Capturing Network Traffic Dynamics
Using the Tools

Rolf Riedi

Dept of Statistics

Stochastic Systems and Modelling in Networking and Finance
Part IV
Dependable Adaptive Systems and Mathematical Modeling
Kaiserslautern, August 2006

LRD estimation

Time domain

Spectral domain

Wavelet domain

Time domain

Auto-covariance

Variance time plots

Rescaled Range Statistic

Auto-correlation

- Auto-covariance of second-order stationary time series $\{X_k\}_{k\geq 1}$

$$\gamma(k) := \mathbb{E}[(X_1 - \mathbb{E}[X])(X_k - \mathbb{E}[X])] = \mathbb{E}[X_1 X_k] - \mathbb{E}[X]^2$$

- Sample auto-covariance from finite data $X_1, ..., X_N$

$$\widehat{\gamma(k)} := \frac{1}{N-k} \sum_{j=1}^{N-k} (X_j - \bar{X})(X_{j+k} - \bar{X}) \qquad \bar{X} = \frac{1}{N} \sum_{j=1}^{N} X_j$$

- Connection to LRD and Hurst exponent:
  - If γ(k) is not summable then LRD
  - Hurst exponent:

$$\gamma(k) \sim k^{2H-2}$$

  - Estimation: 2H-2 = slope of linear fit to log(γ(k)) vs. log (k)

Auto-correlation

Requires sufficient data, estimation of γ at large lag: high error

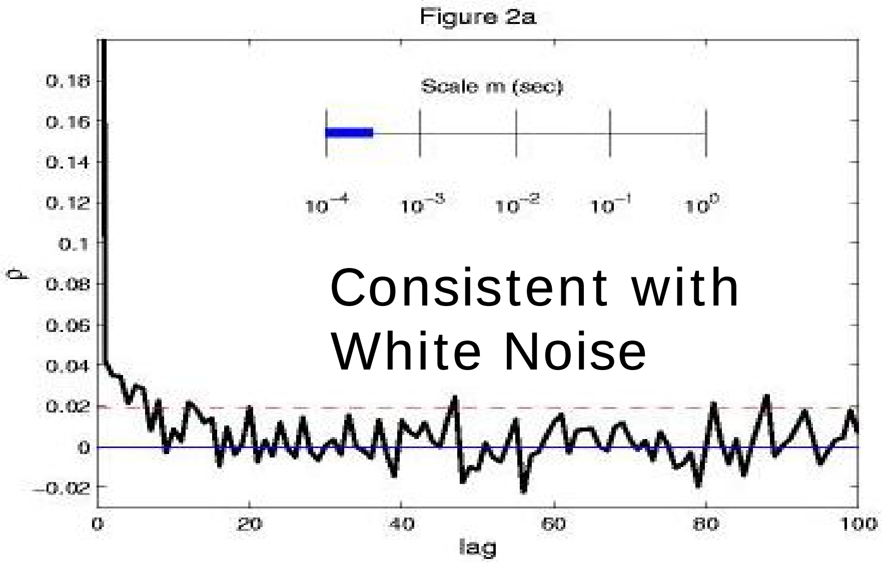

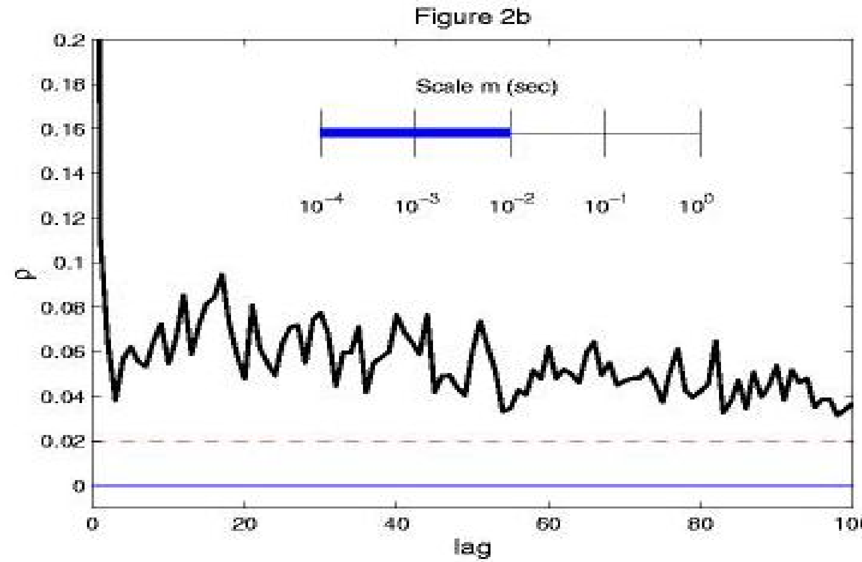

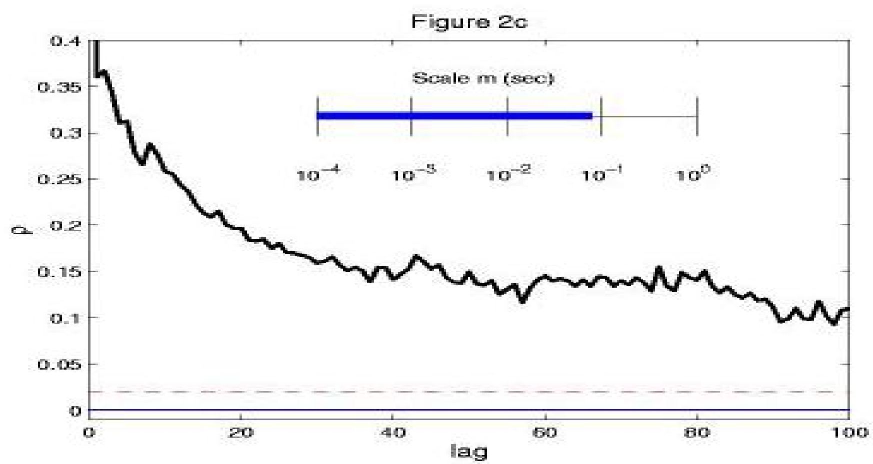

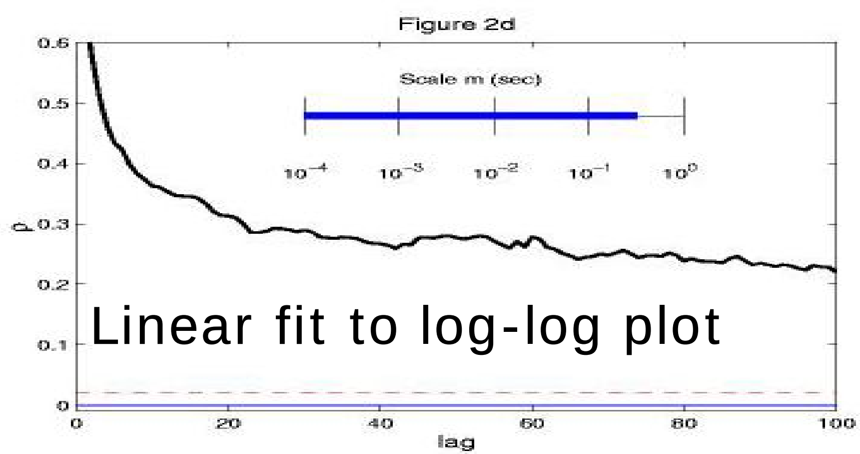

Variance time plots

- Aggregated time series:

$$X_n^{(m)} = X_{(n-1)m+1} + \ldots + X_{nm}$$

- LRD:

$$\gamma(k) \sim k^{2H-2} \text{ iff } \mathrm{Var}X^{(m)} \sim m^{2H}$$

- Independence: H= ½
- Excess variance is indicative of correlation

$$\mathrm{Var}X^{(2m)} = 2\mathrm{Var}X^{(m)} + 2\mathrm{Cov}(X_0^{(m)}, X_1^{(m)})$$

- Estimating the variance

$$\widehat{\mathrm{Var}}\ X^{(m)} = \frac{1}{N/m}\sum_{k=1}^{N/m}(X^{(m)}(k))^2 - \left(\frac{1}{N/m}\sum_{k=1}^{N/m} X^{(m)}(k)\right)$$

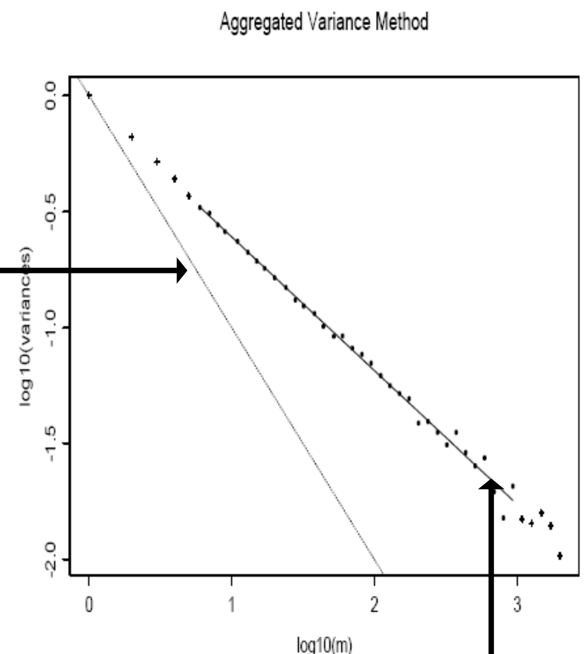

FGN $(H = 0.7)$.

- Inherent difficulty: due to dependence the error of the variance estimate is larger than with iid data

Adjusted Range Statistics

- Sample mean and sample variance after n observations:

$$\bar{X}(n) = (X_1 + \ldots X_n)/n \qquad S(n) = \frac{1}{n-1}\sum_{k=1}^{n}(X_k - \bar{X}(n))^2$$

- Adjusted range statistic:

Mean drain

$$\frac{R}{S}(n) = \frac{1}{S(n)}[\max_{k\le n}(X_1 - k\cdot\bar{X}(n)) - \min_{k\le n}(X_1 - k\cdot\bar{X}(n))]$$

Actual influx until k

- For fGn and FARIMA

$$\mathbb{E}[\frac{R}{S}(n)] \sim n^H$$

- Estimation inprecision!

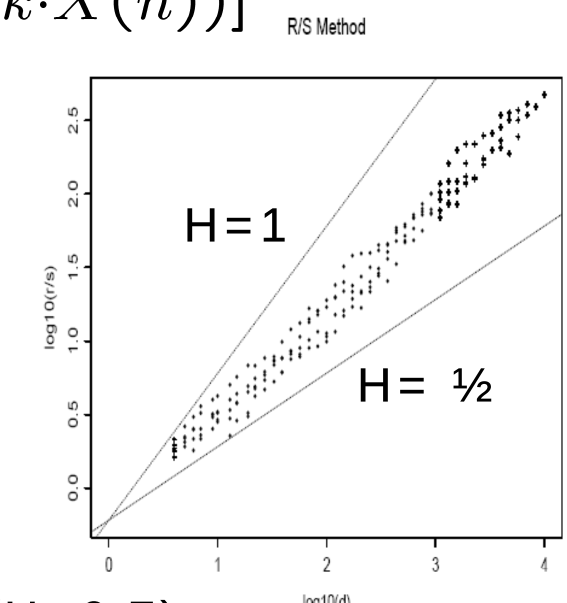

FGN (H=0.7)

Spectral domain

LRD and spectral density

- Assume γ has spectral density f

$$\gamma(k) = \int_{-\pi}^{\pi} e^{i\nu k} f(\nu) d\nu, \ k \in \mathbb{Z},$$

$$f(\nu) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} e^{-i\nu k} \gamma(k), \ \nu \in [-\pi, \pi].$$

- If γ is ultimately monotone then LRD is equivalent with either

$$\begin{array}{llllll}
(i) \sum_{k=-n}^{n} & \gamma(k) & \sim & n^{\alpha} L_1(n), & \text{as } n \to \infty, & 0 < \alpha < 1, \\
(ii) & \gamma(k) & \sim & k^{-\beta} L_2(k), & \text{as } k \to \infty, & 0 < \beta < 1, \\
(iii) & f(\nu) & \sim & |\nu|^{-\gamma} L_3(|\nu|), & \text{as } \nu \to 0, & 0 < \gamma < 1.
\end{array}$$

where $\gamma = \alpha = 1 - \beta$

Spectral estimation

- Fourier trafo of auto-covariance→ errors accumulate
- Wiener Khinchine (wide-sense-stationary process)

Fourier Transform of autocorrelation function = power spectral density (generalization of |Fourier|^2)

- Quick and dirty for L2 signals (instead of processes)

$$\gamma_X(\tau) = \int x(t)x(t-\tau)dt \qquad \mathcal{F}x(\nu) = \int x(\tau)e^{i\tau\nu}d\tau$$

$$\begin{aligned}\mathcal{F}\gamma(\nu) &= \int \gamma(\tau)e^{i\tau\nu}d\tau \\ &= \int\int x(t)x(t-\tau)\overline{e^{i(t-\tau)\nu}}e^{it\nu}dtd\tau \\ &= \overline{\mathcal{F}x(\nu)}\mathcal{F}x(\nu) = |\mathcal{F}x(\nu)|^2\end{aligned}$$

- LRD series $\quad f(\nu) := \sum_k \gamma(k)e^{ik\nu} \simeq \nu^{1-2H} \to \infty \quad (\ \nu \to 0\ )$

Spectral estimation

- Estimate power spectrum via so-called periodogram

$$I(\nu) = \frac{1}{2\pi N} \left| \sum_{j=1}^{N} X_j e^{ij\nu} \right|^2$$

- LRD:

$$I(\nu) \sim \nu^{1-2H} \quad (\nu \to 0)$$

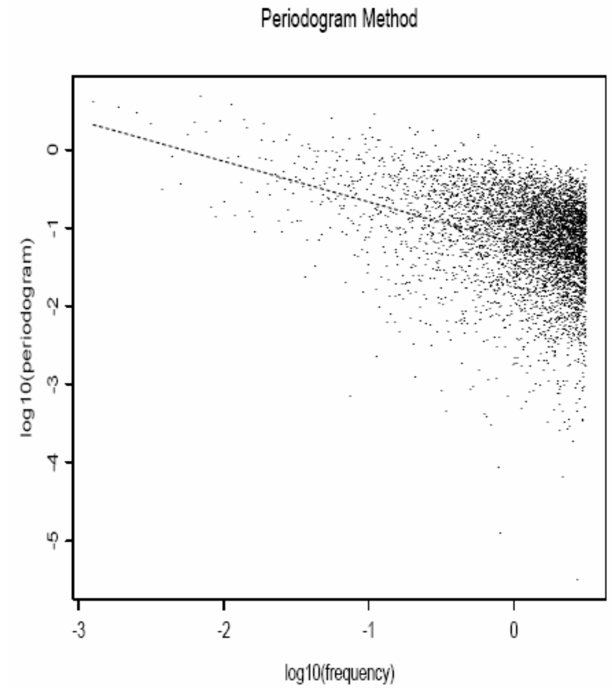

- Estimate H via the slope of a linear fit to the log(periodogram) against log(λ) [see figure]

Estimators applied to FGN

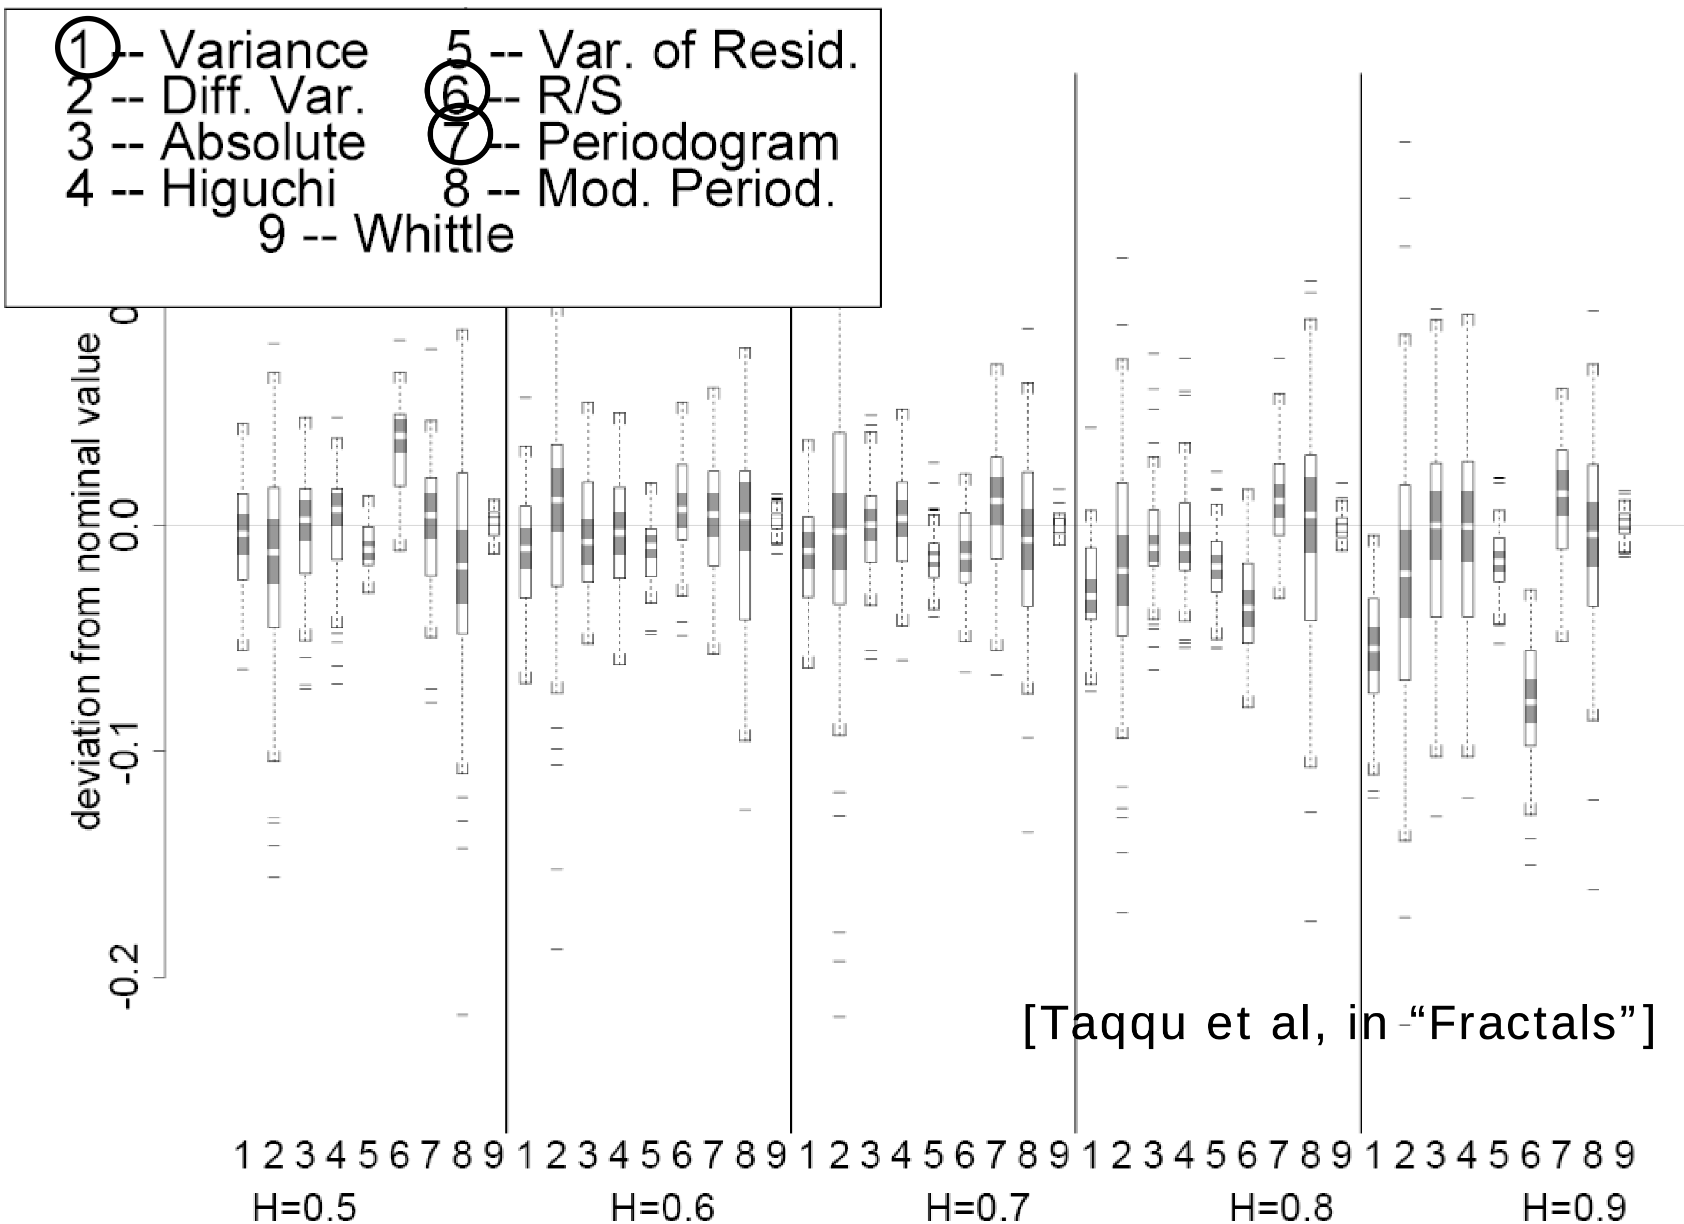

[Taqqu et al, in "Fractals"]

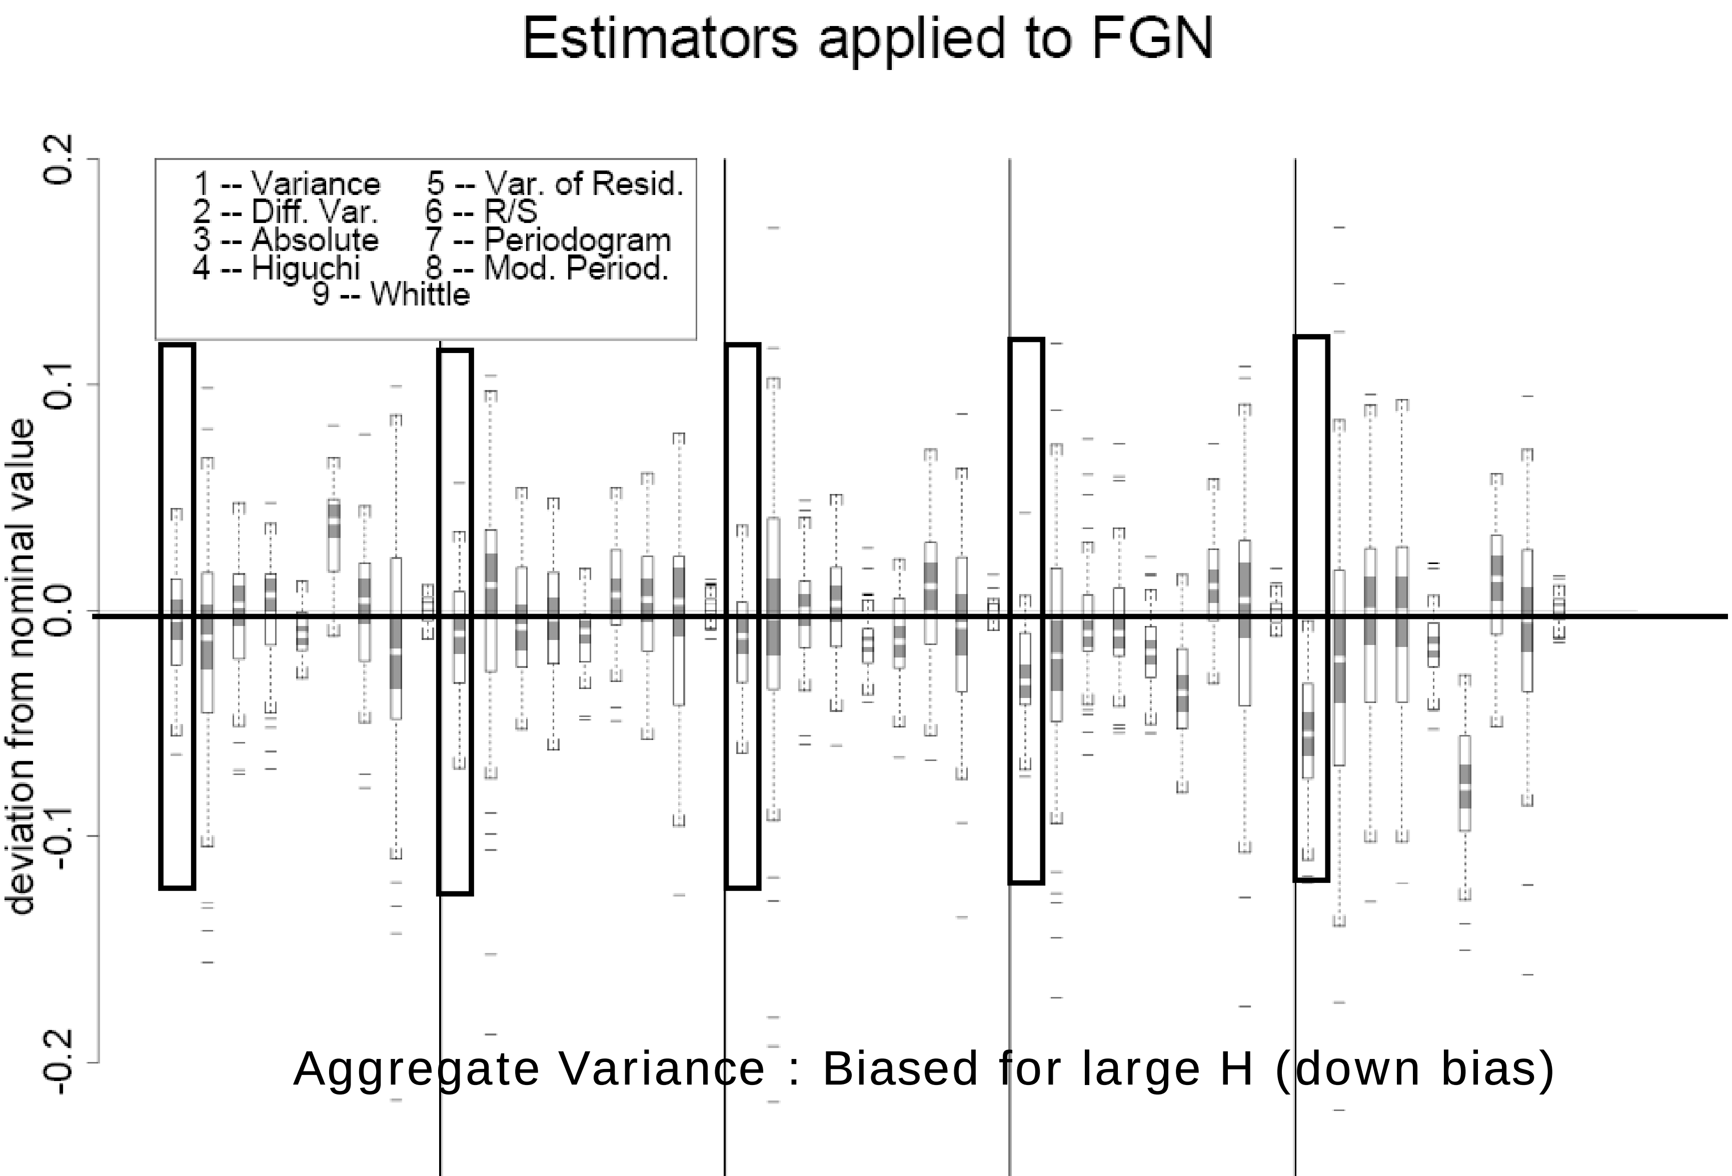
Estimators applied to FGN
deviation from nominal value
0.2
0.1
0.0
-0.1
-0.2
1 -- Variance
2 -- Diff. Var.
3 -- Absolute
4 -- Higuchi
5 -- Var. of Resid.
6 -- R/S
7 -- Periodogram
8 -- Mod. Period.
9 -- Whittle
Aggregate Variance : Biased for large H (down bias)
1 2 3 4 5 6 7 8 9 1 2 3 4 5 6 7 8 9 1 2 3 4 5 6 7 8 9 1 2 3 4 5 6 7 8 9 1 2 3 4 5 6 7 8 9
H=0.5
H=0.6
H=0.7
H=0.8
H=0.9

Estimators applied to FGN

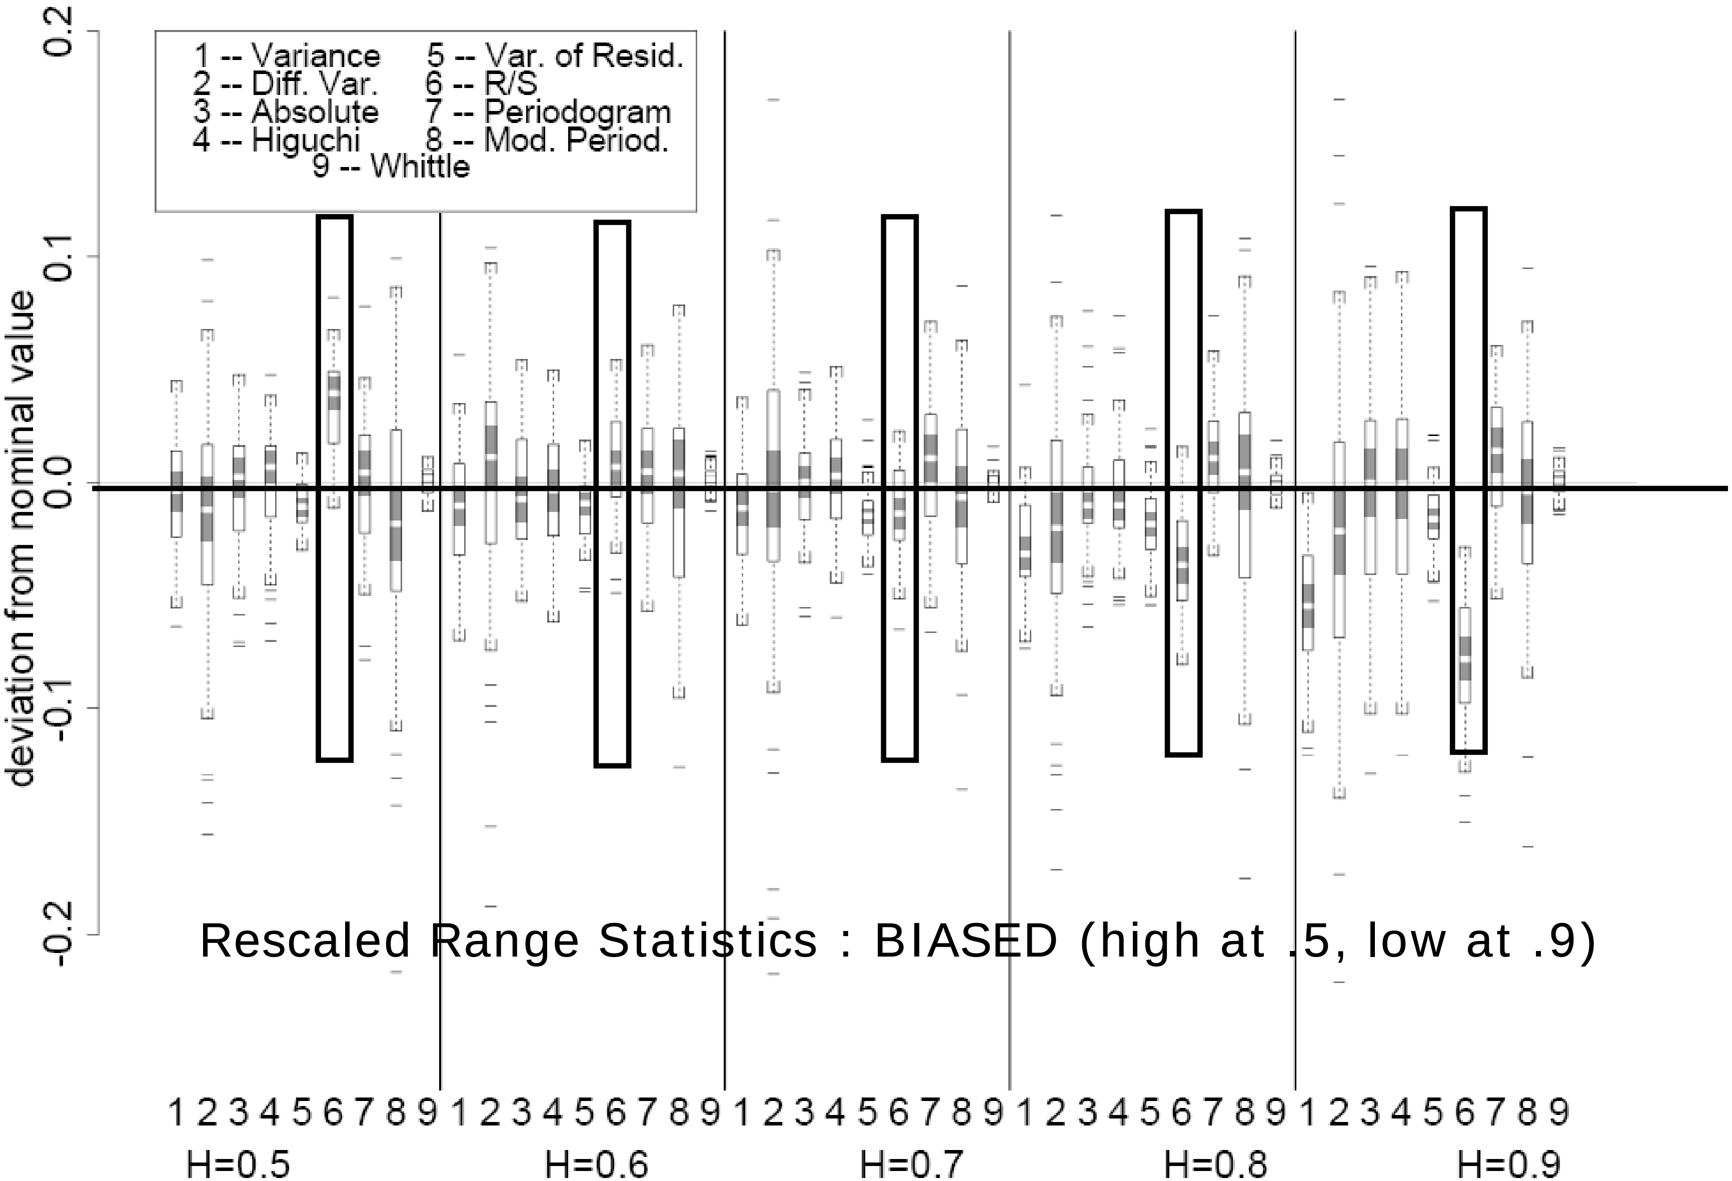

Rescaled Range Statistics : BIASED (high at .5, low at .9)

Estimators applied to FGN

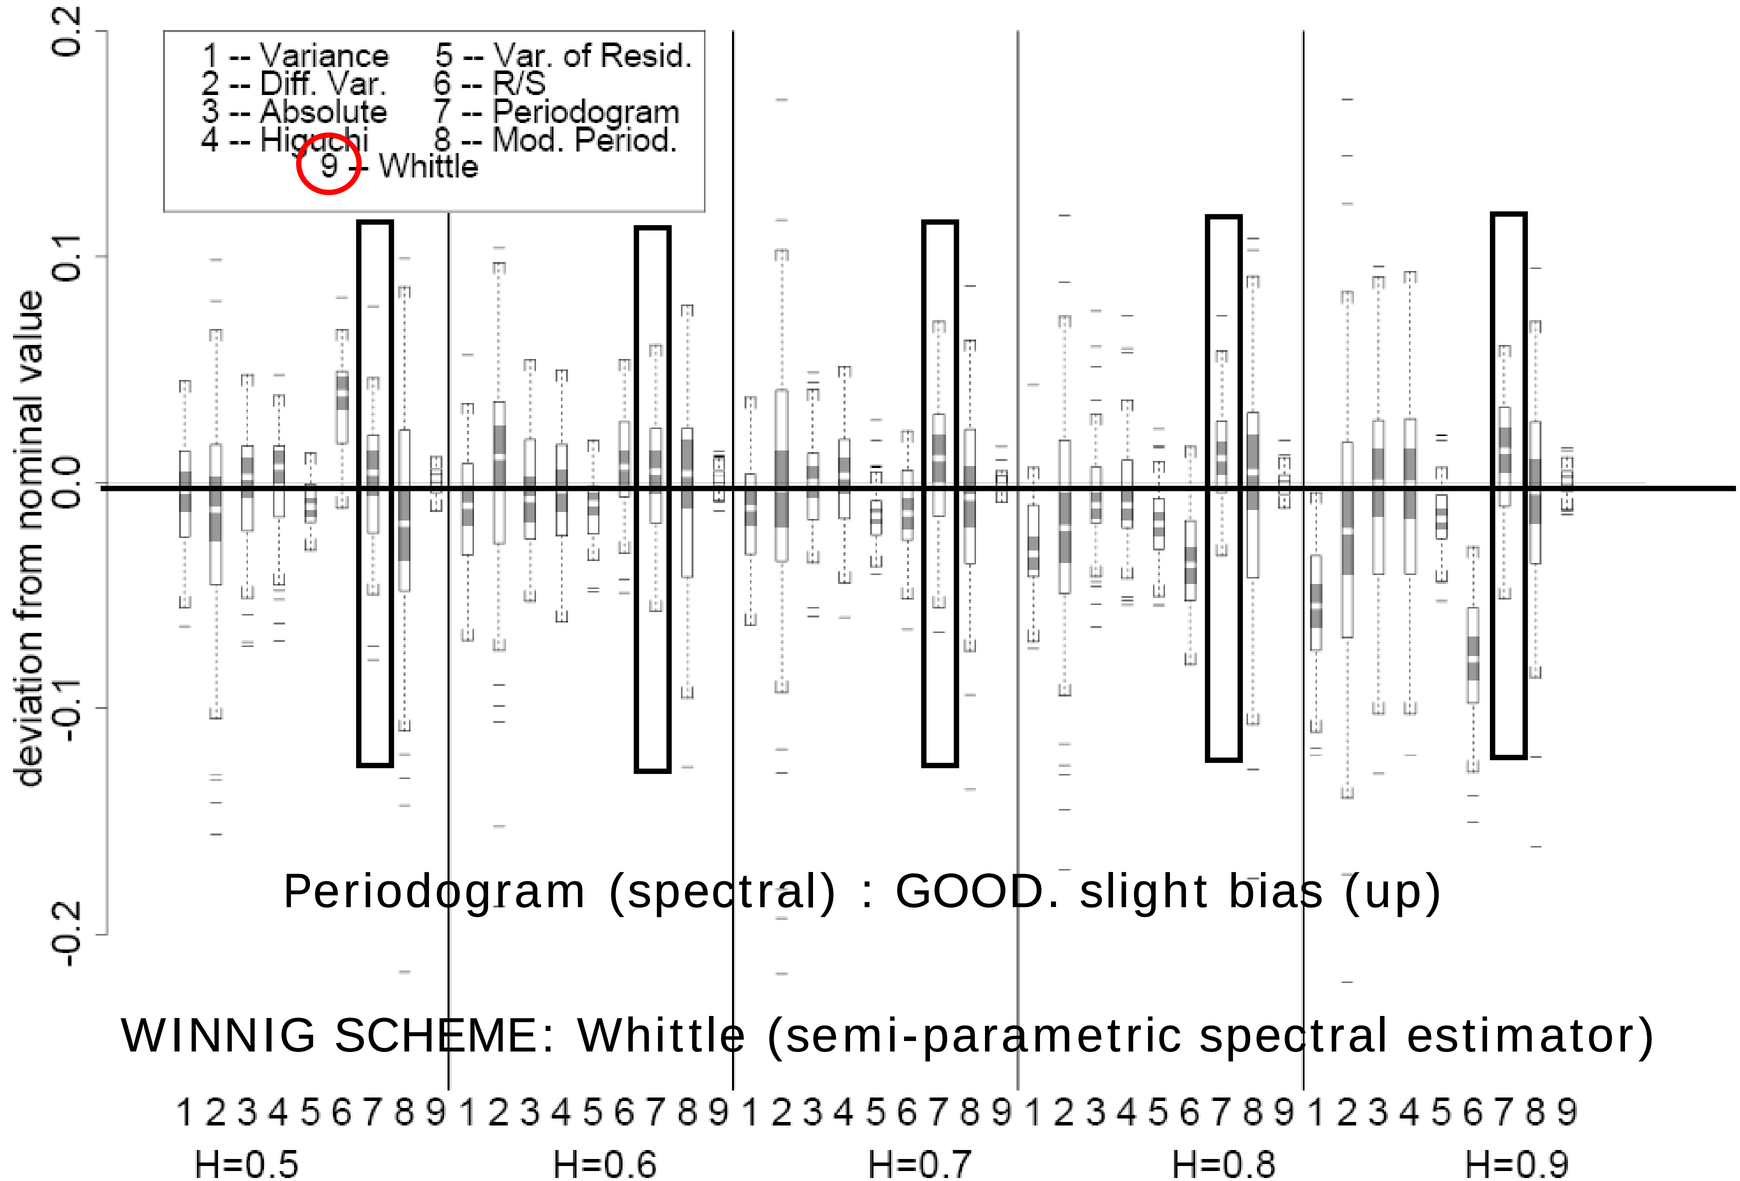

Periodogram (spectral) : GOOD. slight bias (up)

WINNIG SCHEME: Whittle (semi-parametric spectral estimator)

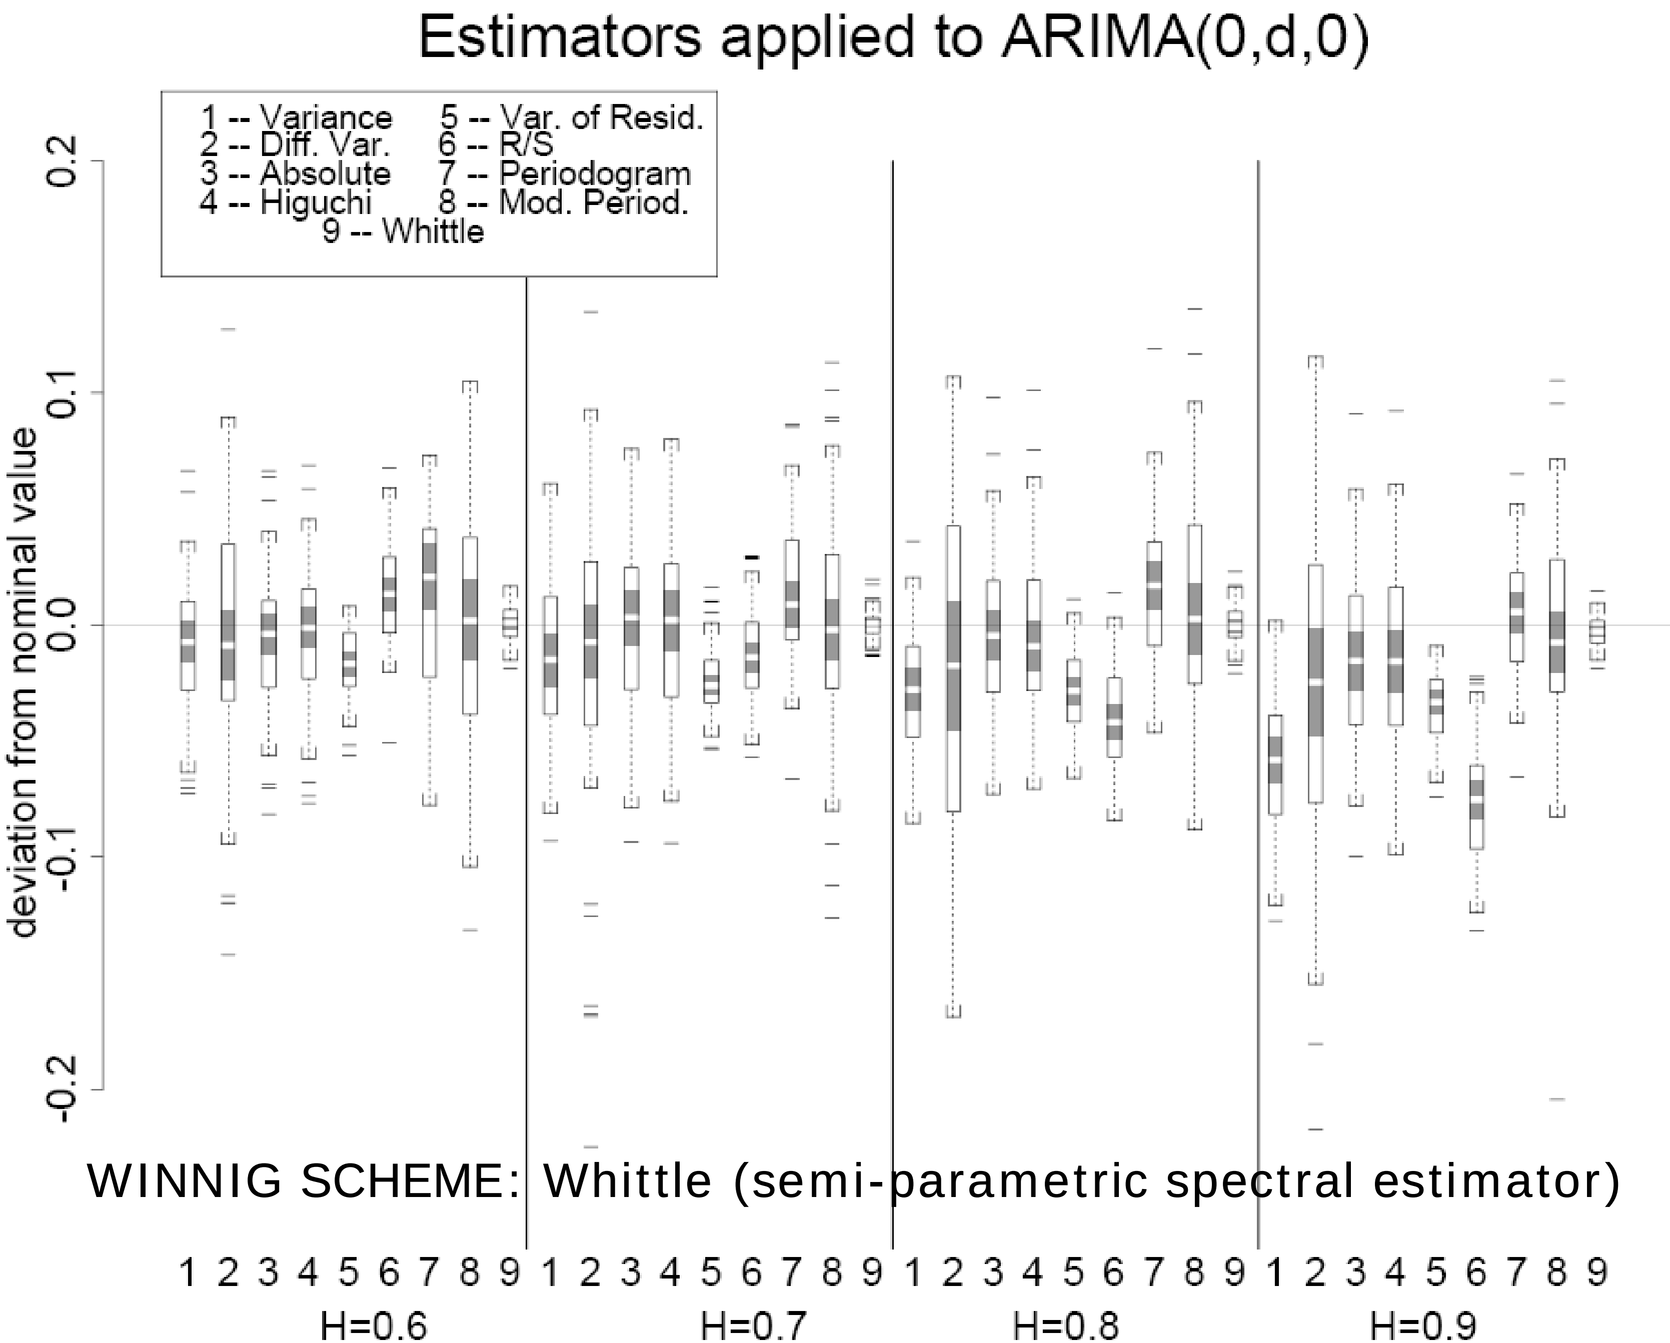

WINNIG SCHEME: Whittle (semi-parametric spectral estimator)

Wavelet domain

Un-biased (!) estimation of LRD

Spectral properties of wavelets

Wavelets:

- localized both in space and frequency

  Power spectrum $|\Psi|^2$ with

$$\Psi(\nu) := \int \psi(t) e^{i\nu t} dt$$

  peaks at characteristic frequency ν_ψ

Continuous wavelet transforms of $2^{-j}\psi(2^j(t-k))$ for $j = 0, 1, 2, 3$

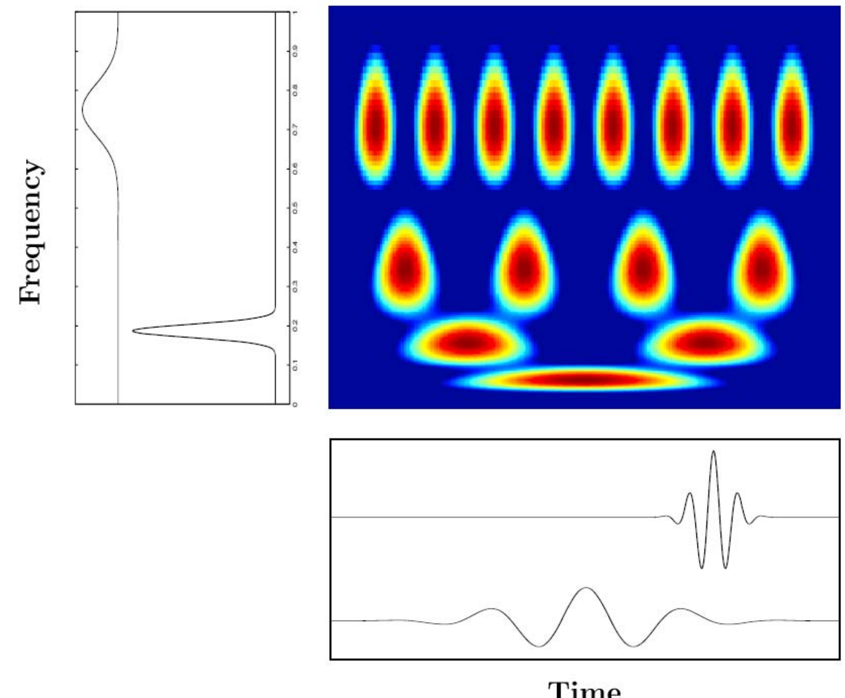

- affine family (see figure):

$$\begin{aligned} \Psi_{j,k}(\nu) &:= \int 2^{-j}\psi(2^j(t-k)) e^{i\nu t} dt \\ &= e^{i\nu k}\Psi(2^{-j}\nu) \end{aligned}$$

$\Longrightarrow$

frequency response of $\psi(2^j(t-k))$ peaks at $2^{-j}\nu_\psi$

(less localized at

small scales/high frequencies)

Wavelet domain

- Frequency response of $\psi(2^j(t-k))$ at $2^{-j}\nu_\psi$
- Motivates to estimate power spectrum via

$$\hat{\Gamma}_x(2^{-j}\nu_0) = \frac{1}{n_j}\sum_k |d_x(j,k)|^2$$

- Estimate H via linear regression on log-log:

$$\log_2(\hat{\Gamma}_x(2^{-j}\nu_0) = \log_2(\frac{1}{n_j}\sum_k |d_x(j,k)|^2) = (2\hat{H}-1)j + \hat{c}$$

Three wave-mirac-lets for fBm

- Wavelet coefficients for fBm are *stationary*
  - Allows for estimation
- ...less correlated
  - Reduces the estimation error of sample variance
- ...yield unbiased estimate of H
  - Bias of wavelet-periodogram-estimator is multiplicative...and does not affect the slope of the log-log data!

For fBm stationary and decorrelated

$$
\begin{aligned}
&\mathbb{E}[W_{jk}W_{jm}] \\
&\stackrel{1}{=} \ \mathbb{E}\int\int \psi_{jk}(t)\psi_{jm}(s)B(t)B(s)dtds \\
&\stackrel{2}{=} \ (\sigma^2/2)\int\int \psi_{jk}(t)\psi_{jm}(s)\left[t^{2H}+s^{2H}-|t-s|^{2H}\right]dtds \\
&\stackrel{3}{=} \ -(\sigma^2/2)\int\int \psi_{jk}(t)\psi_{jm}(s)|t-s|^{2H}dtds \\
&\stackrel{4}{=} \ -(\sigma^2/2)\int |2^j\tau|^{2H}\gamma_\psi(\tau-(k-m))\,d\tau \\
&\stackrel{5}{=} \ -(\sigma^2/2)|2^j|^{2H}\int \frac{1}{|\nu|^{2H+1}}|\Psi(\nu)|^2 e^{-i2\pi\nu(k-m)}d\nu, \\
&\stackrel{6}{\sim} \ O\left(|k-m|^{2H+1-2N}\right) \\
&\qquad (|k-m|\to\infty)
\end{aligned}
$$

Use vanishing moments of wavelets in (3) & (6)

(1) $W_{jk} = \int \psi_{jk}(t)B(t)dt$
(2) $\mathbb{E}[B(u)B(v)] = (\sigma^2/2)[u^{2H}+v^{2H}-|u-v|^{2H}]$
(3) $\int \psi(s)ds = 0 = \int \psi(t)dt$,
(4) $\psi_{jk}(t) = 2^j \int \psi(t/2^j - k)$, $\tau = 2^{-j}(t-s)$,
$\gamma_\psi(\tau) = \int \psi(t)\psi(t+\tau)dt$
(5) Parceval
(6) $\Psi(\nu) \sim |\nu|^N$

Bias of wavelet-periodo-estimator

- Convolutive bias

$$\mathbb{E}\hat{\Gamma}_x(2^{-j}\nu_0) = \int \Gamma_x(\nu) 2^j |\Psi_0(2^j\nu)|^2 d\nu$$

$$\simeq \Gamma_x(2^{-j}\nu_0)$$

Because Ψ is almost a delta distribution

- ...becomes multiplicative bias in the special case for $\Gamma_X(\nu) = c_f \nu^{1-2H}$

$$\begin{aligned}\mathbb{E}\hat{\Gamma}_x(2^{-j}\nu_0) &= c_f |2^{-j}|^{(1-2H)} \int |\nu|^{(1-2H)} |\Psi_0(\nu)|^2 d\nu \\ &= \underbrace{\Gamma_x(2^{-j}\nu_0)}_{c' 2^{-j(1-2H)}} \underbrace{|\nu_0|^{(2H-1)} \int |\nu|^{(1-2H)} |\Psi_0(\nu)|^2 d\nu}_{\text{Bias factor independent of j !}}\end{aligned}$$

Estimator scales as it should:
$\mathbb{E}[\hat{\Gamma}_x(2^{-j}\nu_0)] \sim 2^{-j(1-2H)}$

Summary: Wavelet H-estimation for fBm

- Estimator is meaningful
  - Wavelet coefficients are stationary
- De-correlated → less error
- Un-biased estimate of H

$$\log_2(\hat{\Gamma}_x(2^{-j}\nu_0) = \log_2(\frac{1}{n_j}\sum_k |d_x(j,k)|^2) = (2\hat{H}-1)j + \hat{c}$$

Cross-traffic inference

Probing

- Ideally:

delay spread of packet pair spaced by *T* sec

correlates with

cross-traffic volume at time-scale *T*

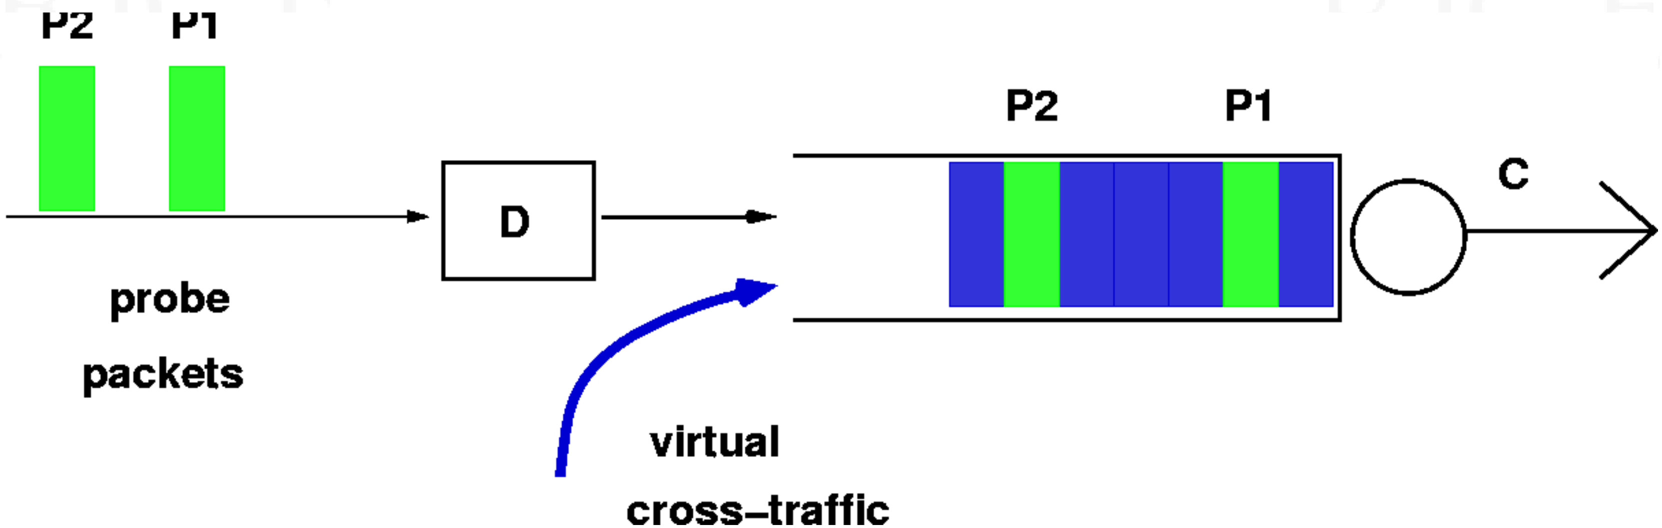

Probing Uncertainty Principle

- Should not allow queue to *empty* between probe packets

- *Small T* for accurate measurements
  - – but probe traffic would disturb cross-traffic (and overflow bottleneck buffer!)

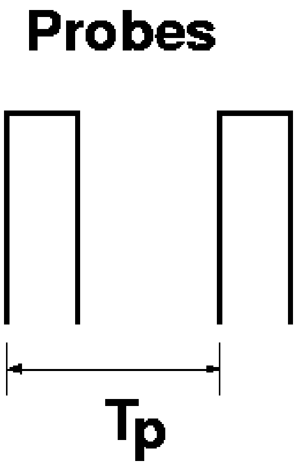

- *Larger T* leads to measurement uncertainties
  - – queue could empty between probes

Multifractal Cross-Traffic Inference

- Model bursty cross-traffic using MWM

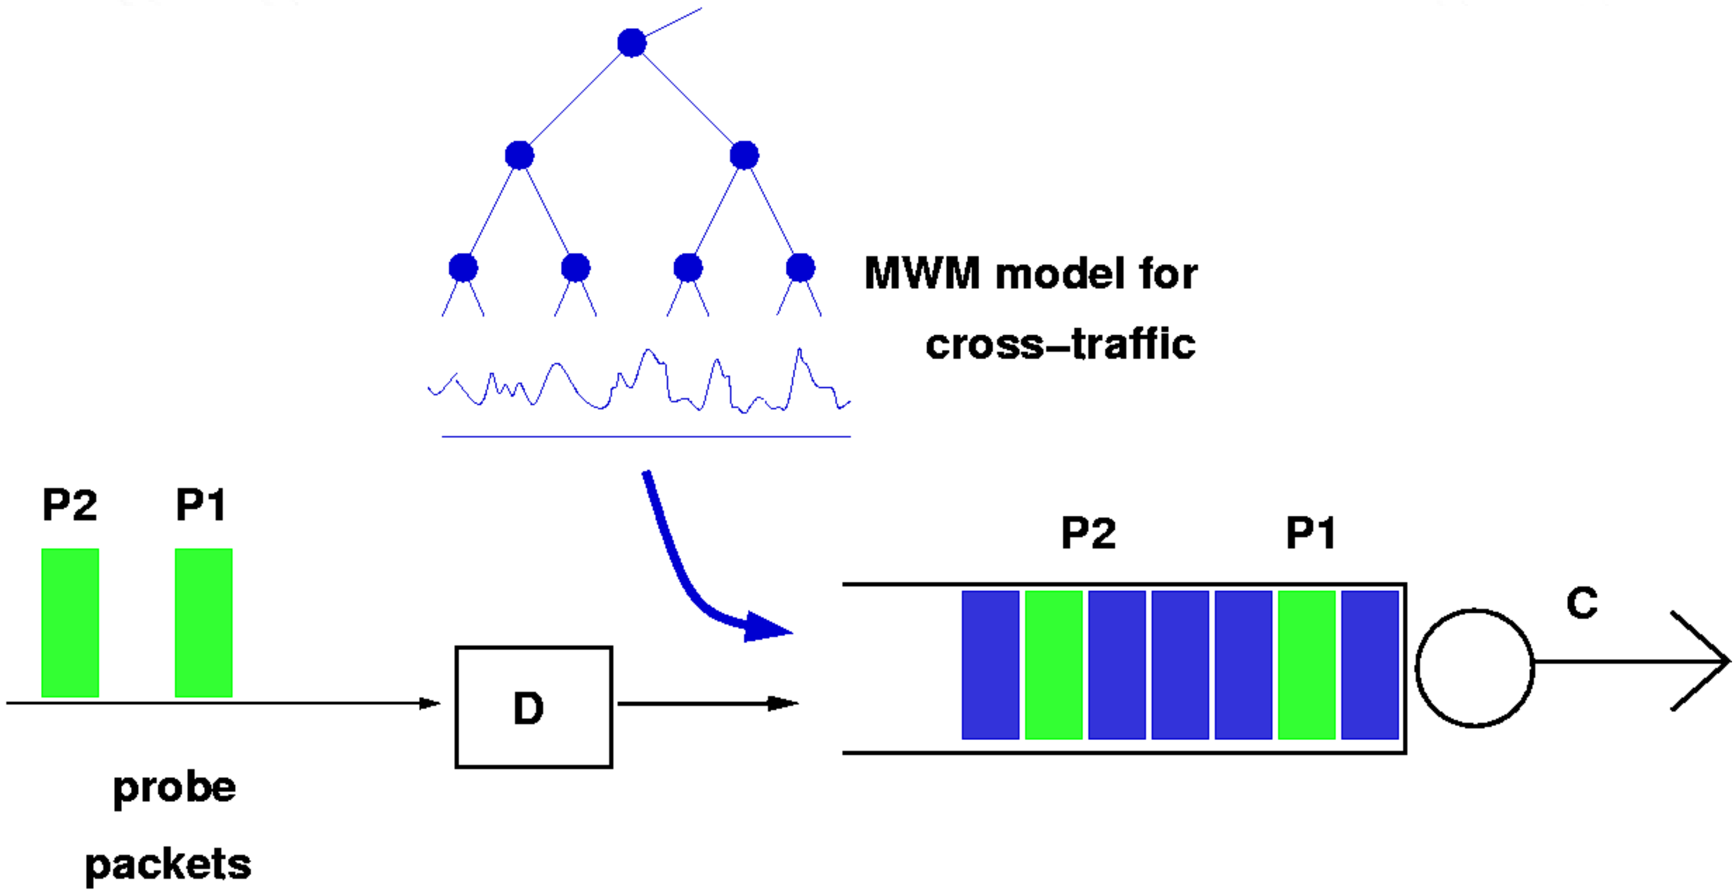

Efficient Probing: exponential spaced

- MWM tree inspires geometric chirp probe
- MLE estimates of cross-traffic at multiple

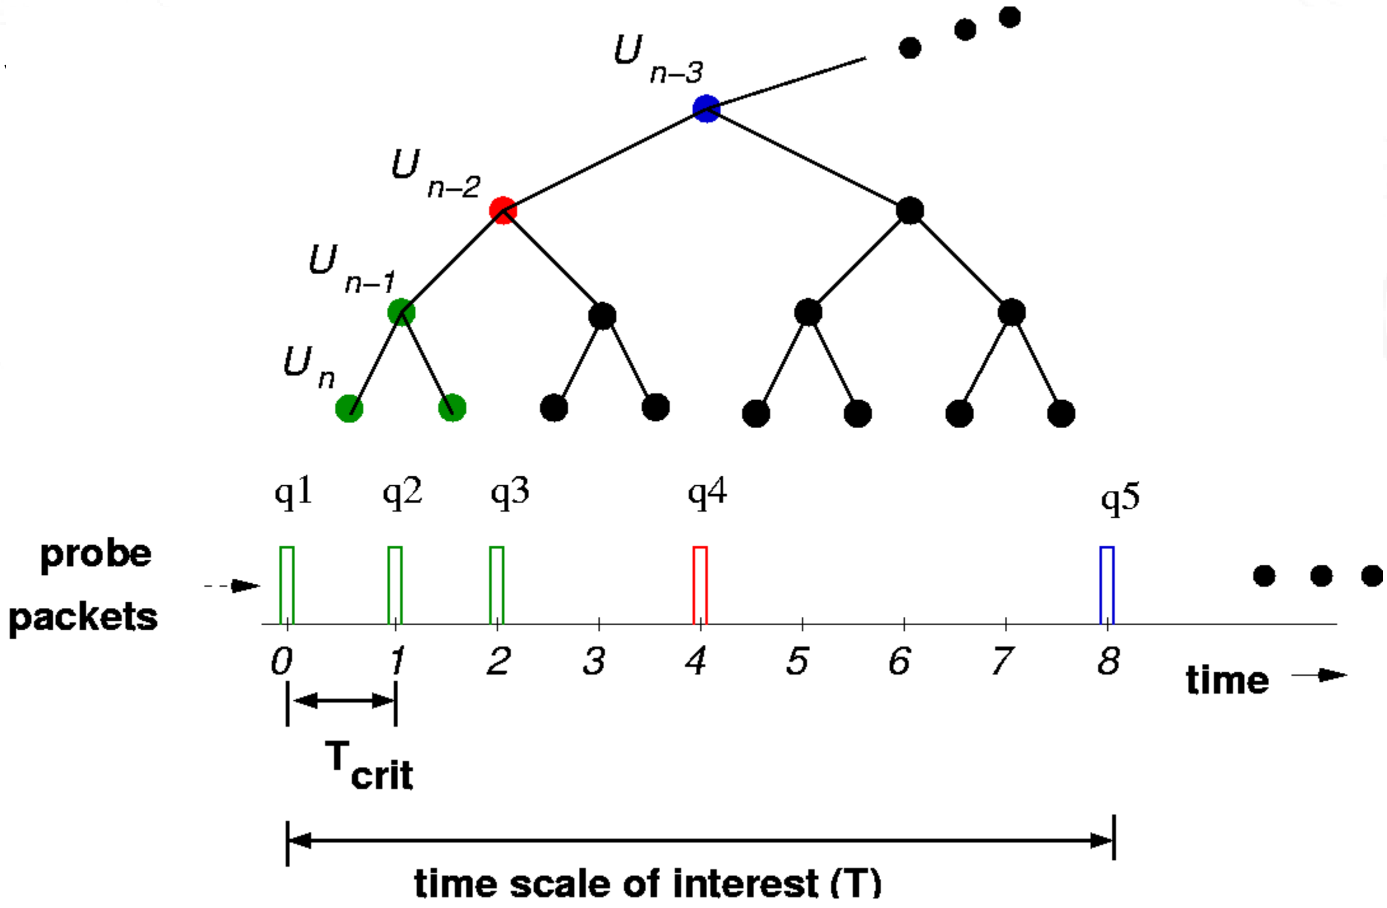

Cross-Traffic Inference

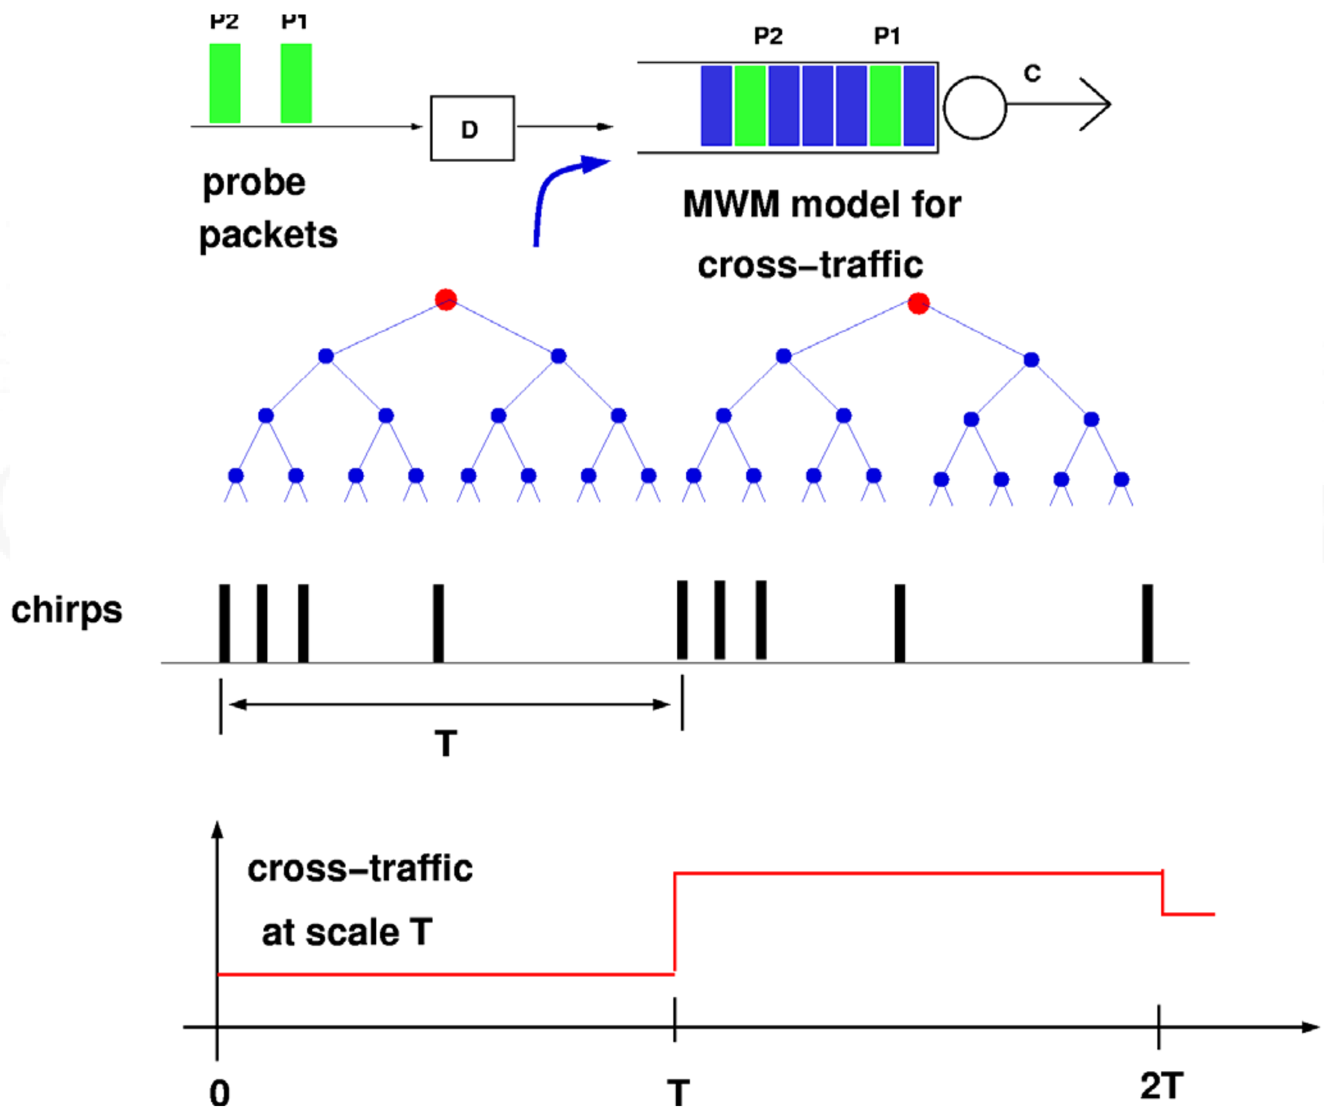

ns-2 Simulation

- Inference improves with increased utilization

Low utilization (39%)

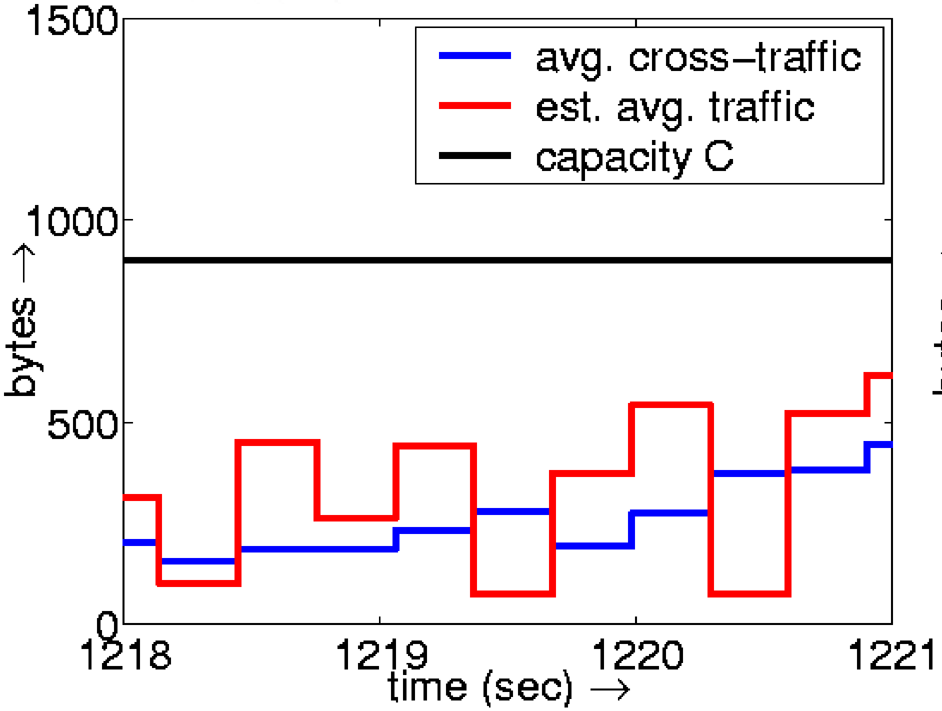

High utilization (65%)

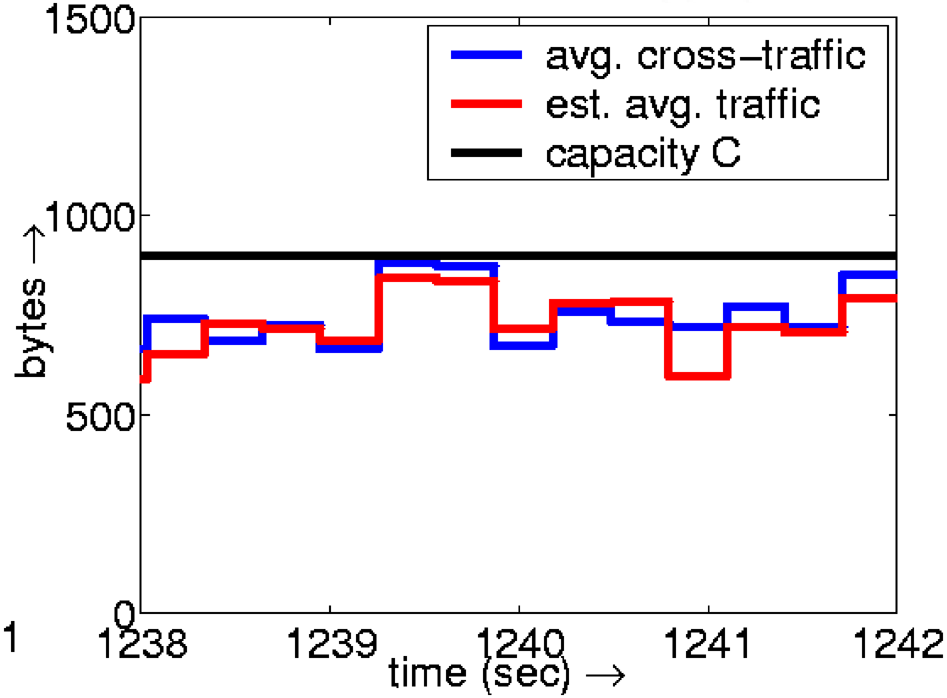

ns-2 Simulation (Adaptivity)

- Inference improves as MWM parameters adapt

MWM parameters

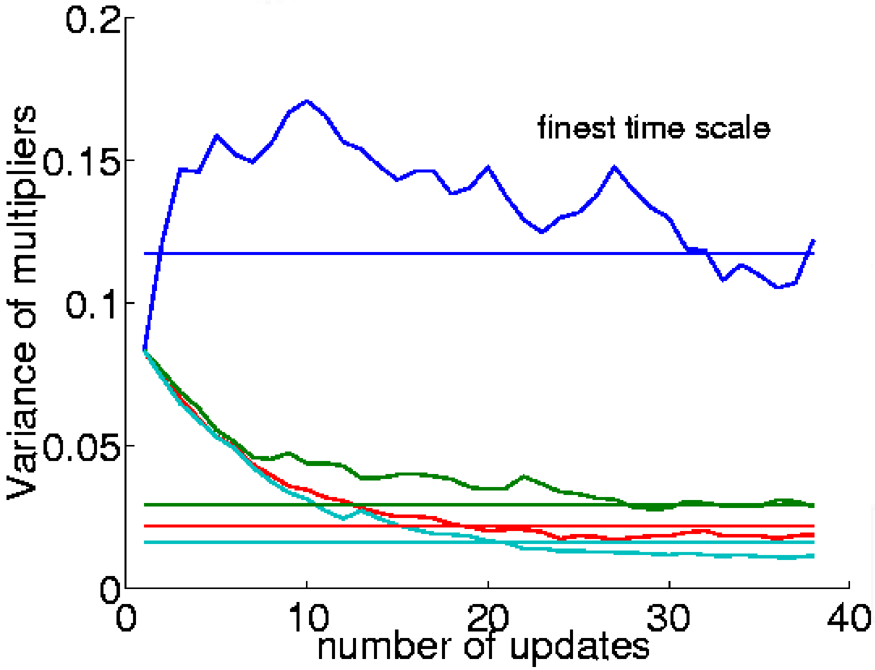

Inferred x-traffic

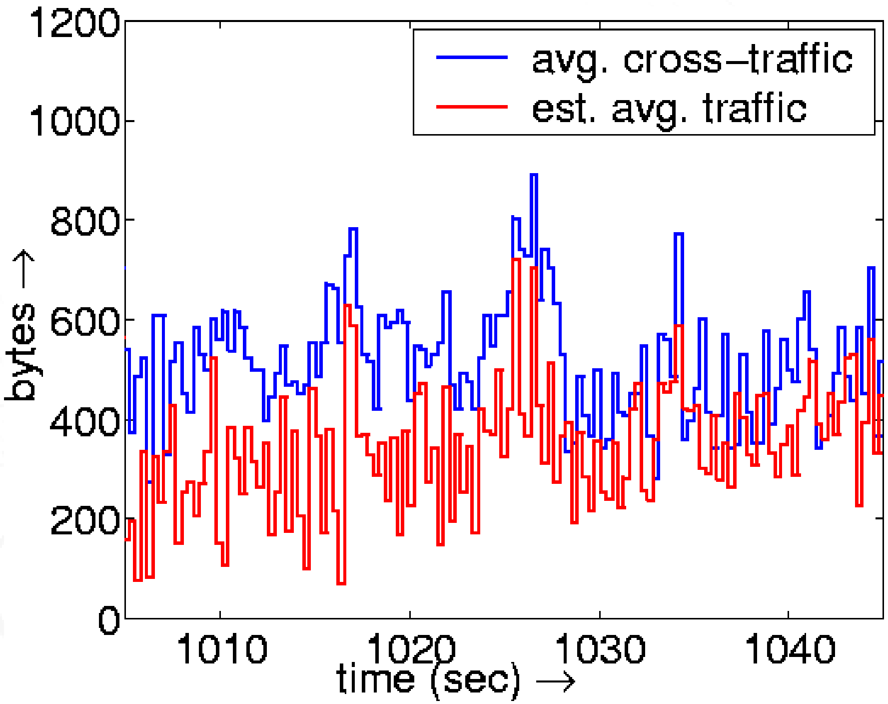

Adaptivity (MWM Cross-Traffic)

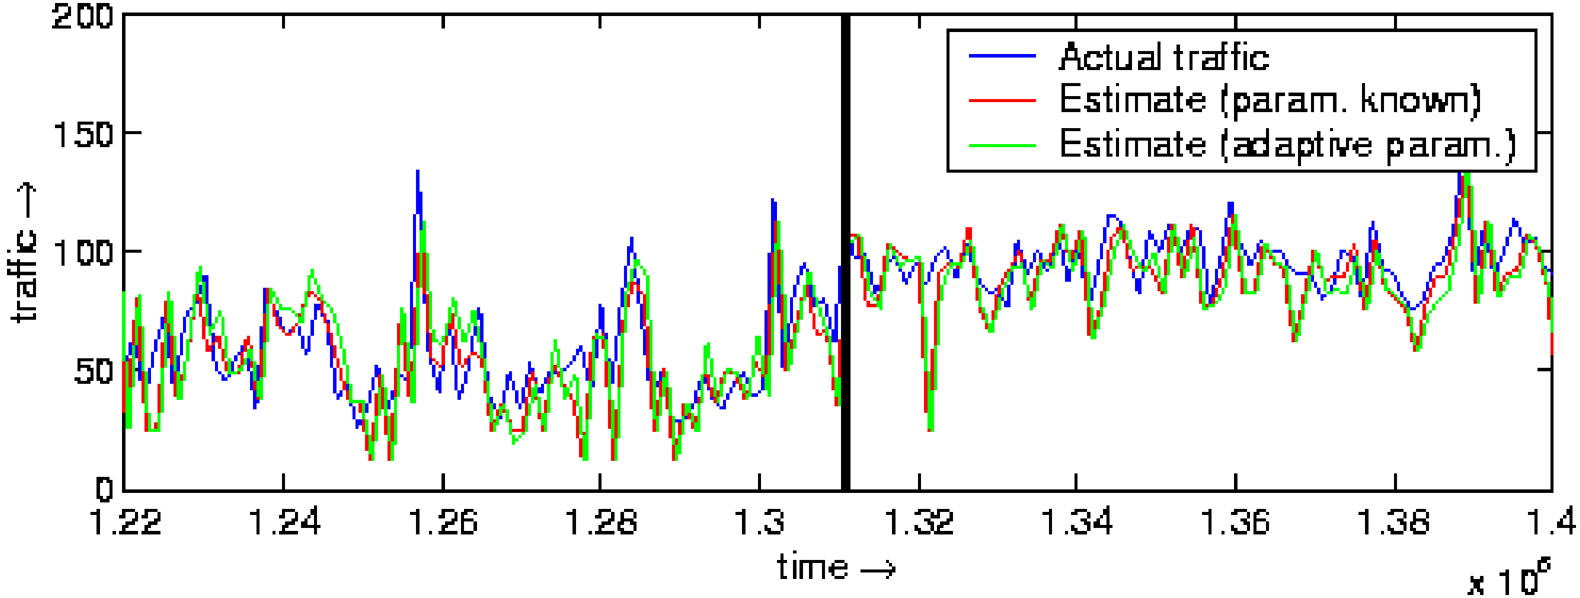

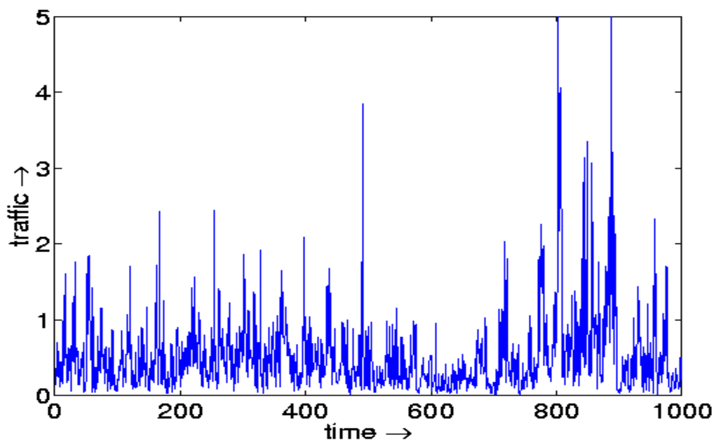

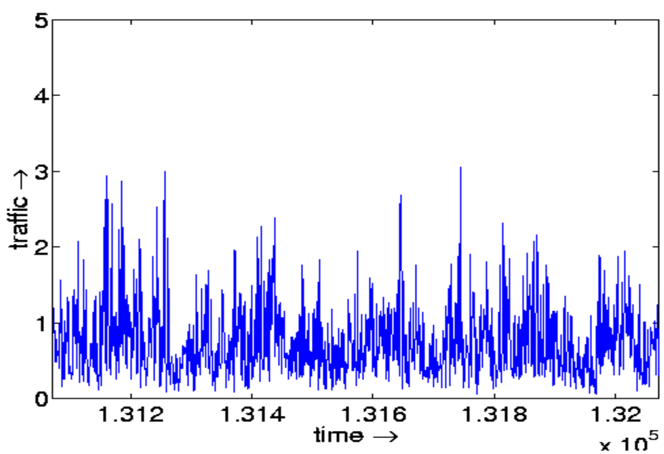

Spatio-Temporal Available Bandwidth Estimation

On-line localization of
the tight link in a path

Key Definitions

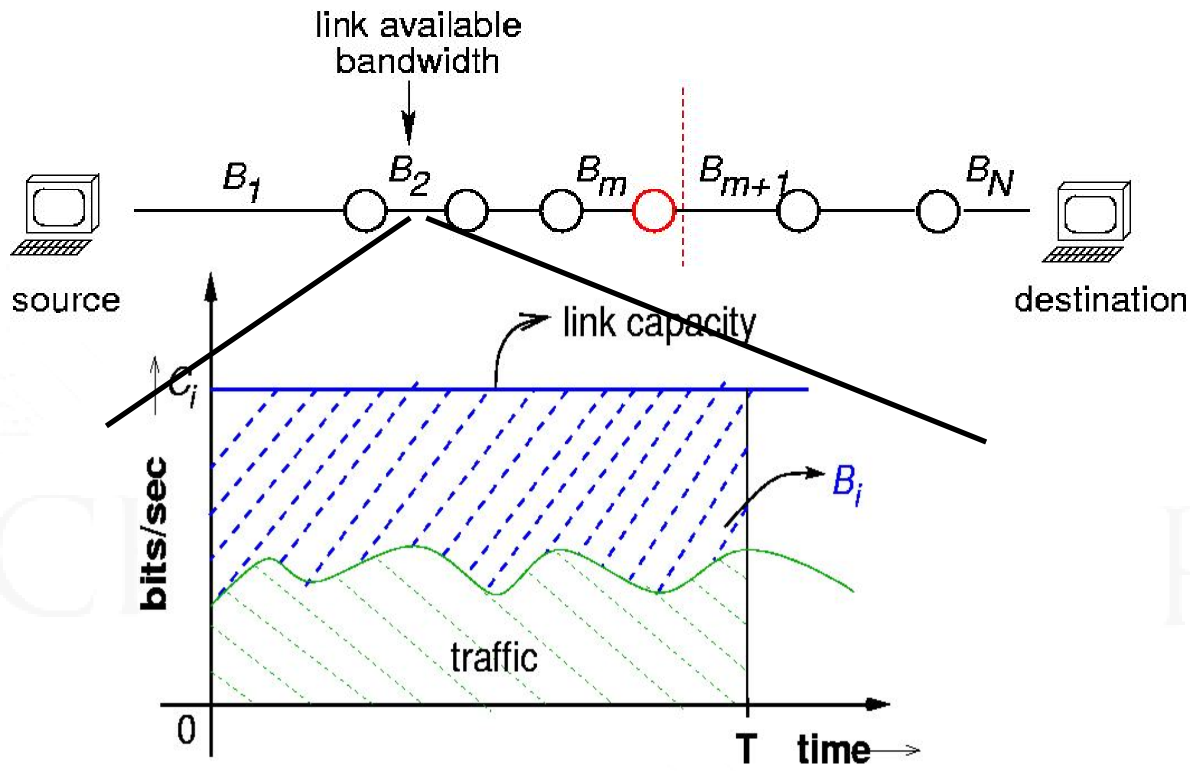

- Available bandwidth: left-over capacity on link
- Tight link: link with least available bandwidth
- **Goal**:
  - locate tight link in space and over time
  - using end-to-end probing

Applications

- Science: *where* do Internet tight links occur and *why*?

- Network aware applicati
  - Server selection
  - Route selection

- Network monitoring
  - locating hot spots

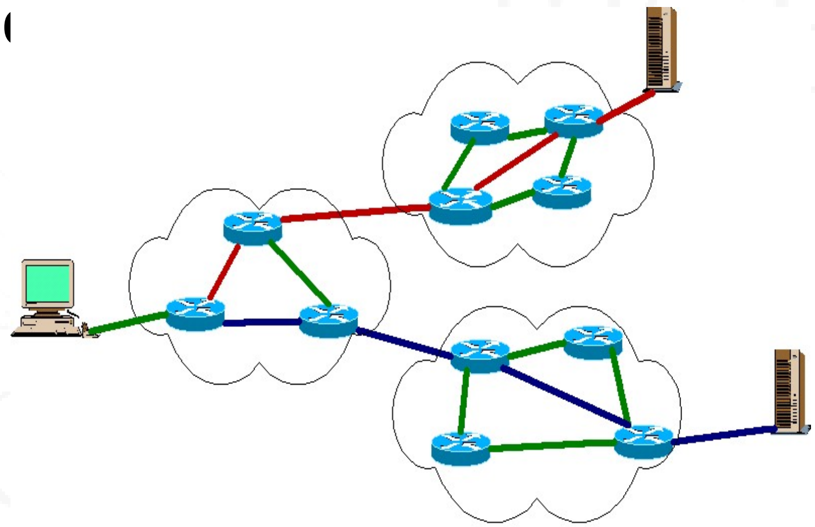

Methodology

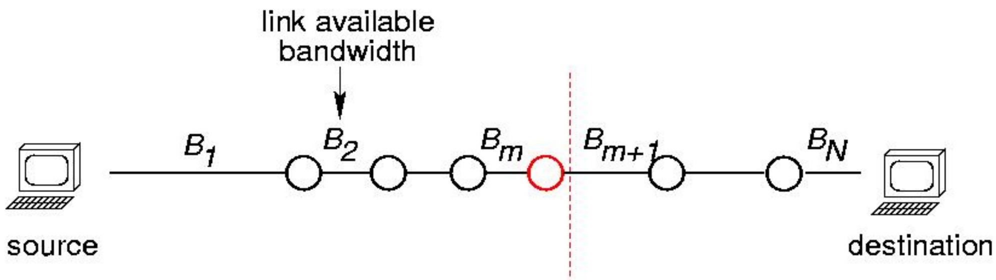

Path available bandwidth

$$\boxed{A = \min_i B_i}$$

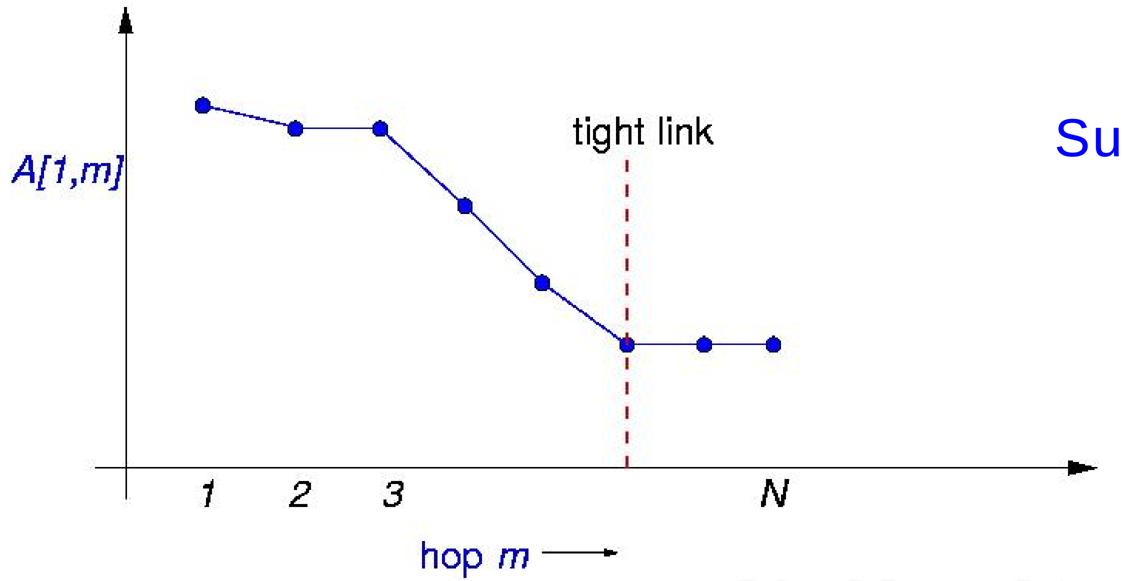

Sub-path available bandwidth

$$\boxed{A[1,m] = \min_{1 \le i \le m} B_i}$$

Methodology:

- For *m*>tight link, *A[1,m]* remains constant

Packet Tailgating

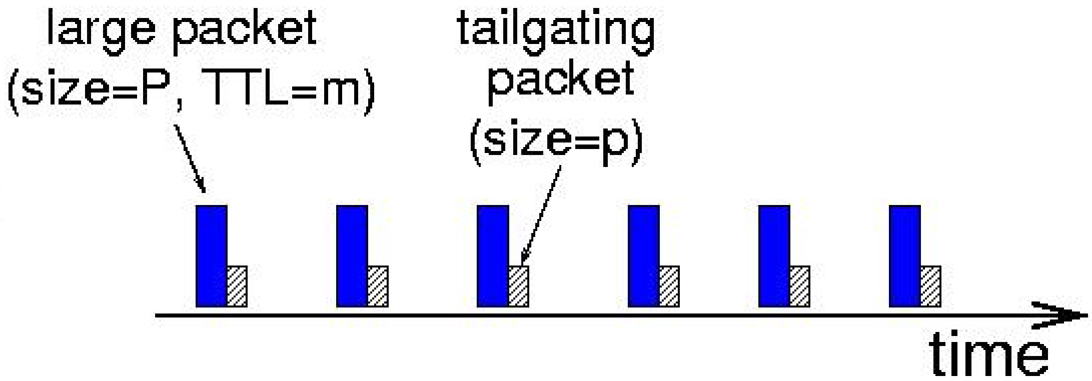

- Packet train contains:
  - Large packets stressing, with *m* hops life time
  - Small packets tailgating, full life time

- Purpose:
  - Large packets "measure" bandwidth via their delay
  - Small packets transport this timing information to the receiver

Efficient probing: *PathChirp*

- Traditional probing paradigm:
  - Produce (light) congestion
  - PacketPair:
    - Sample the traffic
  - Pathload: flood at variable rate
    - intolerable level of congestion
  - TOPP:
    - PacketPairs at variable spacing

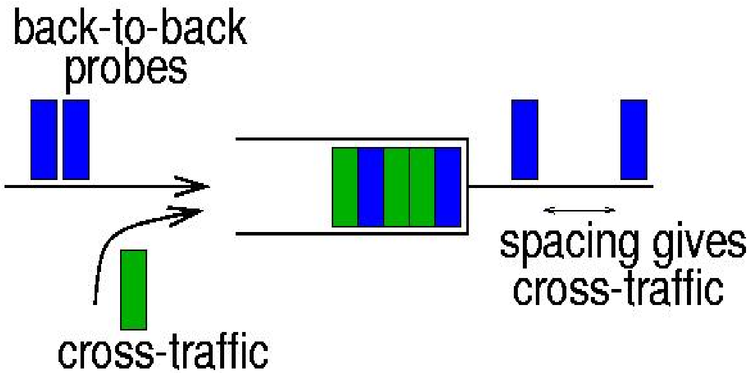

- New:
  - PathChirp:
    - Variable rate within a train of probes
    - More efficient, light

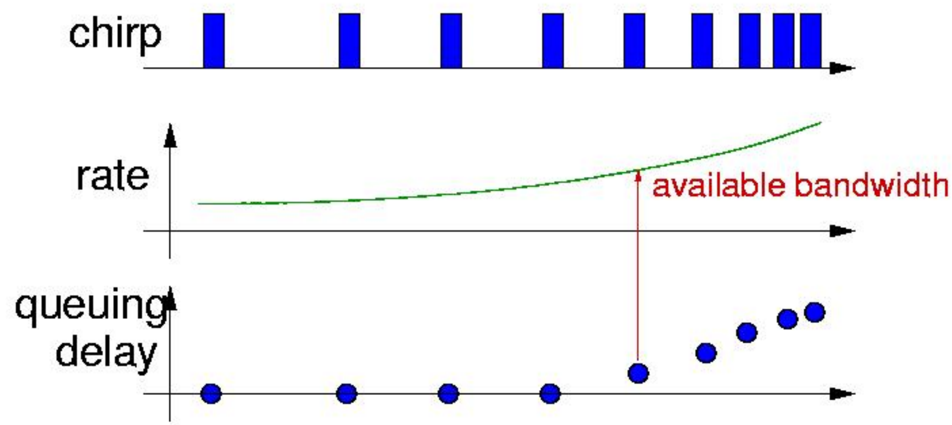

Lite-probing: pathChirp

- real world tool
  - Queuing delay → cross traffic
  - Averaged excursions →available resources
- Light weight
  - Probe at various rates simultaneously
- ...converges in a handful of RTTs

Methodology

Departure pattern

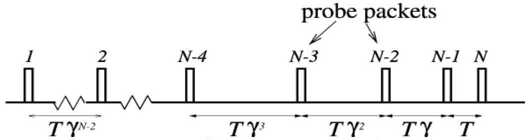

Queuing against departure

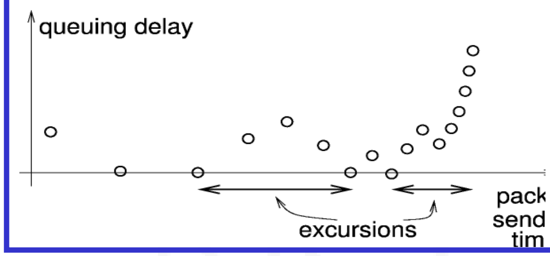

Real world experiments

Internet experiment

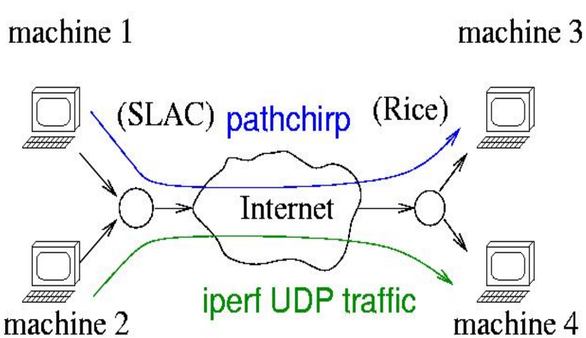

Estimation against true x-traffic

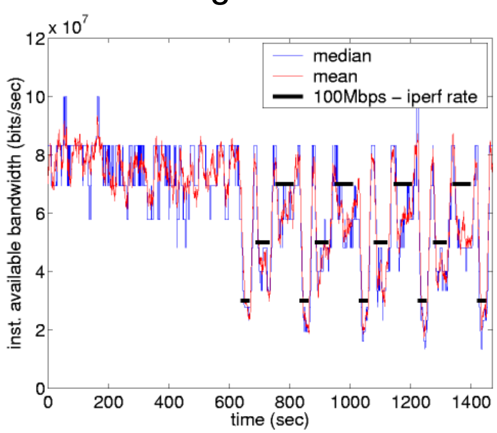

Number of chirps →

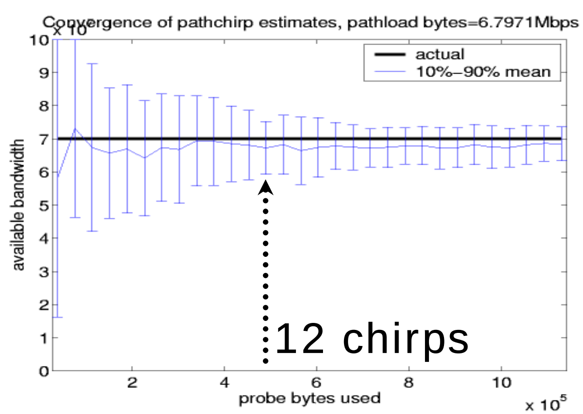

Bandwidth: a Probabilistic entity

- Available bandwidth depends on temporary congestion level of potential tight links

REAL WORLD EXPERIMENTS
UIUC (J. Hou)– Rice
Available sub-path bandwidth

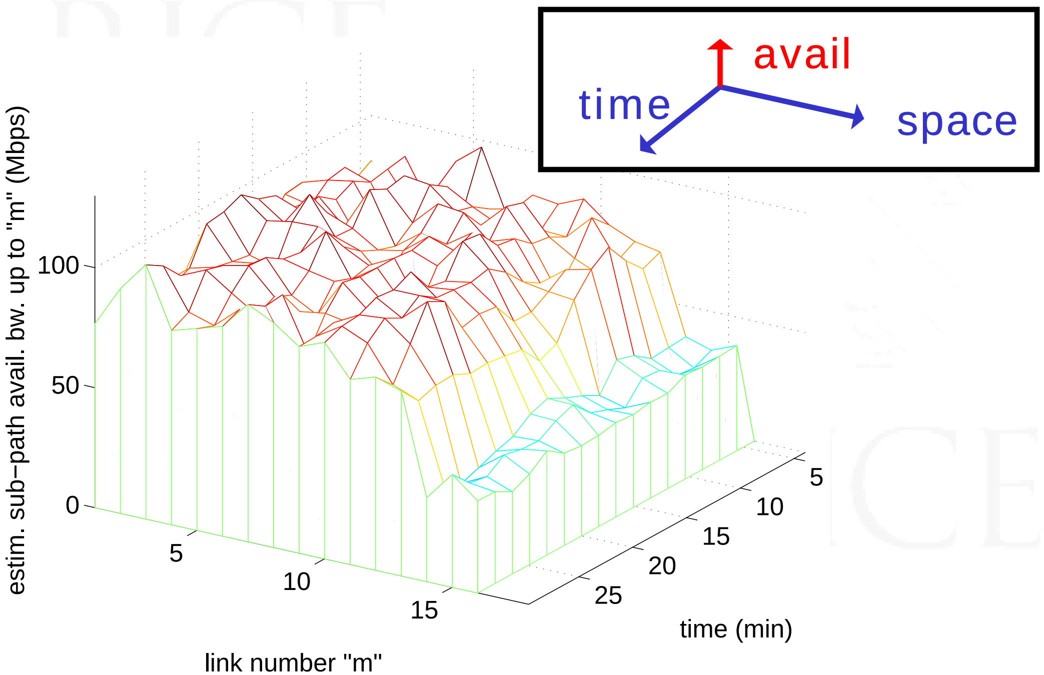

UIUC – Rice
Probability of being tight link
Estimates taken 10 mins apart

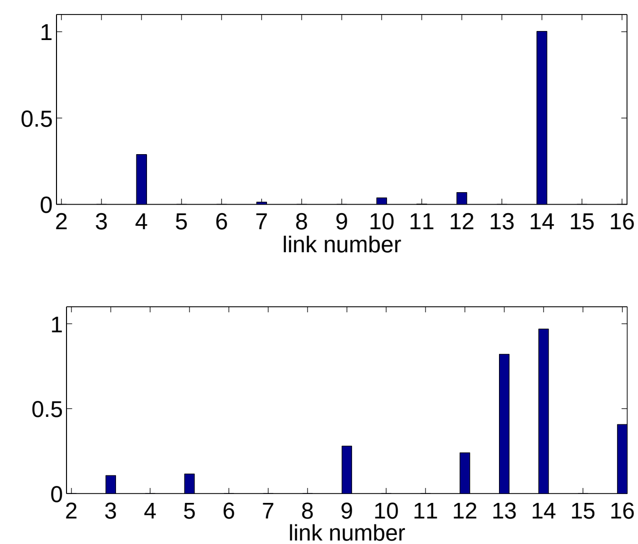

STAB: Spatio Temporal available Bandwidth

- STAB detects new tight link and reduced available bandwidth around 250 secs into simulation

Estimate

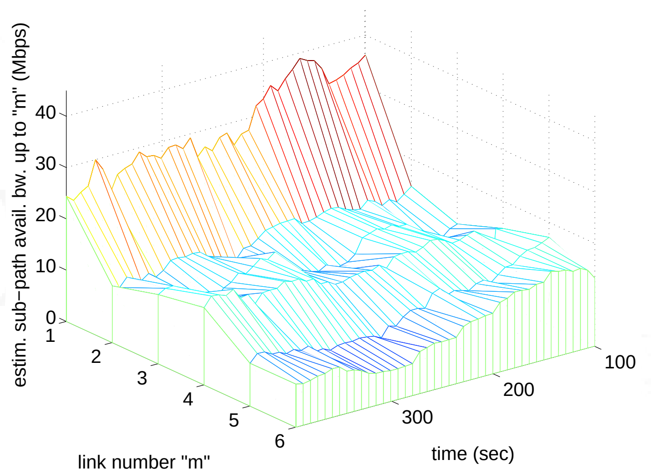

Truth

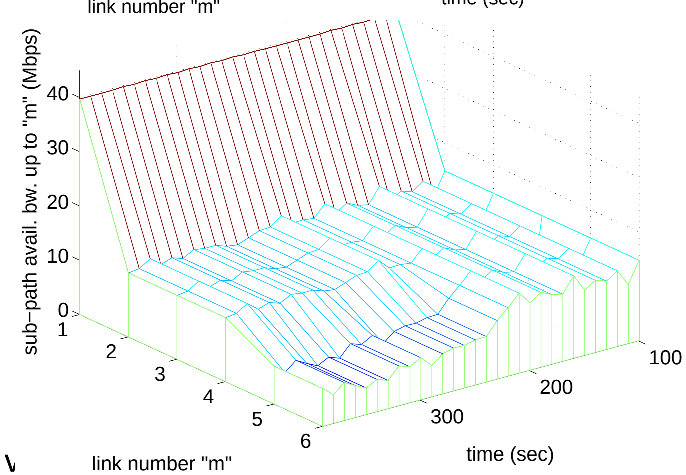

ns2 Simulation setting: Double web farm in ns2 (420 clients, 40 servers)

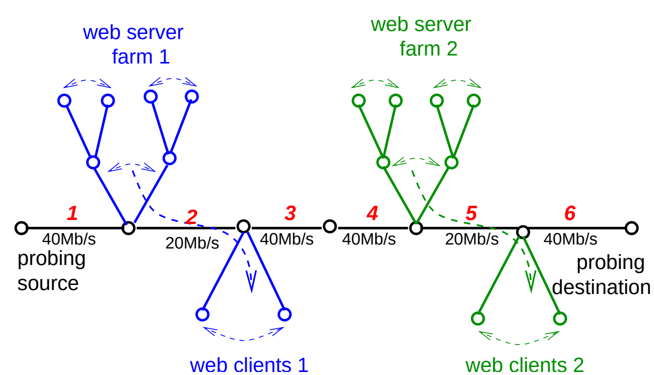

Queuing

Self-similar queuing

Large Deviation queuing

Multi-scale queuing

Queuing 101

- Reich's formula
  - Arriving traffic load

$$K_\tau[t] := \int_{\tau-t}^{\tau} X_\omega \mathrm{d}\omega$$

  - Queue length is (assuming the queue was idle at one point; follows from simple iteration)

$$Q_\tau := \sup_{t>0} \left(K_\tau[t] - ct \right)$$

Large Deviations for Qing

(Duffield-O'Connell, Norros): If λ is smooth

$$\lambda(q) := \lim_{t\to\infty} t^{-v} \log \mathbb{E} \exp(q t^{v-a} X_t)$$

then a Large Deviation Principle holds (LDP):

$$\lim_{b\to\infty} \log b^{-v/a} P[\sup_{t\geq 0} X_t > b] = -\inf_{c>0} c^{-v/a} \lambda^*(c)$$

$X_t := \mathrm{fBm}_H(t) - \mu t$:

$$P[\sup_{t\geq 0} X_t > b] \simeq \exp(-\mathrm{const} \cdot b^{2-2H})$$

(choose $a = 1$, $v = 2 - 2H$)

$H = 1/2$: 'classical' result

Multiscale Queuing

- Delivers non-asymptotic approximations
  - Can be estimated from true traffic using only few observations
  - Can be used to estimate queues of analytical models such as MWM and FGN-tree
- Shows that *exponential* times are *optimal* for computing the queue length
  - In the sense of being sparsest while keeping the correct asymptotic
- Queuing formula from traffic arrivals at dyadic scales (depends on more than LRD!):

$$\mathrm{MSQ}(b) := 1 - \prod_{j=0}^{n} P(K_{2^j} < b + c2^j) \geq P(Q_D > b)$$

$$\approx P(Q_0 > b)$$

MSQ in simulation

- Plot $\log P(Q > b)$ vs. b

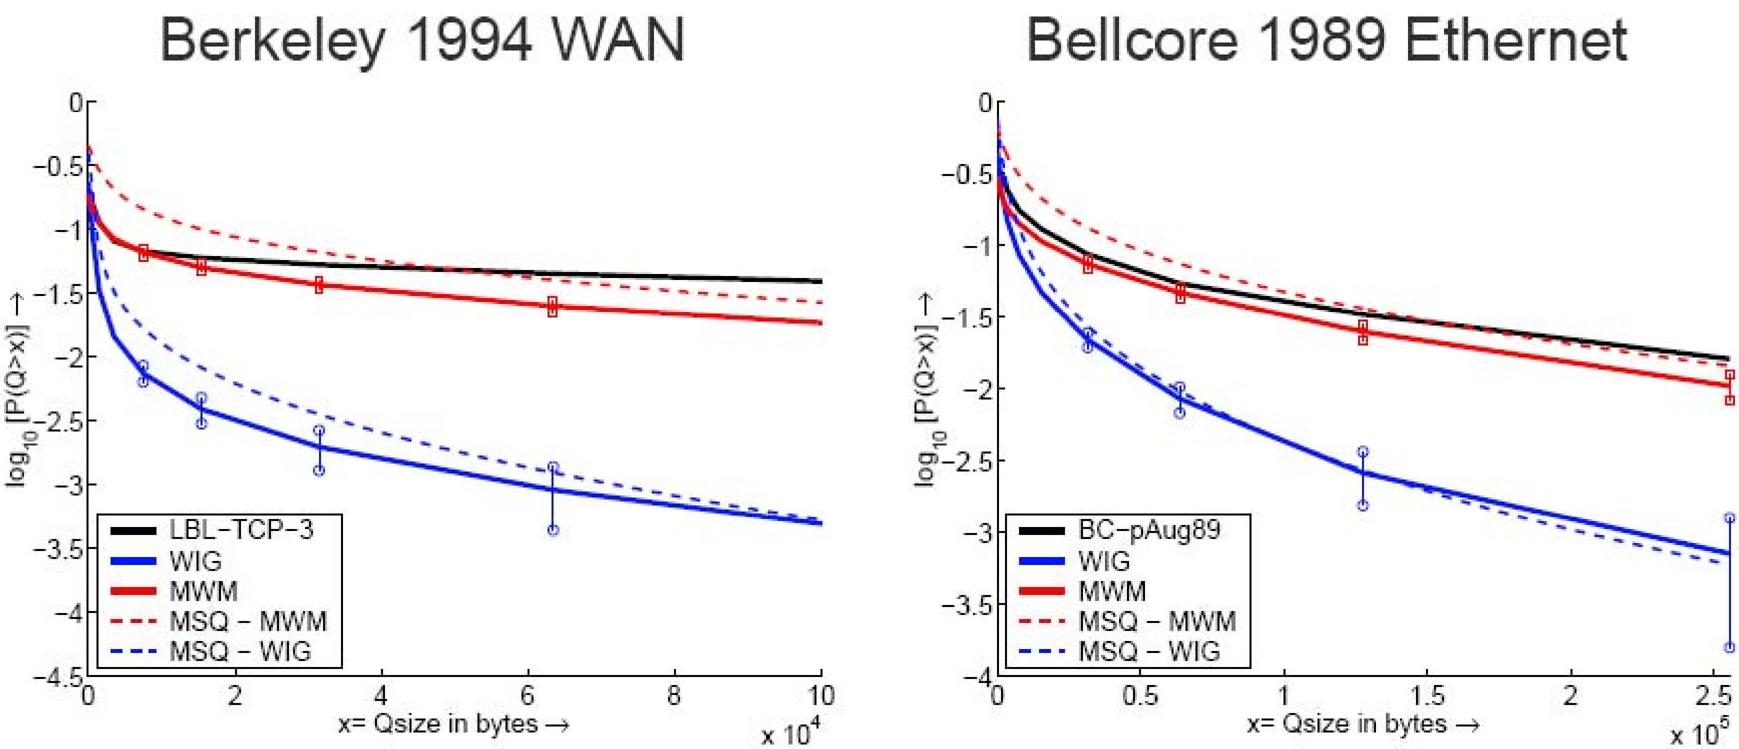

- MSQ is a close approximation

There is more to Qing than LRD

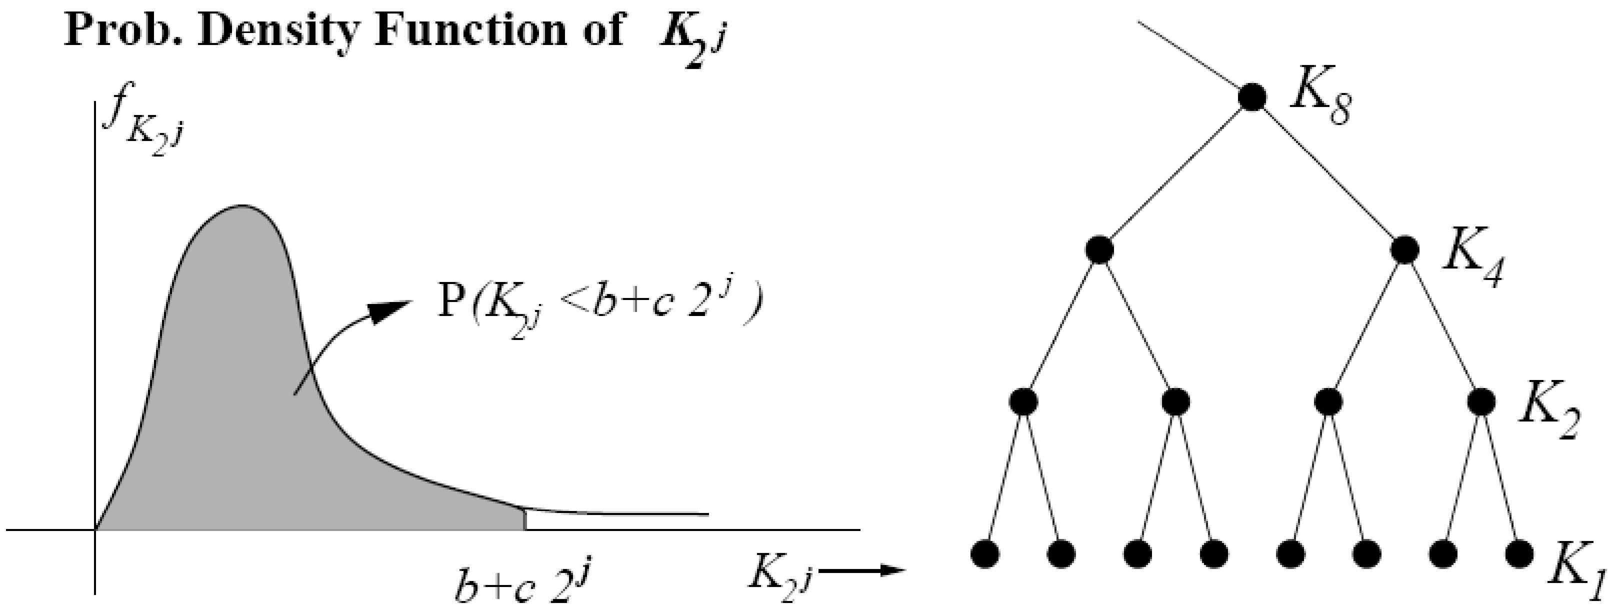

$$\text{MSQ}(b) = 1 - \prod_{i=0}^{n} P(K_{2^j} < b + c2^j)$$

- $P(K_{2^j} < b + c2^j)$ only roughly characterized by $\text{Var}(K_{2^j})$ (LRD)

- marginal tails influence MSQ