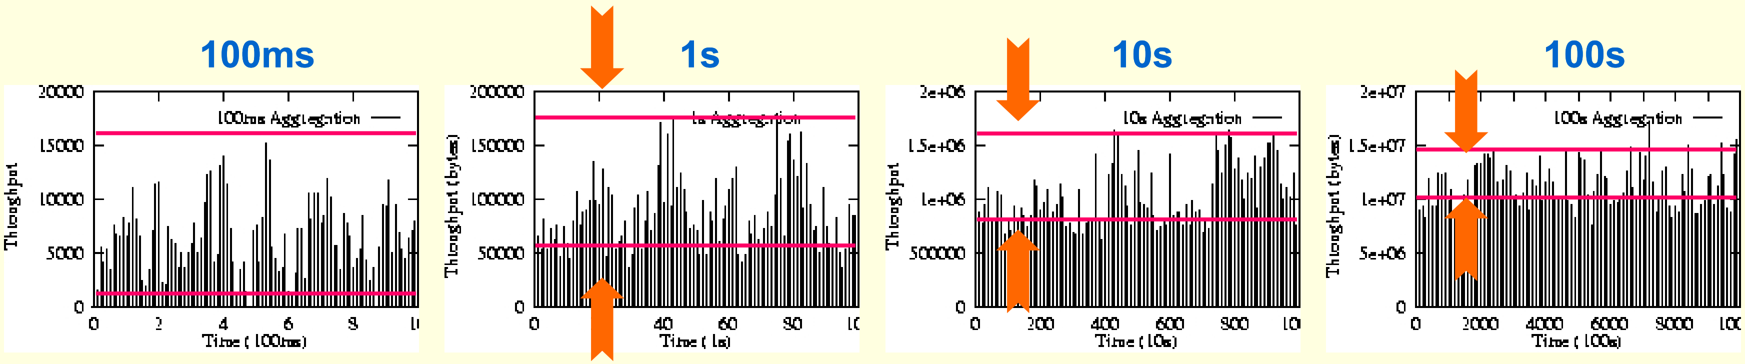

$$X^{(n)} = (X_1 + X_2 + \cdots + X_n)/n$$

- Slower rate of dampening for network traffic

$$V[X^{(n)}] = \sigma^2 n^{-\beta}, \qquad 0 < \beta < 1$$

$$= \sigma^2 n^{2H-2}, \qquad 1/2 < H < 1$$

Empirical Evidence

- LAN traffic: Bellcore (‘89-92)
  - Ethernet
  - Leland *et al.*, SIGCOMM ‘93
- WAN traffic: LBL + others
  - TCP
  - Paxson & Floyd, SIGCOMM ‘94
- Many more
  - Internet Traffic Archive (ita.ee.lbl.gov)
  - NLANR (pma.nlanr.net/PMA)
  - etc.

Key Points

- Internet traffic is bursty over large time scales
- Has potential to affect performance
- Ubiquitous empirical phenomenon
- Causes, impact, and control

Traffic Modeling

Workload Granularity

- Aggregate traffic
- Connection, flow, or session arrival
- Flow duration or lifetime
- Packet arrival within single connection

Workload Granularity

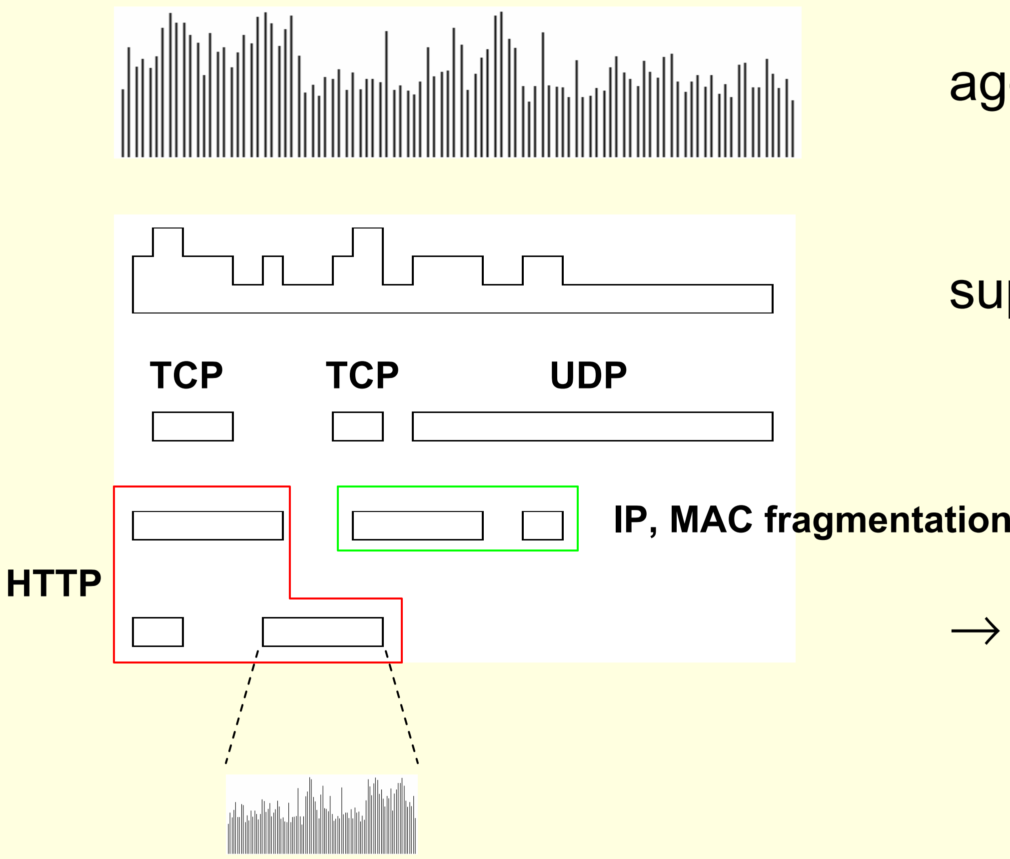

aggregate traffic $X^{(n)}(t)$

superposition of sessions

$\rightarrow$ multi-level

Workload Property: Session Arrival

- Connection arrivals
  - Poisson
  - TCP measurements up to mid-’90s
  - Paxson & Floyd, SIGCOMM ‘94

- Refinement
  - Weibull $\Pr[Z > x] = e^{-ax^c}, \quad 0 < c < 1$
  - Pre- vs. post-WWW TCP interarrivals
  - A. Feldmann, PW 2000

Workload Property: Lifetime

- Connection duration
  - Heavy-tailed $\Pr[Z > x] \approx x^{-\alpha}, \quad 0 < \alpha < 2$
    - $\rightarrow$ large x; regularly varying r.v.
    - $\rightarrow$ infinite variance; if $0 < \alpha < 1$, unbounded mean
  - LAN & WAN measurements up to mid-’90s
  - Paxson & Floyd ’94; Willinger *et al*. ‘95

- Not restricted to network traffic

Workload Property: Lifetime

- UNIX file size distribution
  - File systems research ’80s, Park *et al*., ICNP 96
  - G. Irlam ‘93

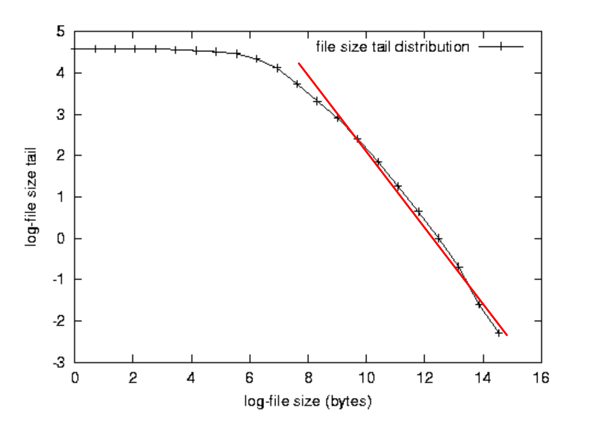

$$\log \Pr[Z > x] = -\alpha \log x$$

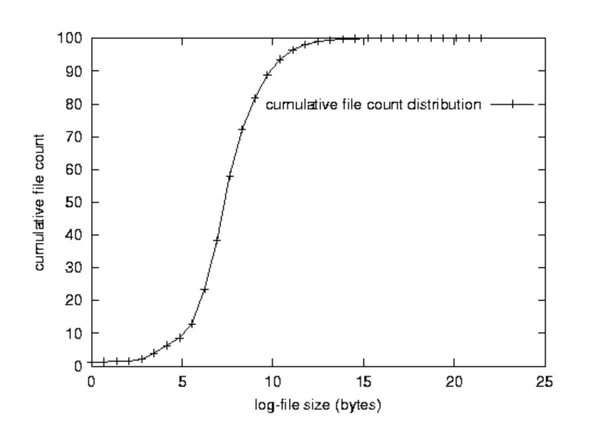

90% are less than 20KB

Workload Property: Lifetime

- UNIX file size distribution
  - File systems research ’80s, Park *et al*., ICNP 96
  - G. Irlam ‘93

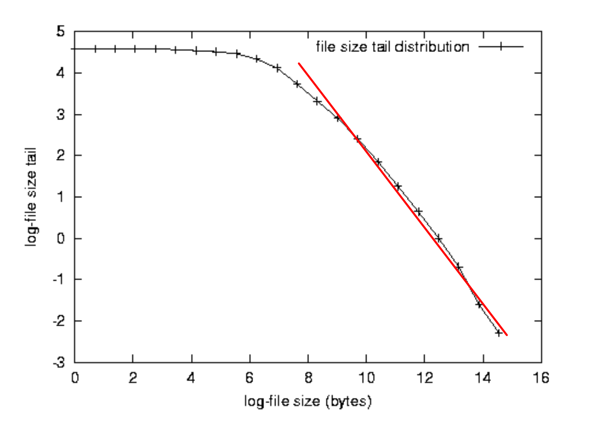

$$\log \Pr[Z > x] = -\alpha \log x$$

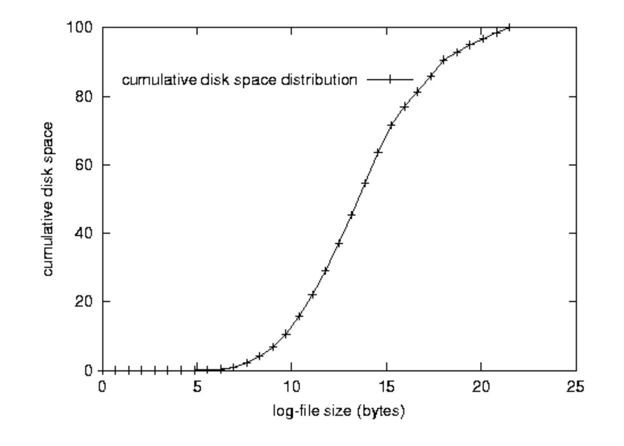

10% take up 90% disk space

“mice and elephants”

Workload Property: Lifetime

- Relevance of UNIX file system
  - Independent non-networking evidence from ’80s
    - $\rightarrow$ different research community, objectives
  - Intimate relationship with traffic burstiness
  - Bellcore measurements: ’89 -’92
    - $\rightarrow$ pre-Web, pre-MPEG video streaming
  - WWW: Crovella & Bestavros, SIGMETRICS ’96
    - $\rightarrow$ C. Cunha; early Web circa ’95

Workload Property: Lifetime

- Telephony: call holding time
  - Heavy-tailed; Duffy *et al*., JSAC ‘94
  - Lognormal; V. Bolotin, JSAC ‘94
  - Call center design; Annals OR, ‘02

→ Past: has not mattered too much due to TDM

… Erlang’s loss formulae: avrg. service time

→ Present: different situation for VoIP

… how much impact?

→ Unclear. Voice: low bit rate real-time CBR

Workload Property: Lifetime

- Key difference with files
  - Long telephone call:

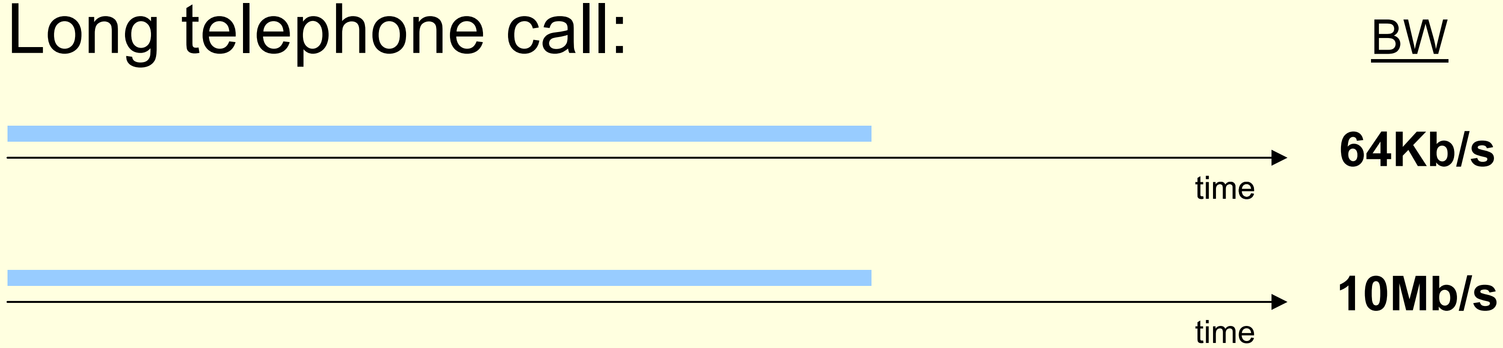

  - Large file:

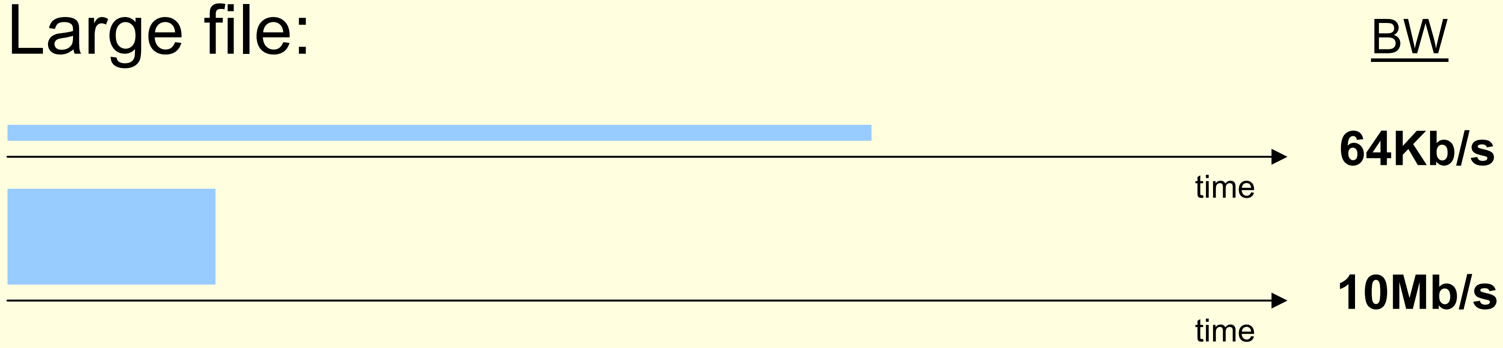

e.g., TCP

“stretching-in-time vs. stretching-in-space”, Park *et al*. ‘96

Workload Property: Lifetime

- UNIX process lifetime
  - Harchol-Balter & Downey, SIGMETRICS '96
  - Process migration: dynamic load balancing

- VBR Video

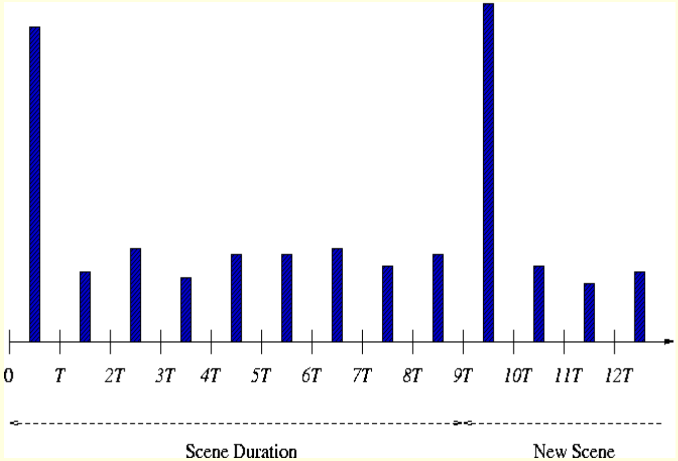

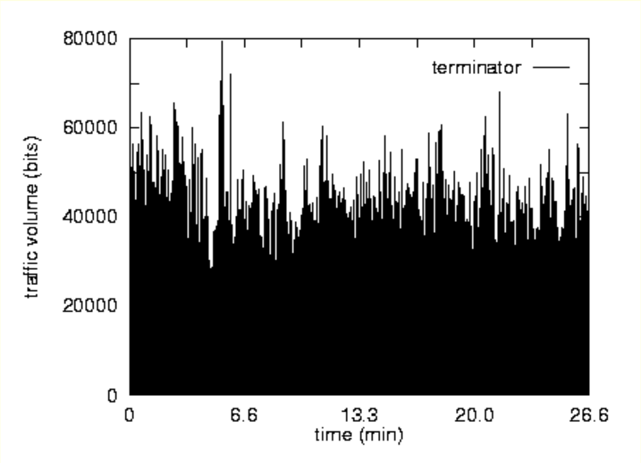

Workload Property: Lifetime

- UNIX process lifetime
  - Harchol-Balter & Downey, SIGMETRICS ‘96
  - Process migration: dynamic load balancing

- VBR Video

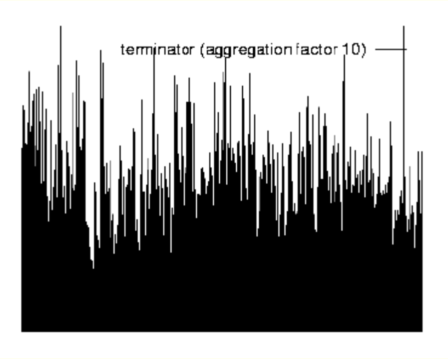

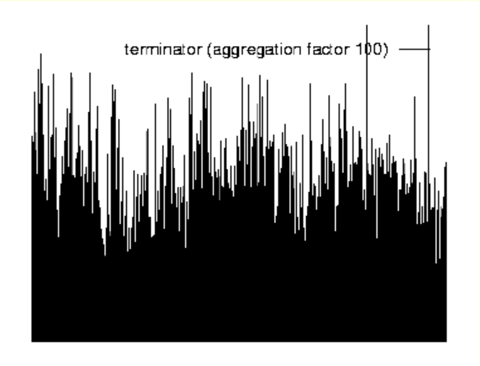

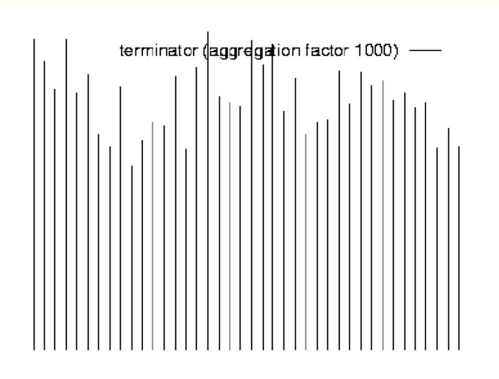

  - Beran *et al*., IEEE Trans. Commun., ’95 + others

Workload Property: Aggregate Traffic

- On/off model
  - Superposition of independent on/off sources
  - Willinger *et al*., SIGCOMM '95

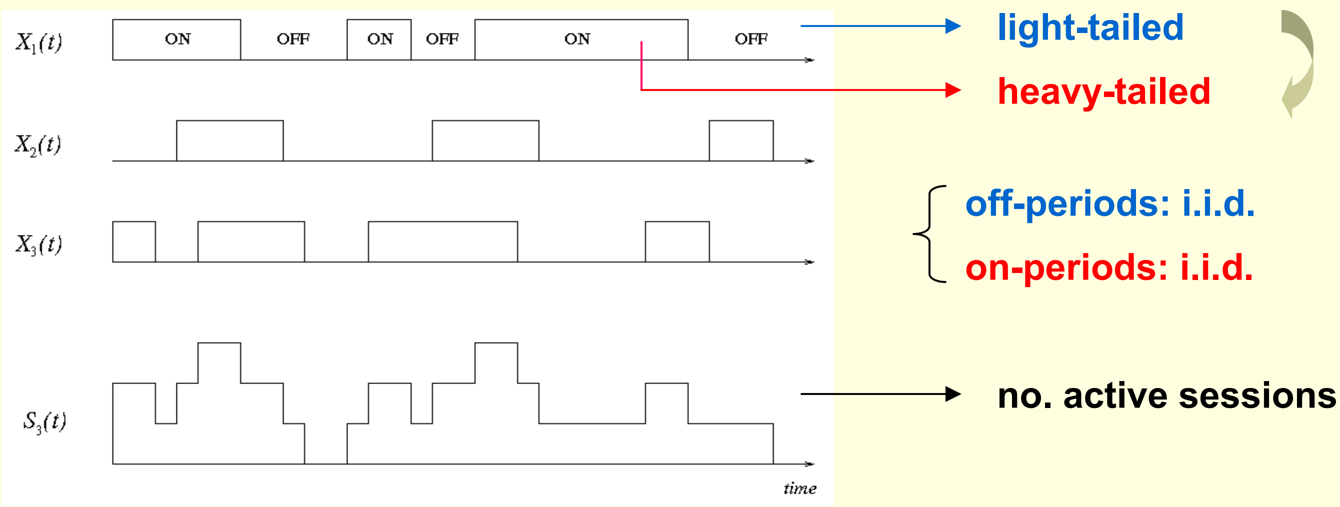

→ fractional Gaussian noise

Some Definitions

- *H-ss*
  - $Y(t)$ is *H-ss* if for all $a > 0, t \geq 0$

$$Y(t) =_d a^{-H} Y(at), \qquad 0 < H < 1$$

- *H-sssi*
  - *H-ss* and stationary increments

$$X(t) = Y(t) - Y(t-1)$$

- Fractional Brownian Motion
  - *H-sssi* and Gaussian
- Fractional Gaussian Noise
  - Increment process of FBM

Some Definitions

- $X(t)$ is exactly second-order self-similar if

$$\gamma(k) = \sigma^2((k+1)^{2H} - 2k^{2H} + (k-1)^{2H})/2, \qquad 1/2 < H < 1$$

$\rightarrow$ autocovariance

- Asymptotically second-order self-similar if

$$\gamma^{(m)}(k) \sim \sigma^2((k+1)^{2H} - 2k^{2H} + (k-1)^{2H})/2$$

$\rightarrow$ autocovariance of aggregated process

- Fact: $\gamma^{(m)}(k) = \gamma(k)$ for all $m \geq 1$

  $\rightarrow$ invariant w.r.t. second-order structure

Workload Property: Aggregate Traffic

- $M / G / \infty$
  - Poisson session arrivals
  - Heavy-tailed connection duration
  - Likhanov *et al*. ’95; Parulekar & Makowski ’96

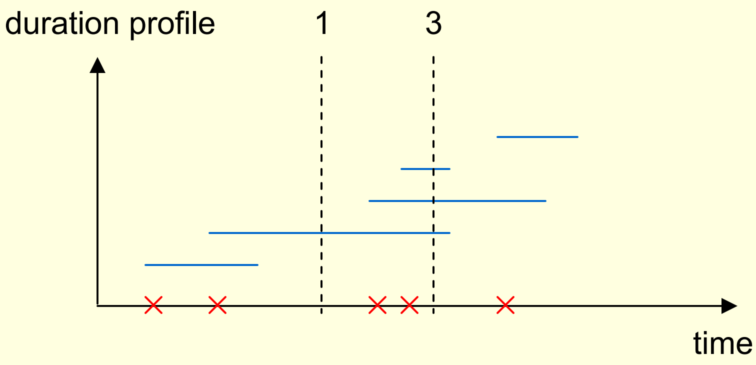

$\Leftrightarrow$ $M / G / \infty$ with *G* heavy-tailed

$\rightarrow$ asymptotically second-order self-similar

Workload Property: Aggregate Traffic

- $M / G / \infty$ and on/off model:

$$H = (3 - \alpha) / 2$$

- Heavy-tailedness parameter α determines Hurst parameter *H*
- Physical model of network traffic
- Versus black-box time series model

Impact of Protocol Stack

- Transport protocols

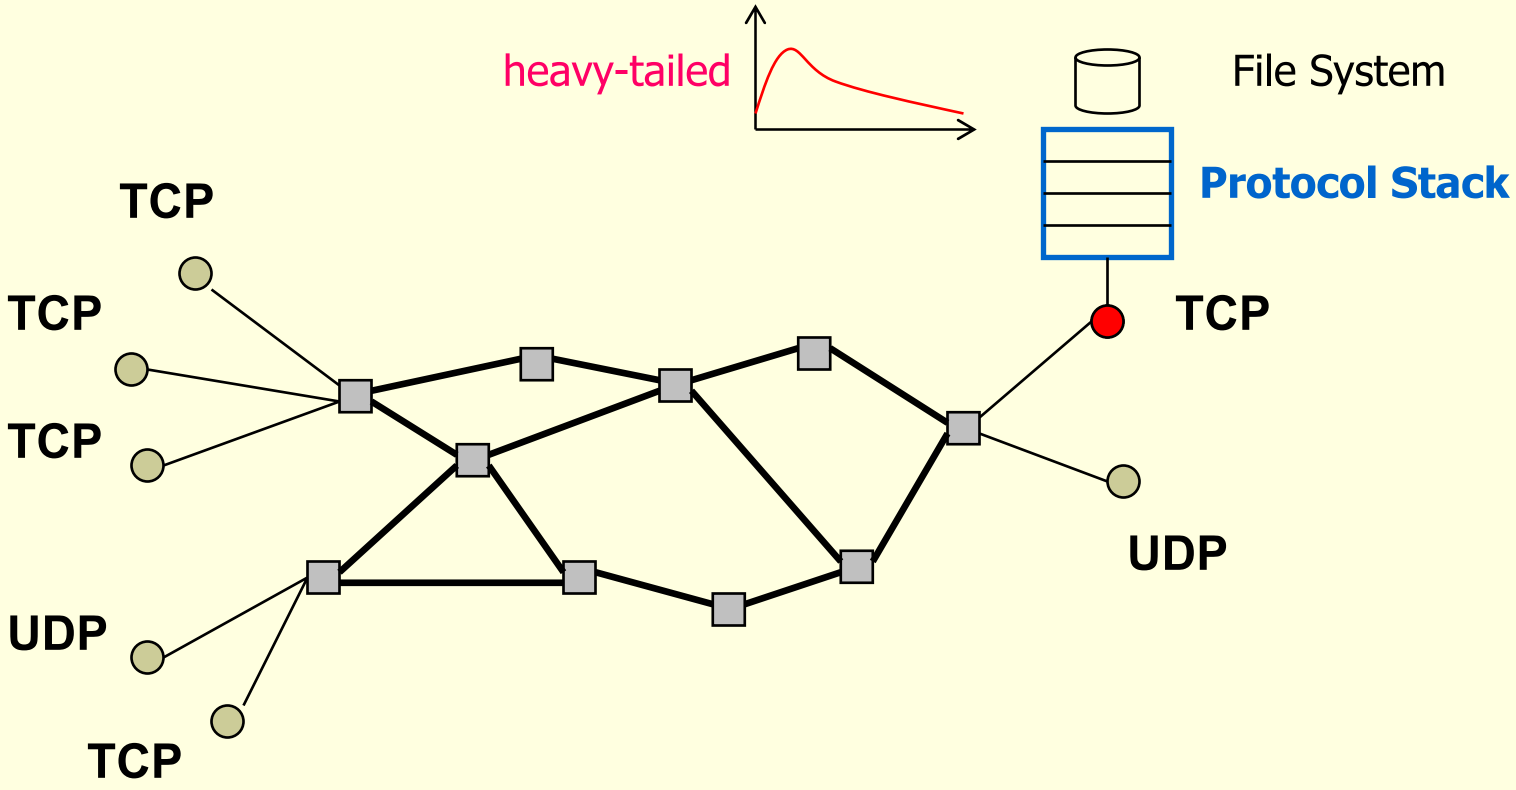

Impact of Protocol Stack

- TCP preserves heavy-tailedness
  - stretching-in-time vs. stretching-in-space
  - conservation law

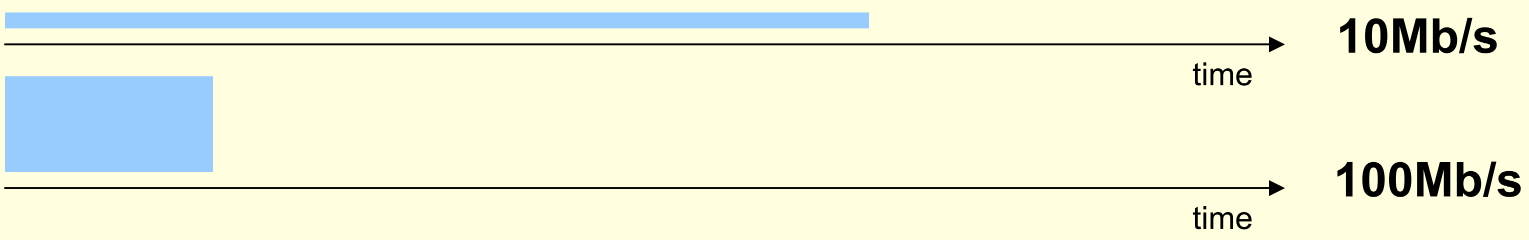

- Trade-off: long-range correlation vs. short-range burst

Impact of Protocol Stack

- Traffic property incorporating feedback control:

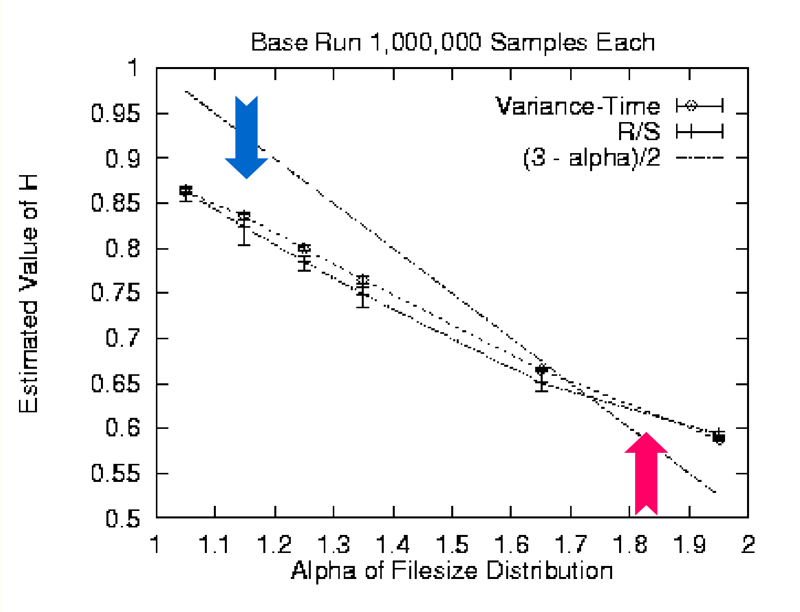

- Slope is less than $H = (3 - \alpha)/2$, why?

Impact of Protocol Stack

- Introduction of non-uniformity

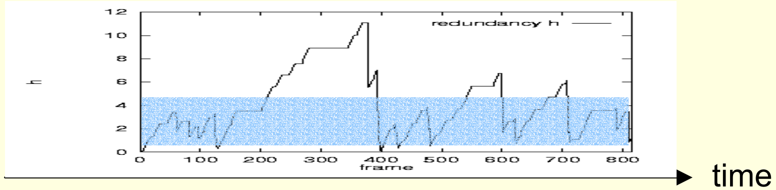

- Introduction of holes

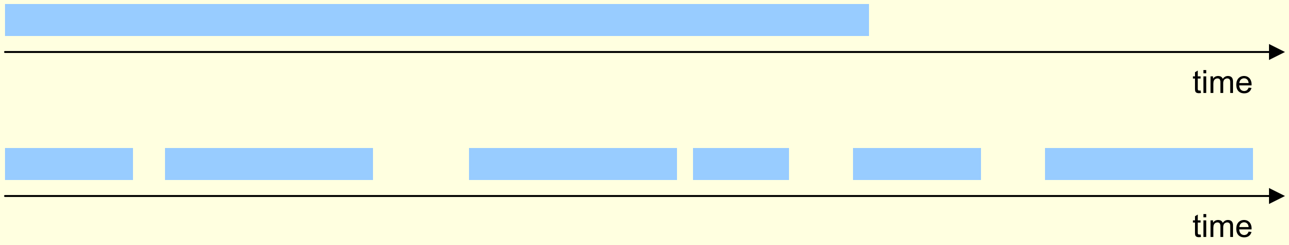

→ lengthening and fragmentation

Impact of Protocol Stack

- Silence periods can be lengthy
  - TCP's exponential back-off
  - Like on/off model with heavy-tailed off periods?

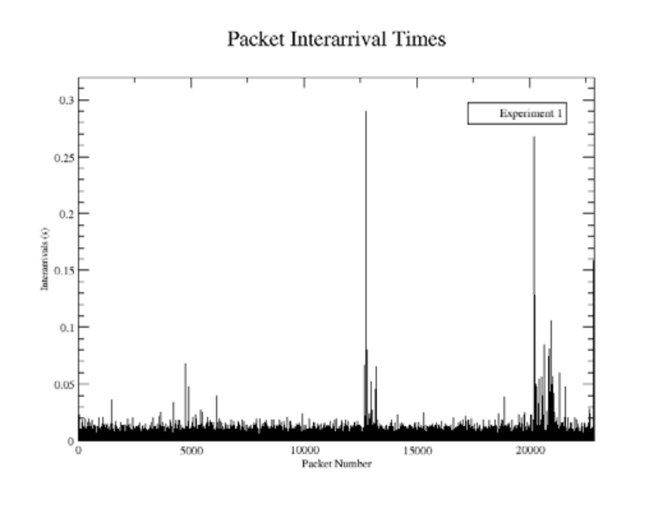

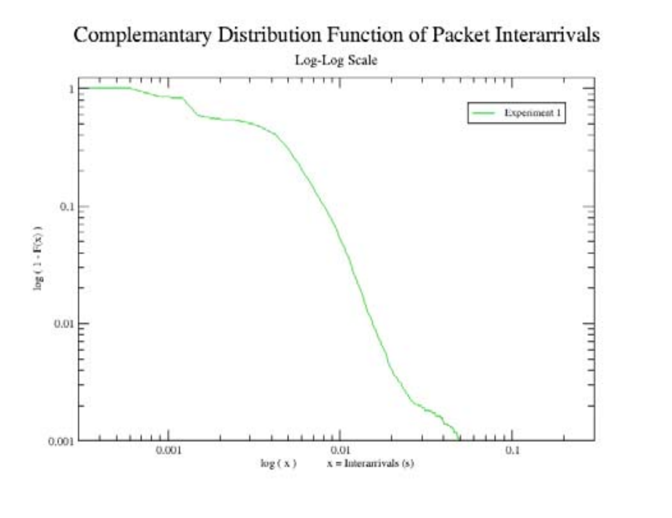

  - And chaotic dynamics? Veres *et al*., '00

Impact of Protocol Stack

- Impact is limited:

on: Pareto / off: Expo; on: Pareto / off: Pareto

on: Expo / off: Pareto

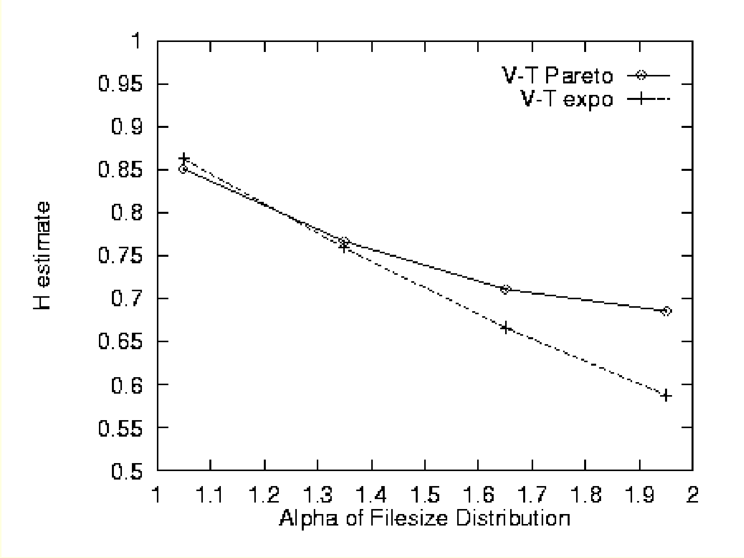

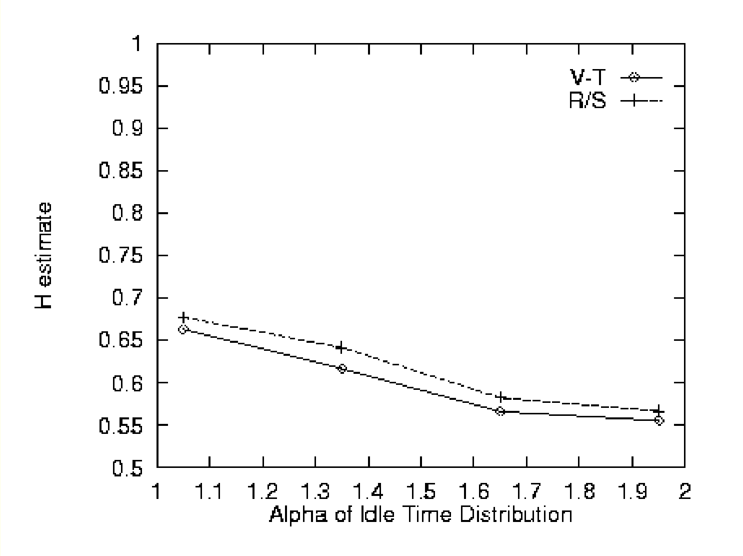

→ can inject correlation; but magnitude secondary

Influence of Topology

- Heavy-tailed Internet connectivity
  - AS graph: Faloutsos *et al*., SIGCOMM '99
  - Web graph: Barabasi's group at Notre Dame Univ.

$$\Pr[\deg(u) > x] \approx x^{-\delta}$$

→ contrast with random graphs: exponential tail

Influence of Topology

- Heavy-tailed Internet connectivity
  - AS graph: Faloutsos *et al*., SIGCOMM '99
  - Web graph: Barabasi's group at Notre Dame Univ.

$$\Pr[\deg(u) > x] \approx x^{-\delta}$$

degree = 601

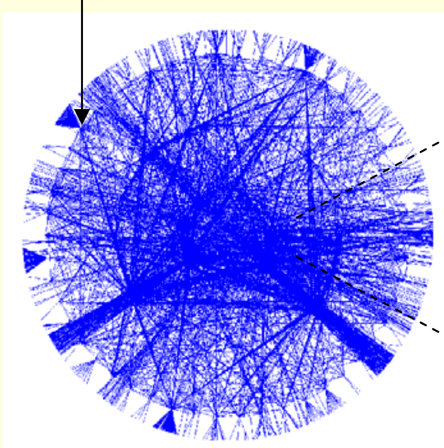

1997 AS graph

random-like

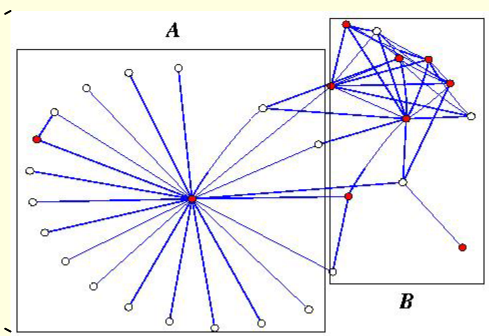

star-like topology

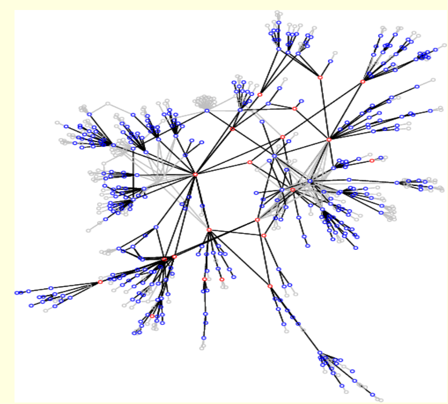

Medium ISP (R. Govindan, USC/ISI)

Influence of Topology

- Connection with network traffic
  - Load of link *e*: no. of paths traversing through *e*
  - $\Pr[\#(e) > x]$ is heavy-tailed

ranked edge load (log-log)

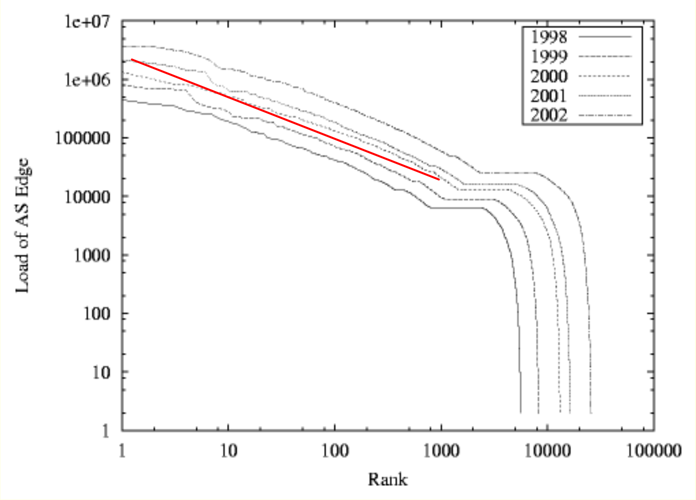

ranked node load (log-log)

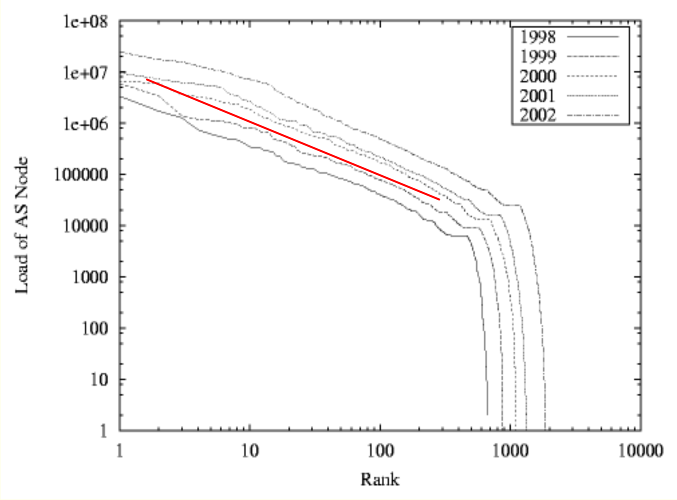

$\rightarrow$ high variability in degree of traffic multiplexing

Influence of Topology

- Coefficient of variation
  - A form of multiplexing gain
  - Dampening reduces impact of time correlation
  - W. Cleveland *et al*., 2000
- Non-uniform bandwidth distribution
  - R. Riedi *et al*., 2001
  - Backbone traffic: spikey “alpha” + Gaussian
  - Multifractal?

→ topology can influence observed traffic

Key Points

- Poisson arrivals, heavy-tailed duration
- Refined workload models
- Structural cause of self-similar burstiness: heavy-tailed file sizes
- Protocol stack and other effects are secondary

Performance Evaluation

Performance Evaluation: Queueing

- Fundamental result:
  - $\rightarrow$ subexponential queue length distribution

- FGN: Weibull
  - I. Norros, Queueing Systems, ‘94

- $M/G/\infty$: polynomial
  - Likhanov *et al*., INFOCOM ‘95

Performance Evaluation: Queueing

- Limited effectiveness of buffering:

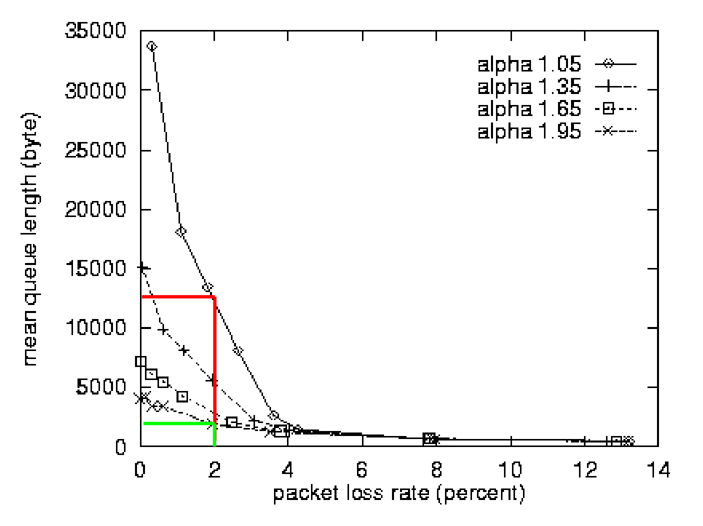

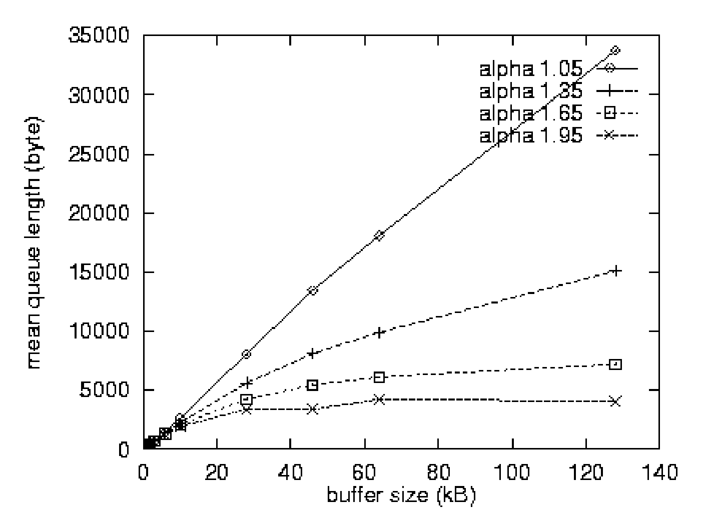

  - Increase buffer capacity
  - QoS trade-off: excessive delay penalty

    → large bandwidth/small buffer policy

Why Heavy-Tailed Queue Tail?

- Consider single on/off process:

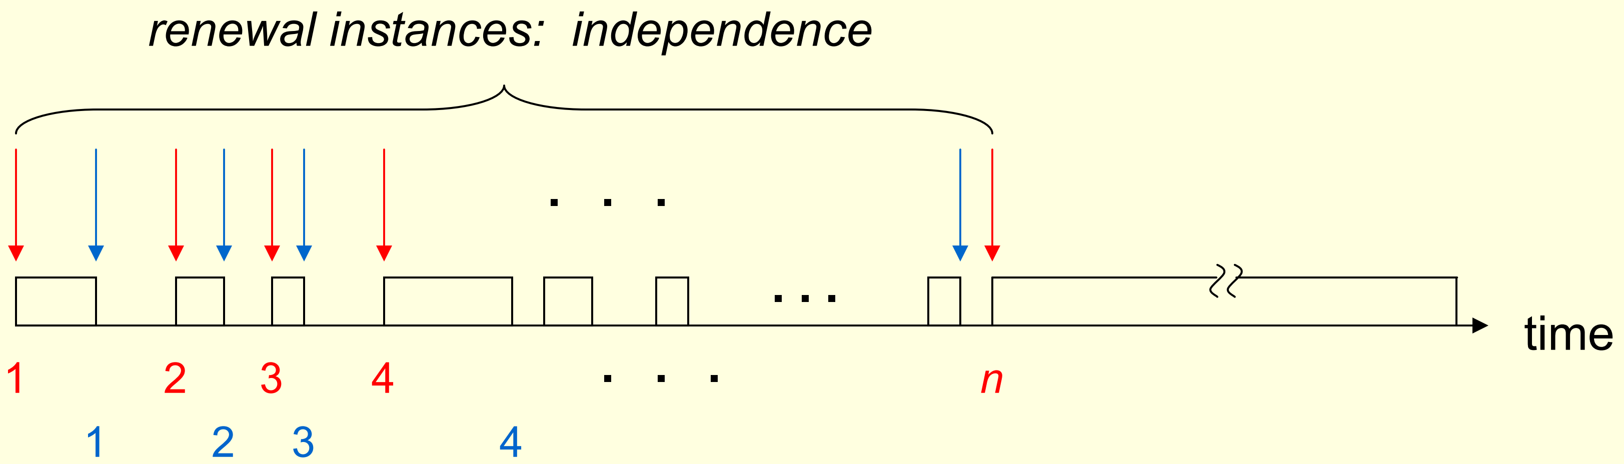

X_1 X_2 X_3 X_4 . . . X_n : i.i.d. heavy-tailed

Y_1 Y_2 Y_3 Y_4 . . . : i.i.d. short-tailed

Why Heavy-Tailed Queue Tail?

- Want to know: $\Pr[Q > b]$?
  - In equilibrium
  - For large buffer level *b*
- Idea:

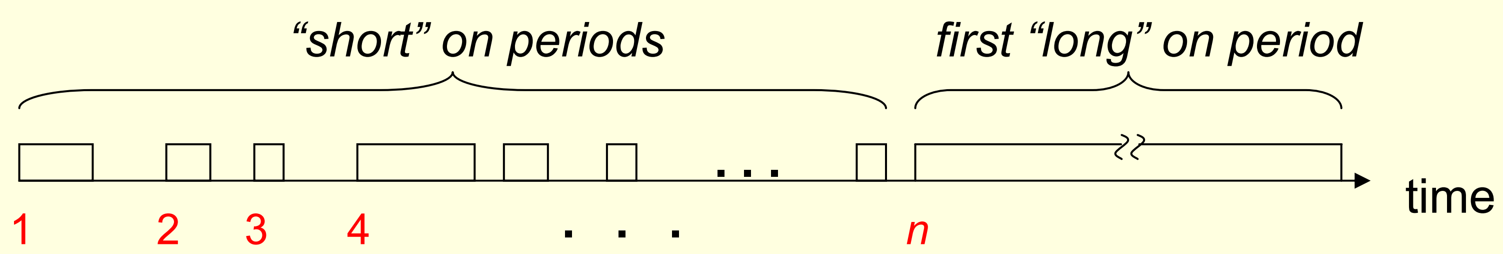

$\rightarrow$ i.e., "long" such that $X_n > b$

Why Heavy-Tailed Queue Tail?

- More precisely: $X_n > (1-\mu)^{-1} b = b'$
  - $\rightarrow$ μ : service rate
- Since: during "long" on-period X_n
  - $\rightarrow$ queue build-up is at least (1-μ) *b'*
  - $\rightarrow$ i.e., event $Q_t > b$ occurs
- In the following ignore *b*, *b*' distinction

Why Heavy-Tailed Queue Tail?

- Ignore queue dynamics before long on-period

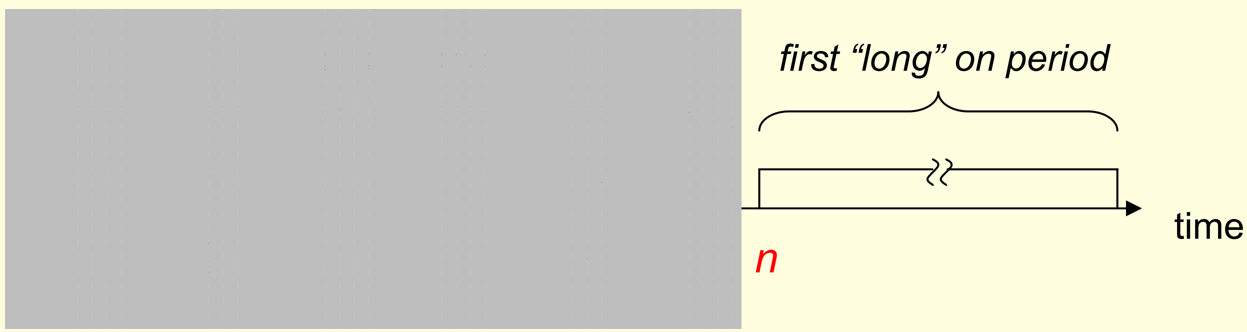

Q_t : **don’t care** **care**

- Hence, lower bound on $\Pr[Q > b]$

Why Heavy-Tailed Queue Tail?

- Need to estimate total time $T = S + L$

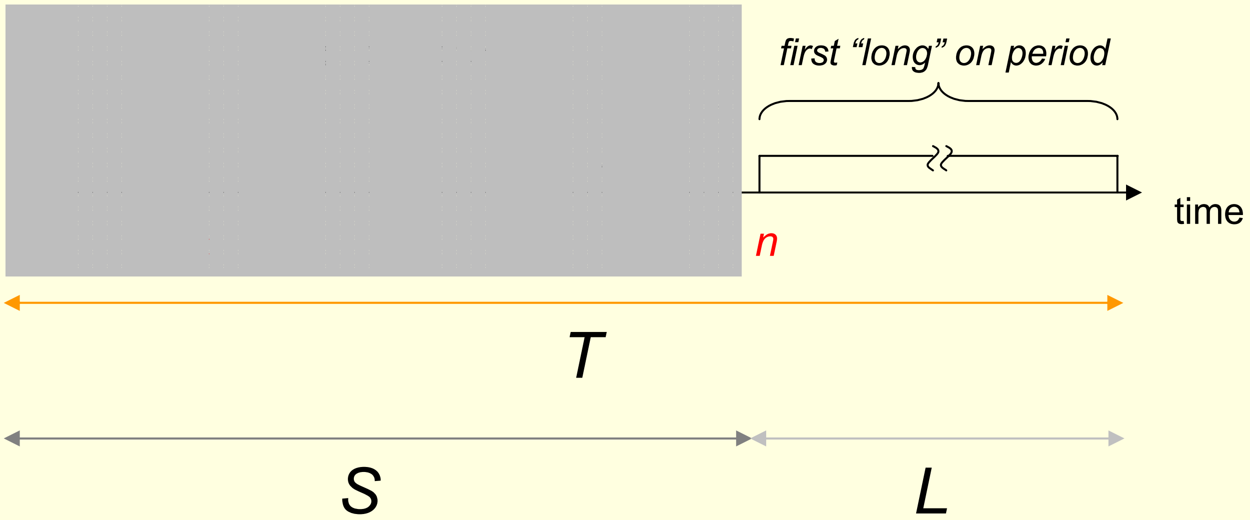

- If ergodic, estimate $\Pr[Q > b] \geq \dfrac{(1-\mu)E[L]}{E[S]+E[L]}$

Why Heavy-Tailed Queue Tail?

- Assume exponential off period, Pareto on-period
  - Exponential λ_{off}
  - Pareto: shape parameter α, location parameter *k*

$\rightarrow$ pdf $f(x) = \alpha k^{\alpha} x^{-(1+\alpha)}$

- Assume stability
- For large *t* and *b*: $$E[S] \approx (n-1)\left(\frac{1}{\lambda_{off}} + \frac{k\alpha}{\alpha - 1}\right)$$

Why Heavy-Tailed Queue Tail?

- To estimate expected *L*, note
  - *L* is conditioned on *n* such that *L* > *b*
  - $E[L] = E[X_n \mid X_n > b]$
- Easy to check

$$E[X_n \mid X_n > b] = \int_0^\infty (x+b) \underbrace{\frac{\alpha b^\alpha}{(x+b)^{\alpha+1}}}_{\text{conditional probability}} dx = \frac{\alpha}{\alpha - 1} b$$

Why Heavy-Tailed Queue Tail?

- Almost done.
- Note: since $\Pr[X_n > b] = (b/k)^{-\alpha}$

$$\therefore \quad n \approx (b/k)^{\alpha}$$

- Combining everything

$$\Pr[Q > b] = \Omega(b^{1-\alpha})$$

$\rightarrow$ heavy-tailed on time leads to heavy-tailed queue tail

Remarks

- Demonstrates connection between heavy-tailed workload and heavy-tailed queue length
- Similar ideas apply to $M/G/\infty$ and more general on/off input models
- Sketch of key ideas; not a proof
- Applies to upper bound: "short" & "long" picture

$$\Pr[X_1 + \cdots + X_n > b] = \Pr[\max\{X_1, \ldots, X_n\} > b]$$

$\rightarrow$ key property of heavy-tailed, i.e., regularly varying r.v.

Impact of SRD vs. LRD Traffic

- Relative importance of short-range dependent vs. long-range dependent model
- Debate: LRD needed for traffic modeling and performance evaluation?
  - SRD models can effectively capture input traffic
  - Well-understood performance evaluation
- Finite time scale and resource dimensioning

Impact of SRD vs. LRD Traffic

- SRD vs. LRD packet loss:

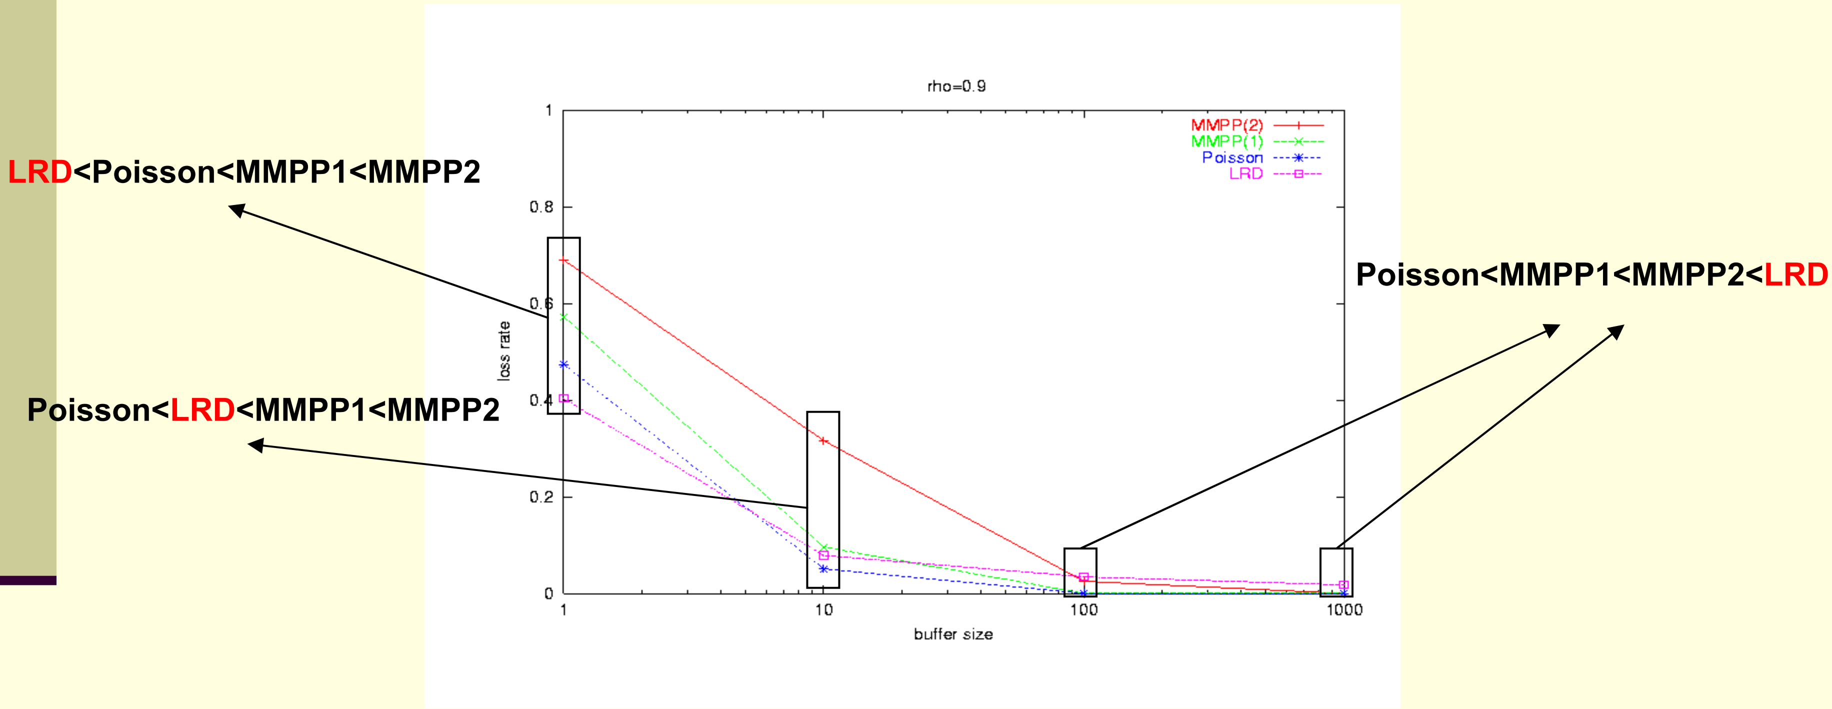

  - Depends on details of resource configuration
  - SRD vs. LRD debate: no uniform answer

Impact of SRD vs. LRD Traffic

- Main differences:
  - Physical vs. “black box” time series modeling
    - → pros & cons
    - → depends on objectives
  - Physical models are useful for closed-loop traffic control evaluation
    - → time series models: open-loop
  - For queueing application: little essential difference

Self-Similar Burstiness and Jitter

- Focus has been on first-order performance measure
- Relevance of second-order measure
  - Real-time multimedia data
    - $\rightarrow$ small loss rate: insufficient
  - Packet-level forward error correction
    - $\rightarrow$ correlated losses: the enemy

Self-Similar Burstiness and Jitter

- Block loss
  - Network traffic $X(t)$: sequence of packets
  - Block size N

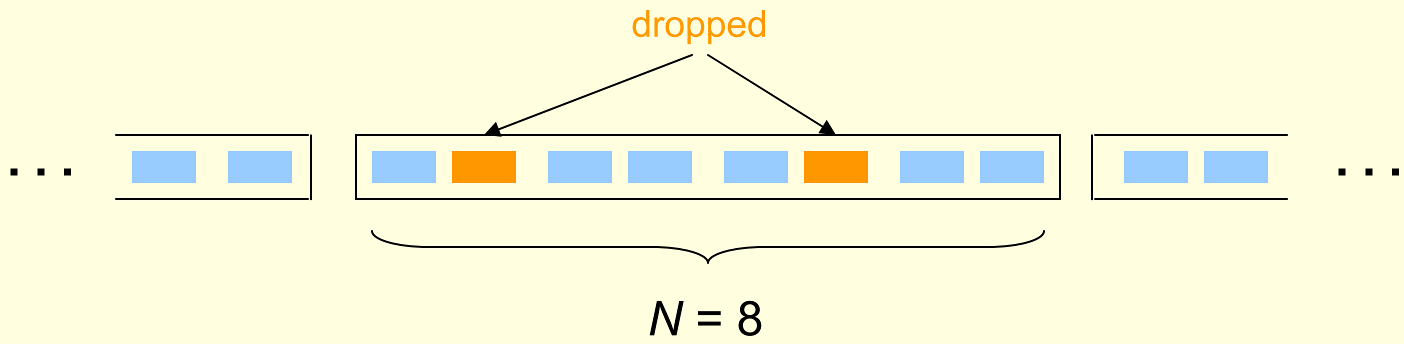

- Block loss process $B(n)$

  $\rightarrow$ no. of losses in n'th block

Self-Similar Burstiness and Jitter

- Block loss distribution in steady state

  $\rightarrow$ $B \equiv B(\infty)$: r.v. with B $\in$ {0, 1, …, *N*}

- Normalized block loss distribution

  $\rightarrow$ B / N : r.v. with values in [0,1]

- Assuming FEC satisfying *k*-out-of-*N* property is used

  $\rightarrow$ the heavier the tail $\Pr[B > x]$, $x > k$, the less effective FEC is

Self-Similar Burstiness and Jitter

- Block loss behavior

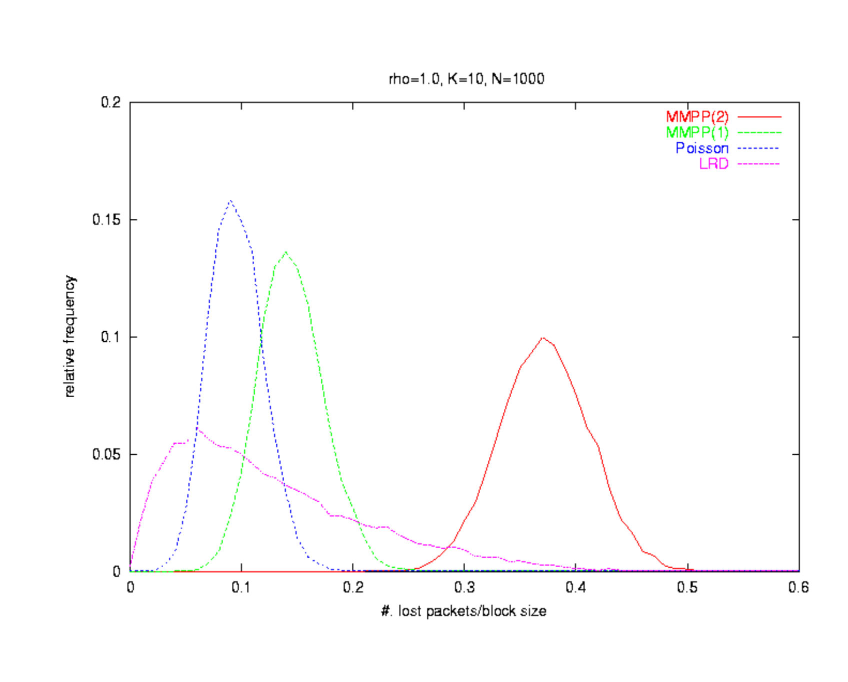

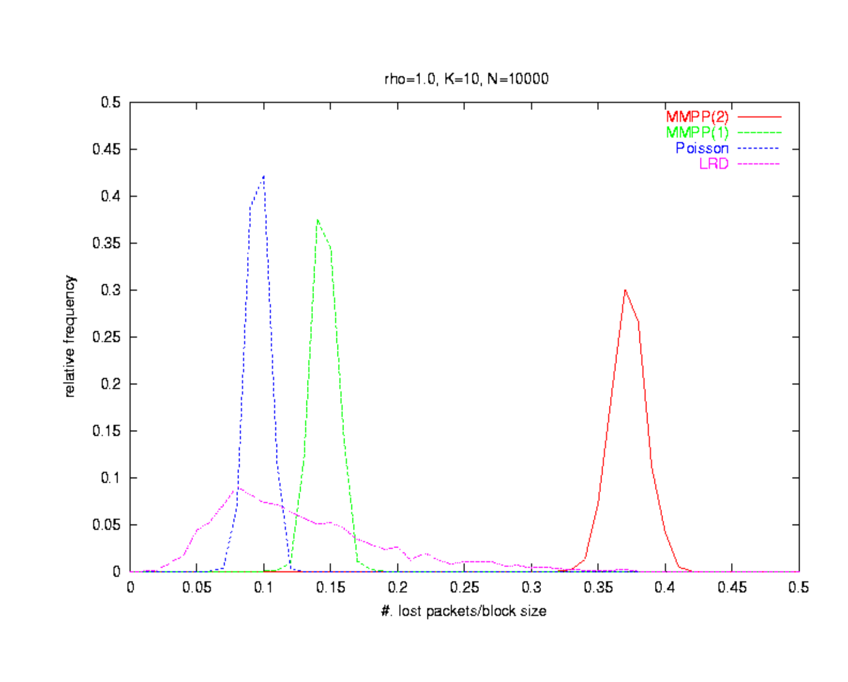

- Loss rate: LRD < Poisson < MMPP1 < MMPP2
- Block loss variance: LRD dominant

Slow Convergence and Simulation

- Sampling from heavy-tailed distribution
  → slow convergence of sample mean to population mean

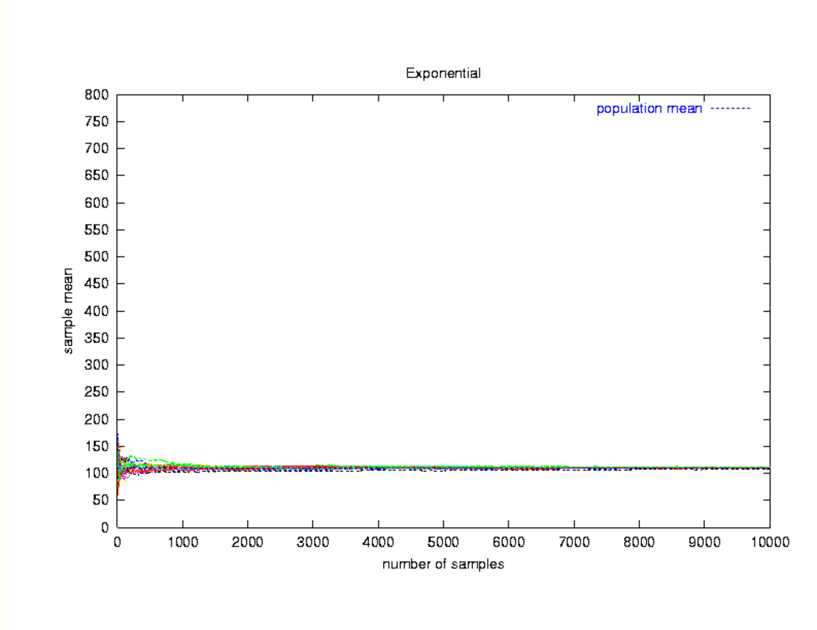

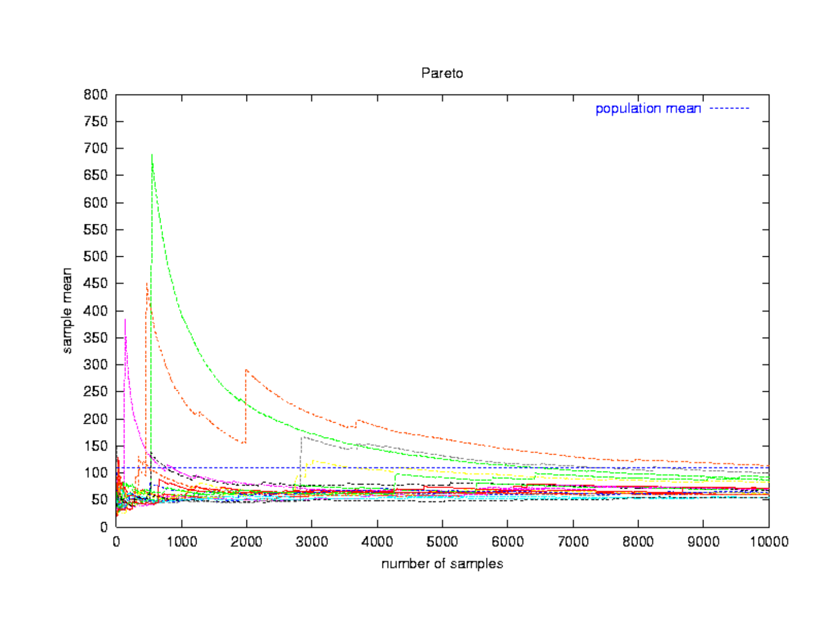

→ running mean of 20 sample paths

Slow Convergence and Simulation

- Sampling from heavy-tailed distribution
  → slow convergence of sample mean to population mean

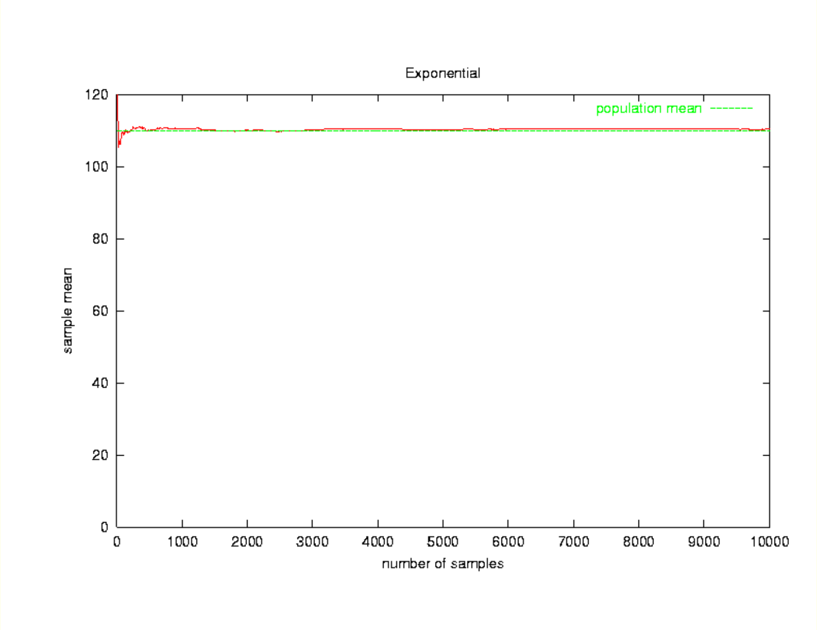

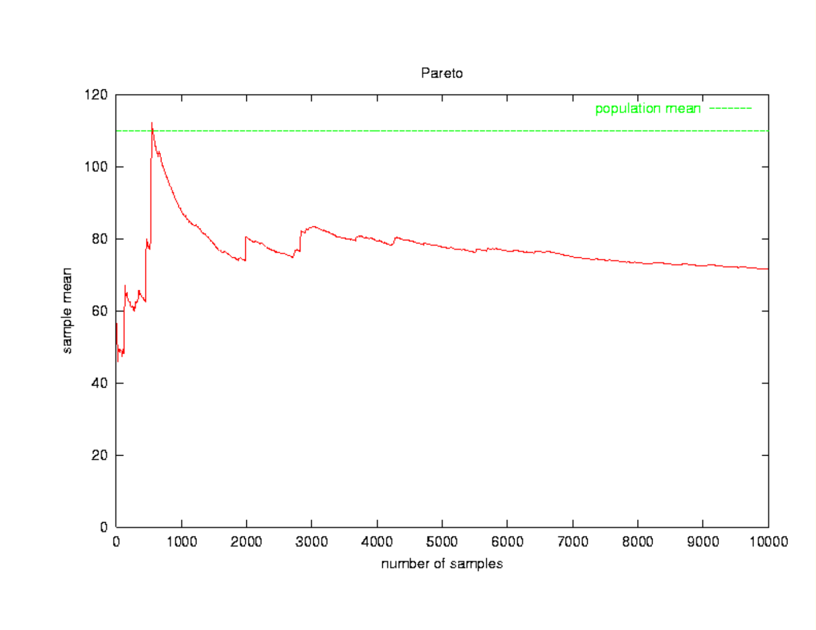

→ average of 20 sample path running means

Slow Convergence and Simulation

- Approximating the population mean:
  - Pareto: pdf $f(x) = \alpha k^{\alpha} x^{-(1+\alpha)}$
  - Shape parameter α, location parameter k
- Population mean of Pareto r.v. Z

$$E[Z] = \underbrace{\int_k^y x f(x) dx}_{A(y)} + \underbrace{\int_y^\infty x f(x) dx}_{B(y)} = \frac{k\alpha}{\alpha - 1}$$

Slow Convergence and Simulation

- Want *y* such that $B(y)/E[Z] < \varepsilon$

$$\therefore \quad y > y_0 = k\left(\frac{1}{\varepsilon}\right)^{\frac{1}{\alpha-1}}$$

- Thus $\Pr[Z > y_0] = \varepsilon^{\frac{\alpha}{\alpha-1}}$, and

$$\therefore \quad \text{no. of samples} \approx \left(\frac{1}{\varepsilon}\right)^{\frac{\alpha}{\alpha-1}}$$

Slow Convergence and Simulation

- For truncated sampling: $Z_i > y_0 \Rightarrow Z_i = 0$
- Sample mean within accuracy ε:

$$\frac{\bar{Z}_n}{E[Z]} \geq 1 - \varepsilon$$

- Likely occurrence for

$$n \geq const \times \left(\frac{1}{\varepsilon}\right)^{\frac{\alpha}{\alpha-1}}$$

Slow Convergence and Simulation

- For example:
  - $\alpha = 1.2$; $H = (3 - \alpha) / 2 = 0.9$
  - $\varepsilon = 0.01$
  - sample size greater than 10 billion
- Practically:
  - Brute-force is problematic
  - Speed-up methods required
  - Rare event simulation

Key Points

- Heavy-tailed workload and queue tail
- Impact of LRD on loss performance is mixed
- Impact of LRD on second-order performance measure (“jitter”) is more clear cut
- Convergence and simulation pose significant challenges

Traffic Control

Workload-Sensitive Traffic Control

- Self-similar burstiness
  - Bad news: queueing
    - $\rightarrow$ heavy-tailed queue tail

  - Good news: predictability
    - $\rightarrow$ facilitates traffic control

Workload-Sensitive Traffic Control

- Approach:
  - → exploit predictability in the workload

  - Simple: heavy-tailed life time distribution
    - → optimistic congestion control

  - More complex: large time scale traffic correlation
    - → predictive control

Heavy-Tailed Life Time

- Heavy-tailedness implies predictability
  - Assume *Z* is heavy-tailed (e.g., Pareto)
  - *Z* represents life time or connection duration

- Easy to check:

$$\Pr[Z > \tau + h \mid Z > \tau] = \left(\frac{\tau}{\tau + h}\right)^{\alpha}$$

$\rightarrow$ as τ increases $\Pr[Z > \tau + h \mid Z > \tau] \rightarrow 1$

$\rightarrow$ conditioning on past helps

… compare with exponential r.v. where $e^{-\lambda h}$

Heavy-Tailed Life Time

- Can make prediction error arbitrarily small by conditioning on longer past

- Expected conditional life time duration

$$E[Z \mid Z > b] = \frac{\alpha}{\alpha - 1} b$$

- For example: if $\alpha = 1.2$, $E[Z \mid Z > b] = 6b$;
  if $\alpha = 1.1$, expected future lifetime 11*b*

Heavy-Tailed Life Time

- Applications:
  - Dynamic load balancing; Harchol-Balter & Downey '96
  - Routing stability; Shaikh *et al*. '99
  - Task scheduling; Crovella *et al*., '99
  - Optimistic congestion control; Park *et al*., '02

Large Time Scale Predictability

- Condition future on past traffic level

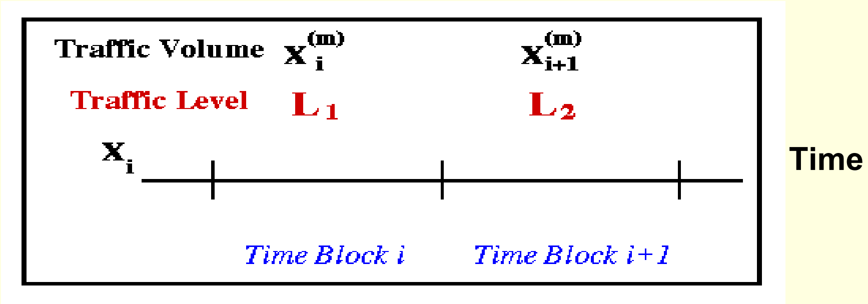

- Conditional expectation estimator

$$E[X_{i+1}^{(n)} \mid X_i^{(n)} = x]$$

- Quantized estimator $E[L_{i+1}^{(n)} \mid L_i^{(n)} = c],\ c = 1, \ldots, c_{\max}$

Large Time Scale Predictability

- Example:
  - LRD traffic

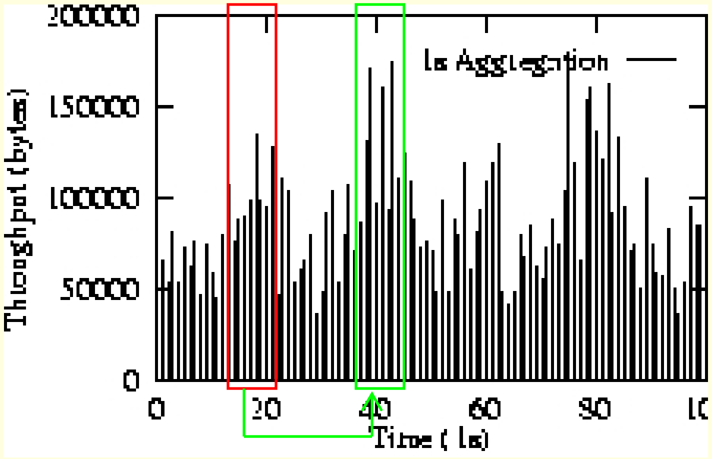

  - Conditional probability

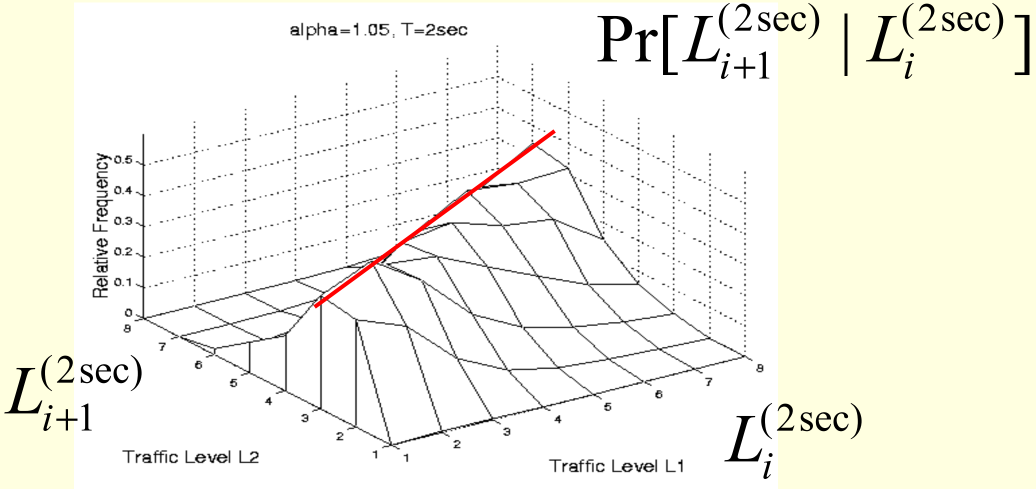

→ **conditioning helps**

Large Time Scale Predictability

- Example:
  - SRD traffic

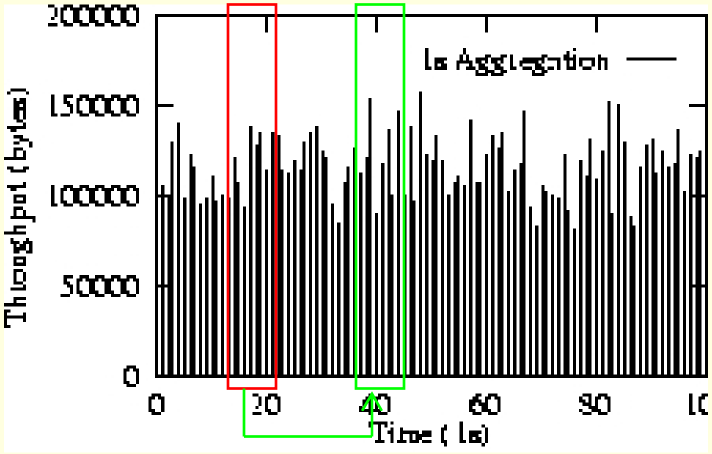

  - Conditional probability

$\Pr[L_{i+1}^{(2\,\mathrm{sec})} \mid L_i^{(2\,\mathrm{sec})}]$

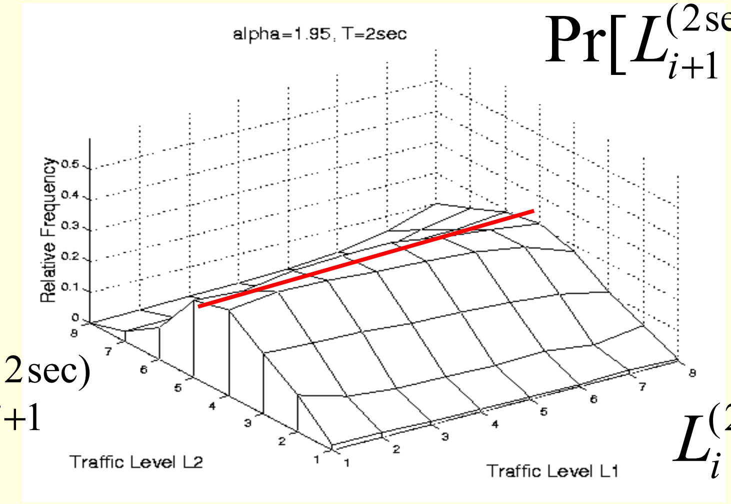

$L_{i+1}^{(2\,\mathrm{sec})}$ $L_i^{(2\,\mathrm{sec})}$

→ **conditioning doesn't help**

Selective Slope Control

Congestion control: TCP and rate-based

Idea:

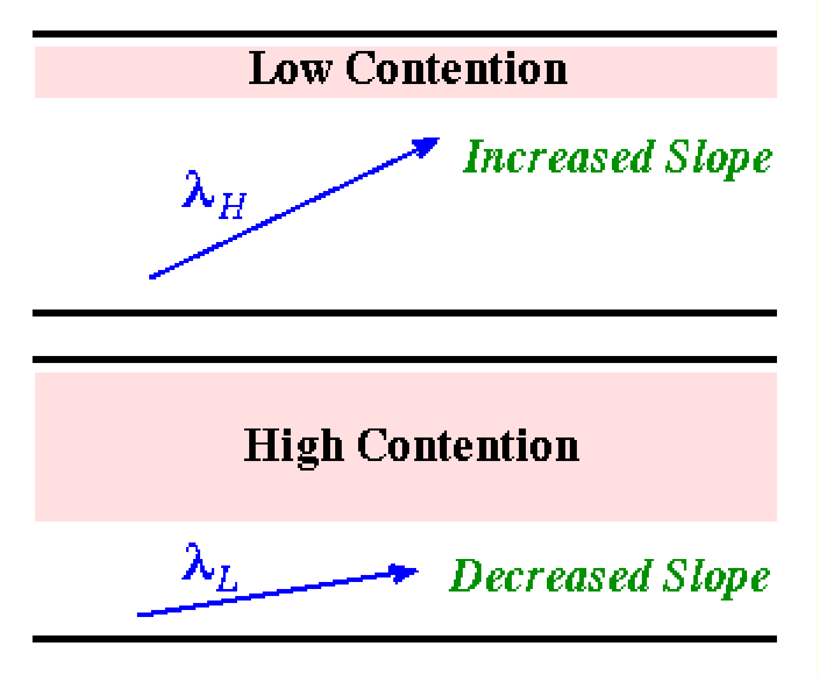

"aggressive"

→ **modulate slope of linear increase phase in AIMD**

"conservative"

Selective Slope Control

- Linear increase phase of AIMD:

$$cwnd \leftarrow cwnd + \frac{1}{cwnd} \quad \Rightarrow \quad cwnd \leftarrow cwnd + \frac{A}{cwnd}$$

- *A* is a control variable:

$$A = A(E[L_{i+1}^{(n)} \mid L_i^{(n)} = c])$$

  - Selective aggressiveness schedule
  - Monotone

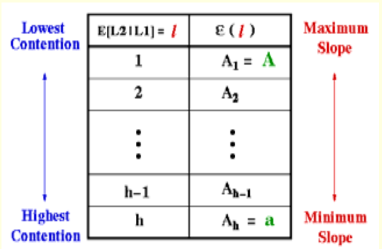

E[L2 | L1] = l	ε (l)
1	A_1 = A
2	A_2
⋮	⋮
h−1	A_{h-1}
h	A_h = a

Structure of TCP-MT: Modularity

- Multiple time scale TCP: TCP-MT

Large time scale
selective slope modulation

TCP feedback control

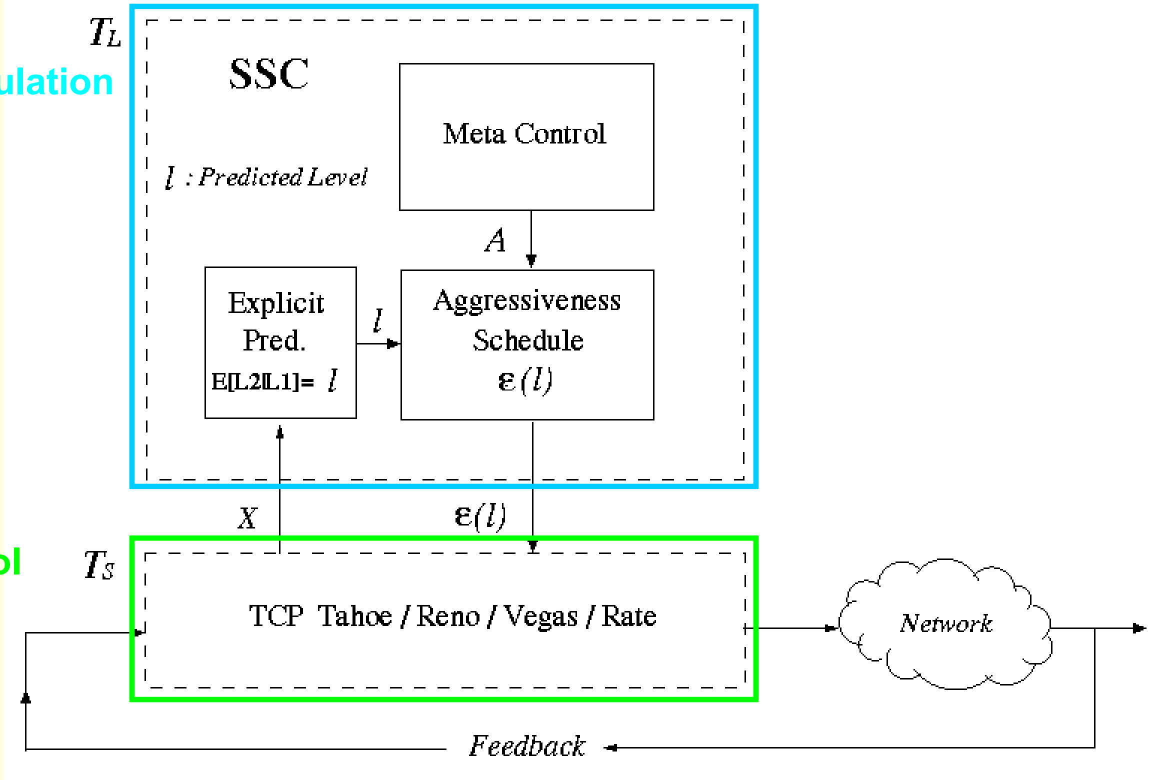

Available Bandwidth Estimation

- Passive probing:
  - Use output behavior of TCP sender
  - Coupled with background traffic

    → tracking ability

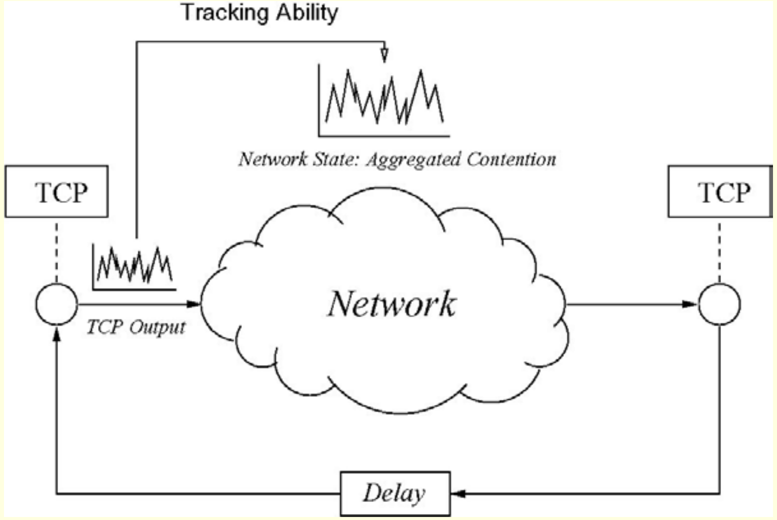

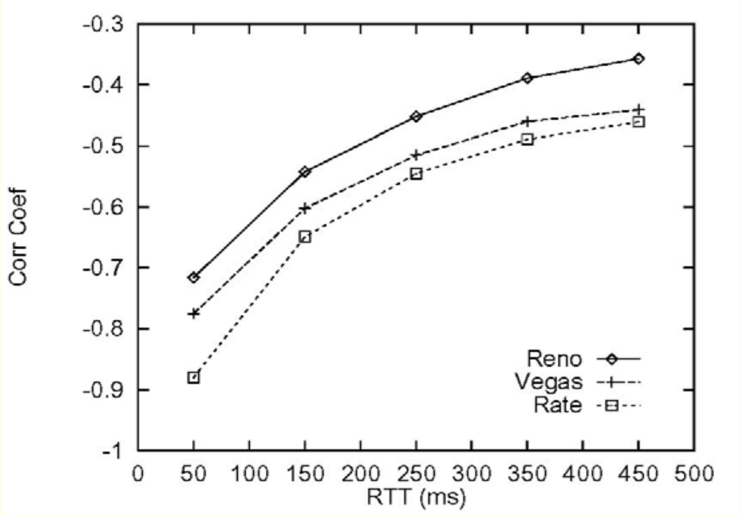

→ negative coupling

→ corr. coeff. as ftn. of RTT

Application

- Mitigate reactive cost of feedback control:
  - → large delay-bandwidth product
  - Broadband WANs
  - Satellite networks

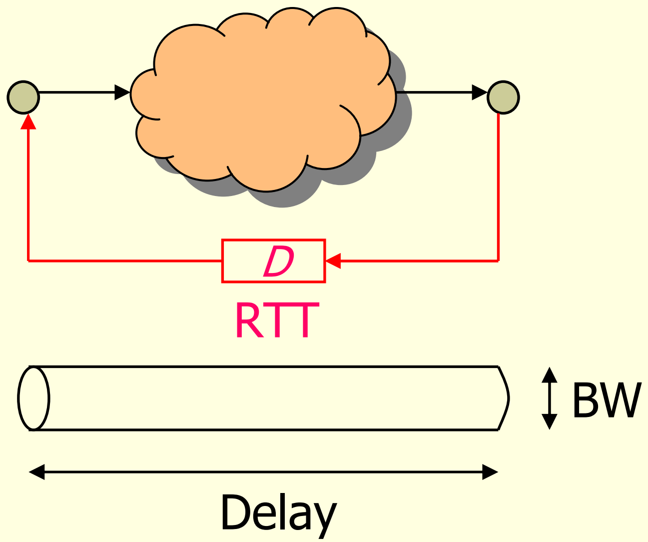

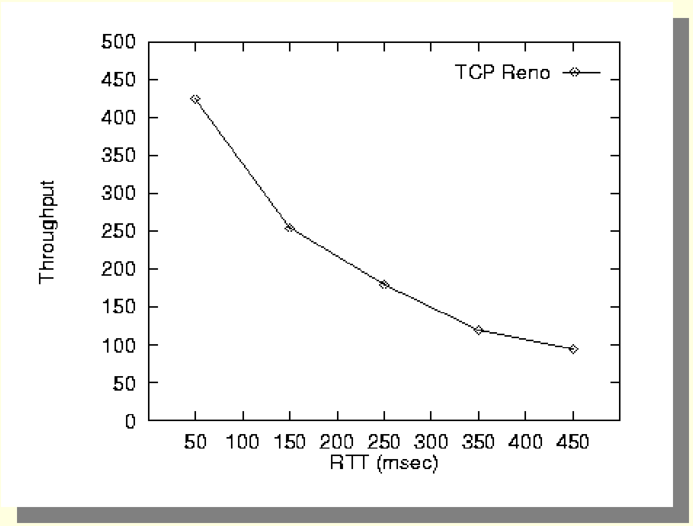

Application

- Delayed feedback
  - Outdated information & control action
  - Stability condition limitation

    $\rightarrow$ delay or functional differential equation

$$0 < \varepsilon \cdot D < \infty$$

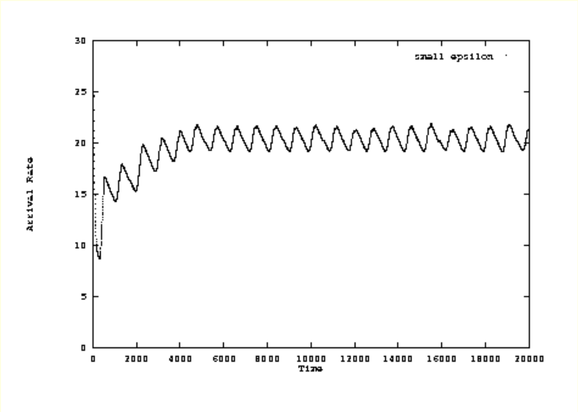

small ε

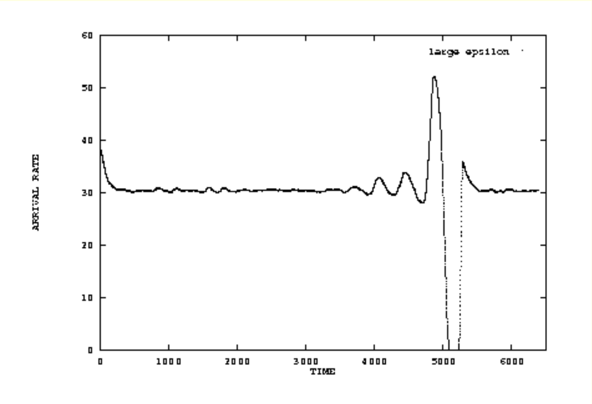

large ε

$\rightarrow$ bounded oscillation

$\rightarrow$ unbounded oscillation: instability

Application

- Large time scale predictability
  - Time scale >> RTT
  - Bridge timeliness barrier

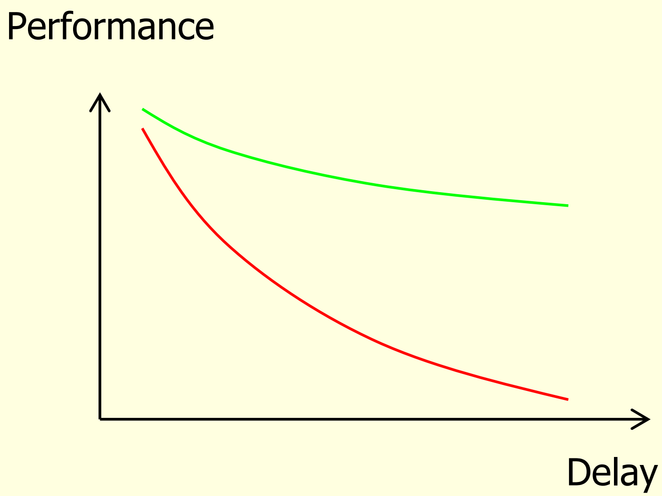

LRD time scale » RTT

Performance Gain

- TCP-MT: Performance gain as function of RTT

$$\frac{\gamma_{TCP} - \gamma_{TCP-MT}}{\gamma_{TCP}}$$

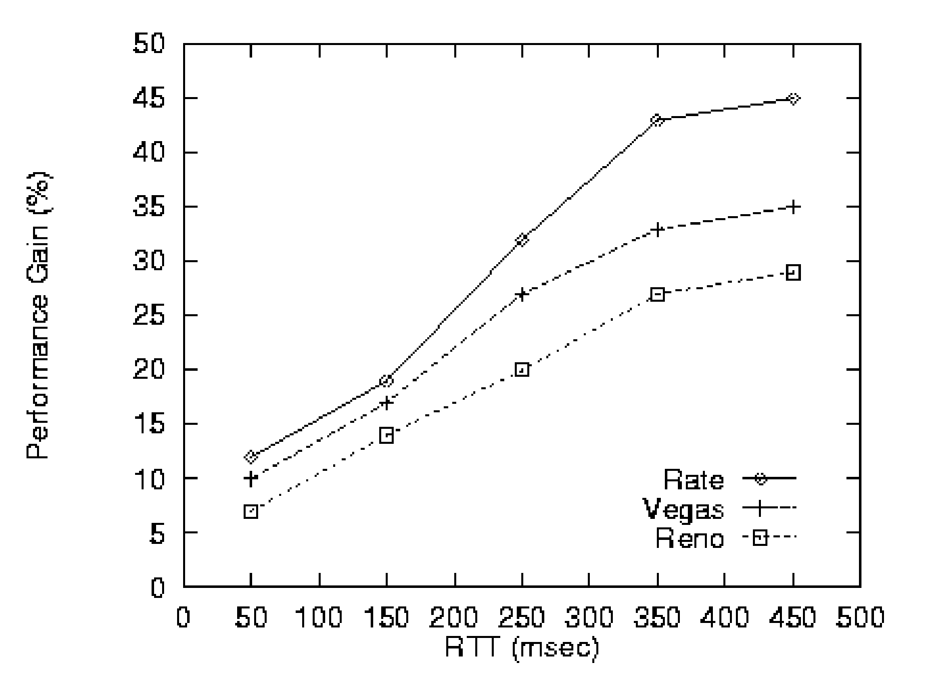

Large Time Scale Predictability

- Multiple time scale redundancy control

  → packet-level FEC: real-time multimedia traffic

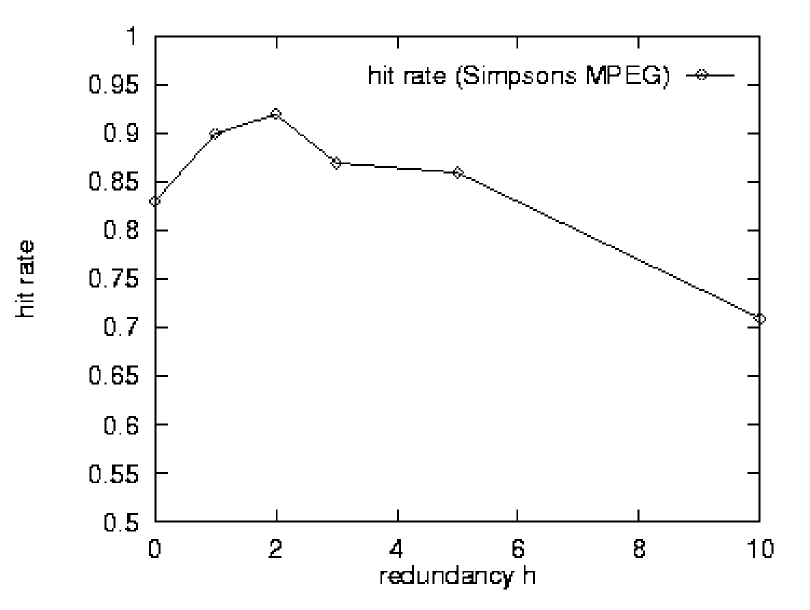

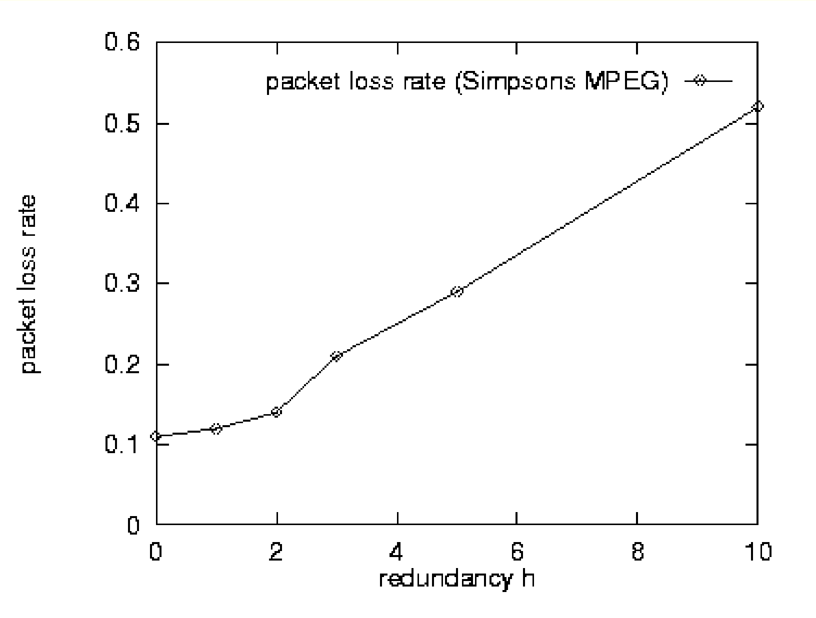

→ **static FEC**

AFEC: Adaptive FEC

- Dynamically adjust redundancy:
  - As a function of network state
  - To achieve target QoS
- Hit rate:
  - Fraction of timely decoded frames
  - Loss or delay

AFEC: Adaptive FEC

- Optimal feedback control problem:
  - User-specified QoS
  - Caveat: too much redundancy counter-productive

    → "shoot oneself in the foot"

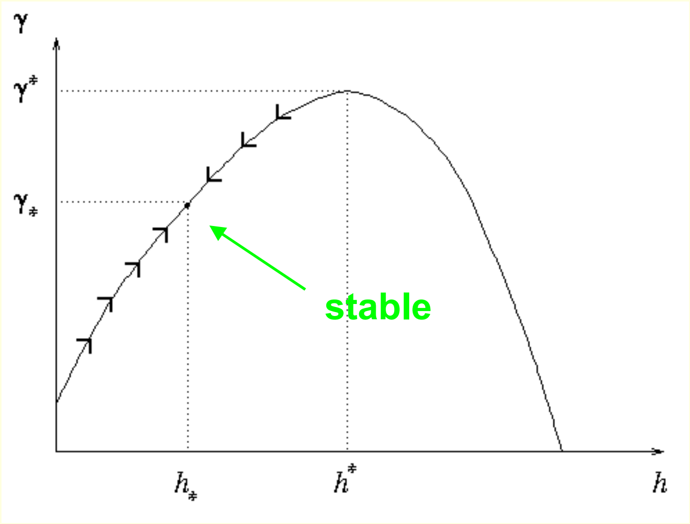

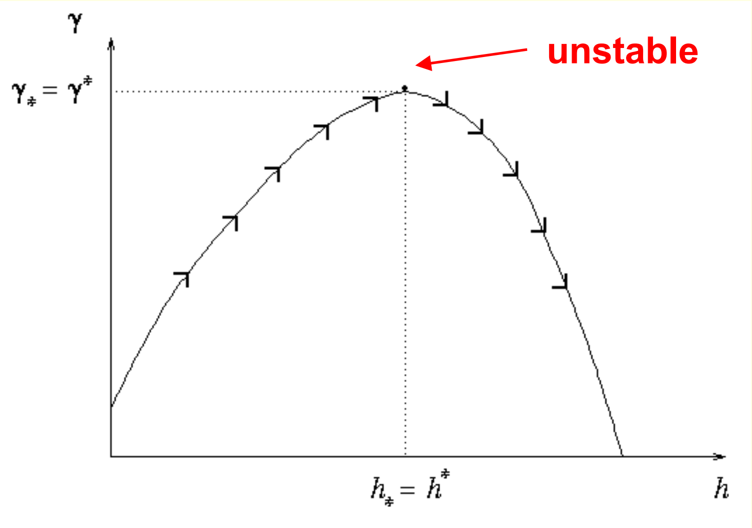

AFEC: Adaptive FEC

- Intrinsic problem of AFEC:
  - Maximum hit rate operating point
    - → unstable: exponential back-off
  - Efficiency vs. QoS trade-off
    - → reduce redundancy when network state good
    - → pay QoS penalty when network state turns bad
    - → sensitive to transients

Selective Level Control

- Level control:

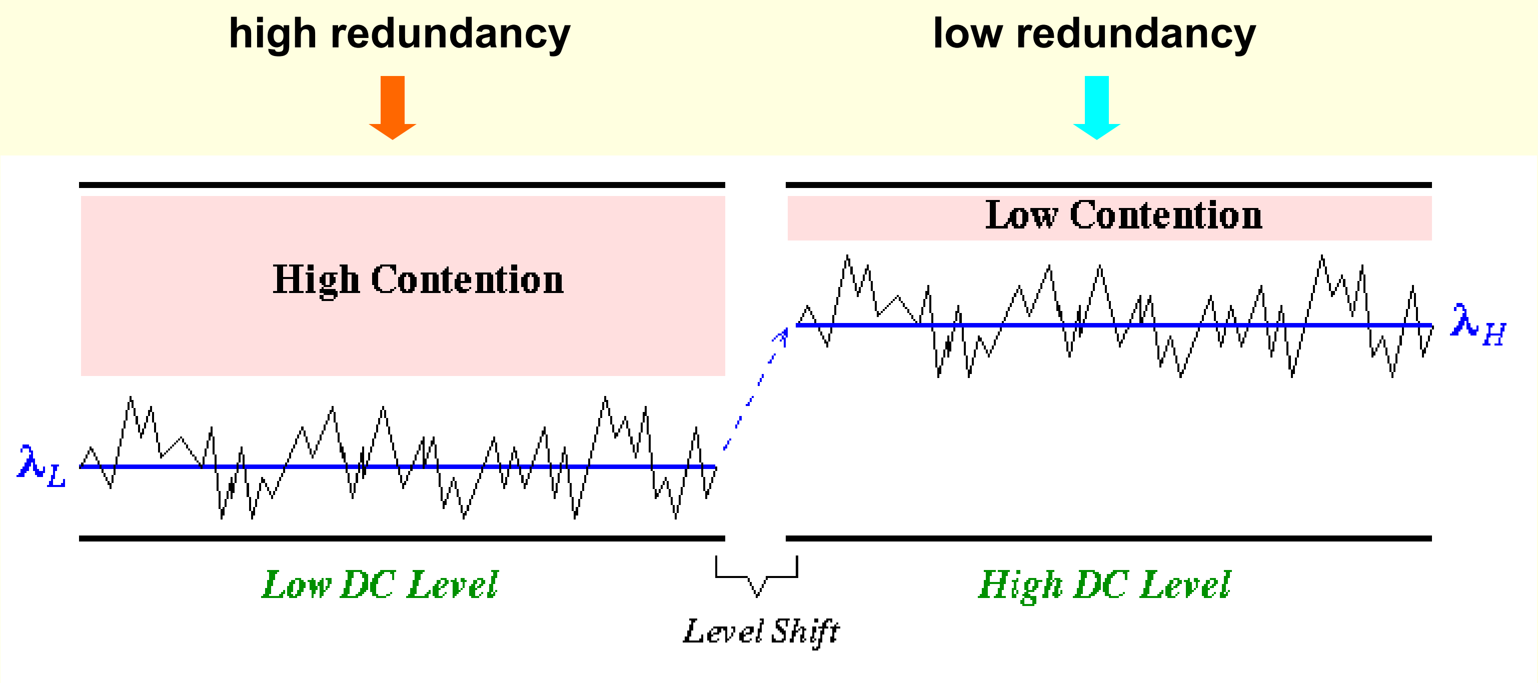

Structure of AFEC-MT: Modularity

- Multiple time scale AFEC: AFEC-MT

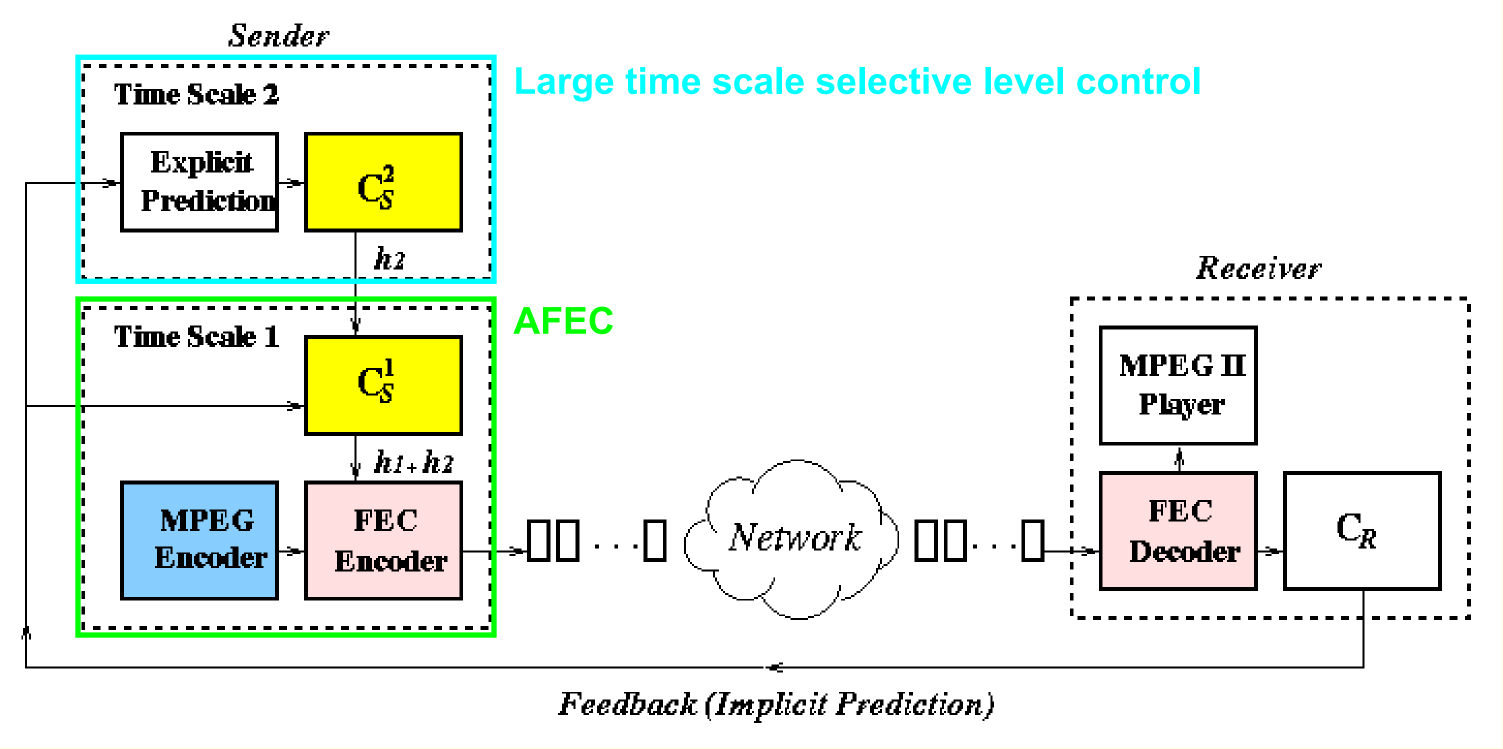

Performance Evaluation: Hit Rate

hit trace:

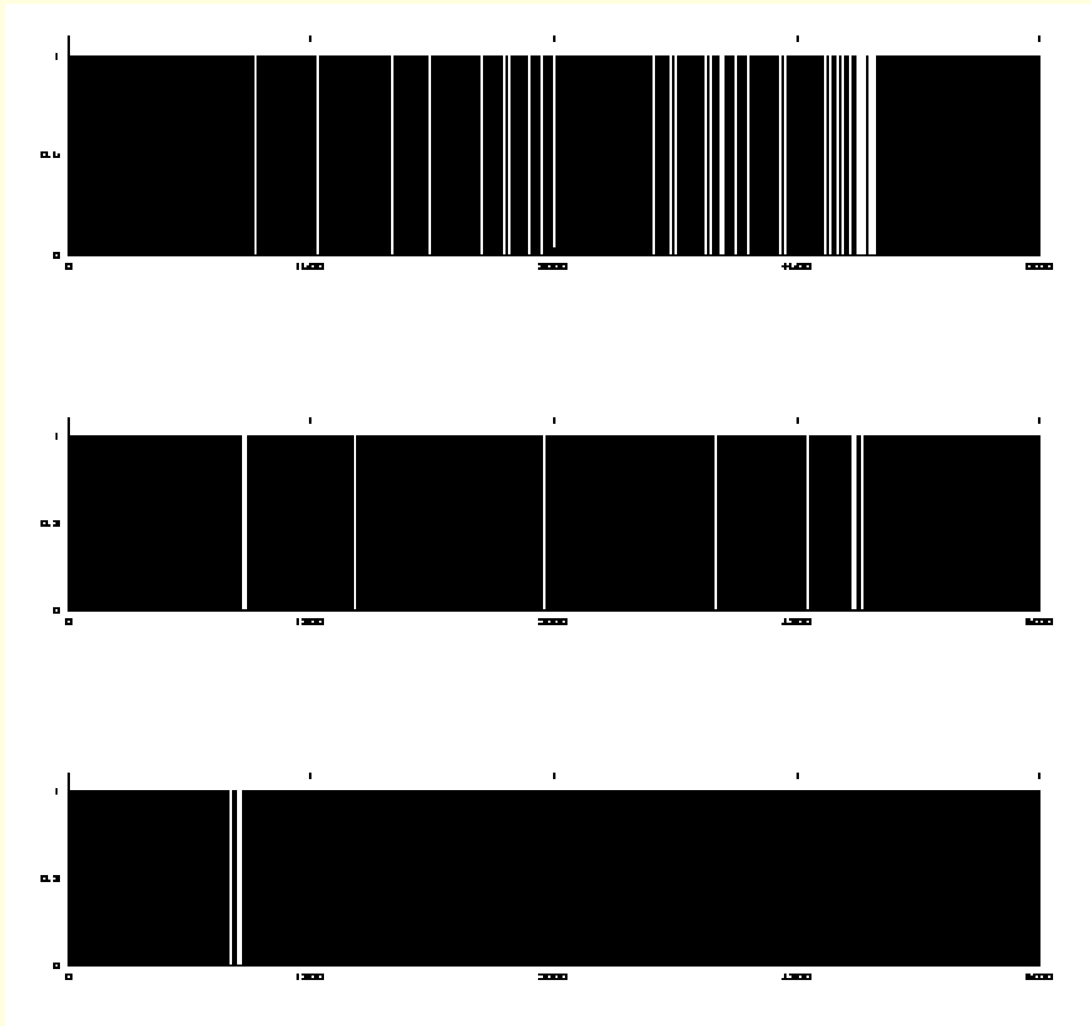

static FEC

AFEC

AFEC-MT

Performance Evaluation: Hit Rate

- Long-term hit rate trace:

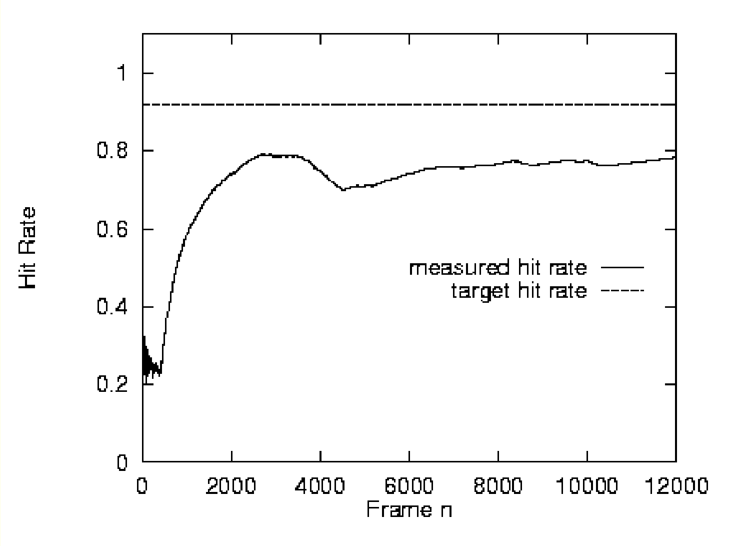

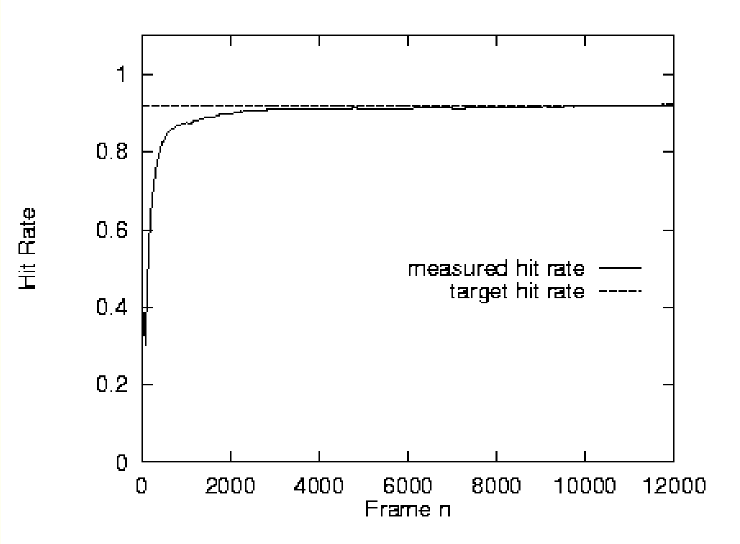

→ AFEC

→ AFEC-MT

Some References

- Tuan & Park, Performance Evaluation ’99
  - → rate-based congestion control
- Tuan & Park, INFOCOM ‘00
  - → AFEC-MT
- Park & Tuan, ACM TOMACS ‘00
  - → TCP-MT
- Östring *et al*., IEEE Trans. Commun. ’01
  - → router assisted rate-based

Key Points

- Workload can be exploited for traffic control
- Heavy-tailed life time: simple
- LRD traffic: more complex
- Significant performance gain possible
- Delay-bandwidth product problem mitigation

Open Problems

… even in the parking lot

Challenges and Open Problems

- Relevance of second-order performance measures
- Rare event simulation with heavy-tailed workload
- Workload-sensitive traffic control
- Finite resource dimensioning

Book Plug

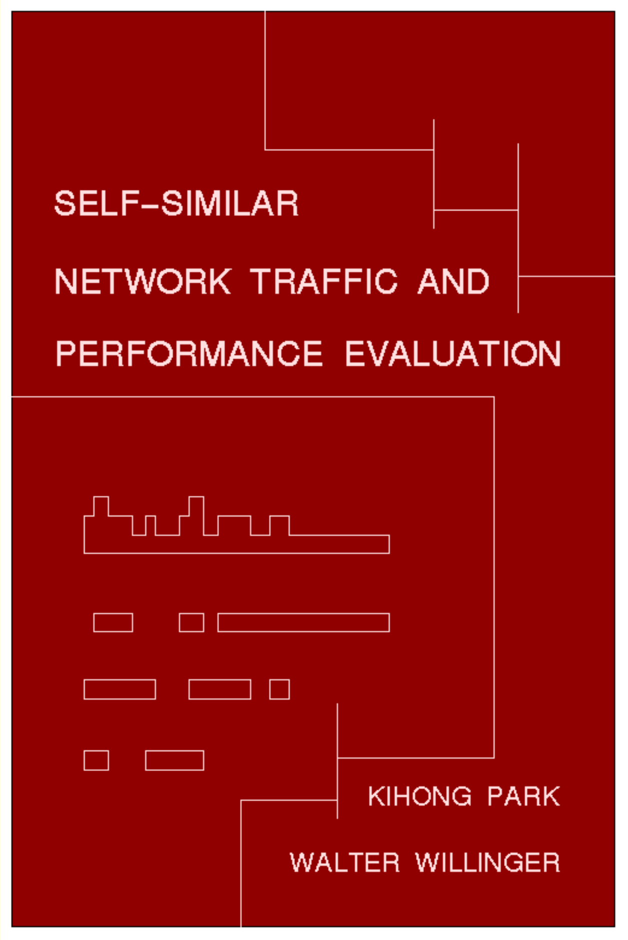

- Wiley-Interscience 2000
- Collection of chapter contributions
- Landscape circa ’00
- Additional references