Fractals in Networking: Modeling and Inference

Rolf Riedi

R. Baraniuk
A. Keshavarz-Haddad, V. Ribeiro, S. Sarvotham
spin.rice.edu

Fractal2004, Vancouver, April 2004

www.caida.org

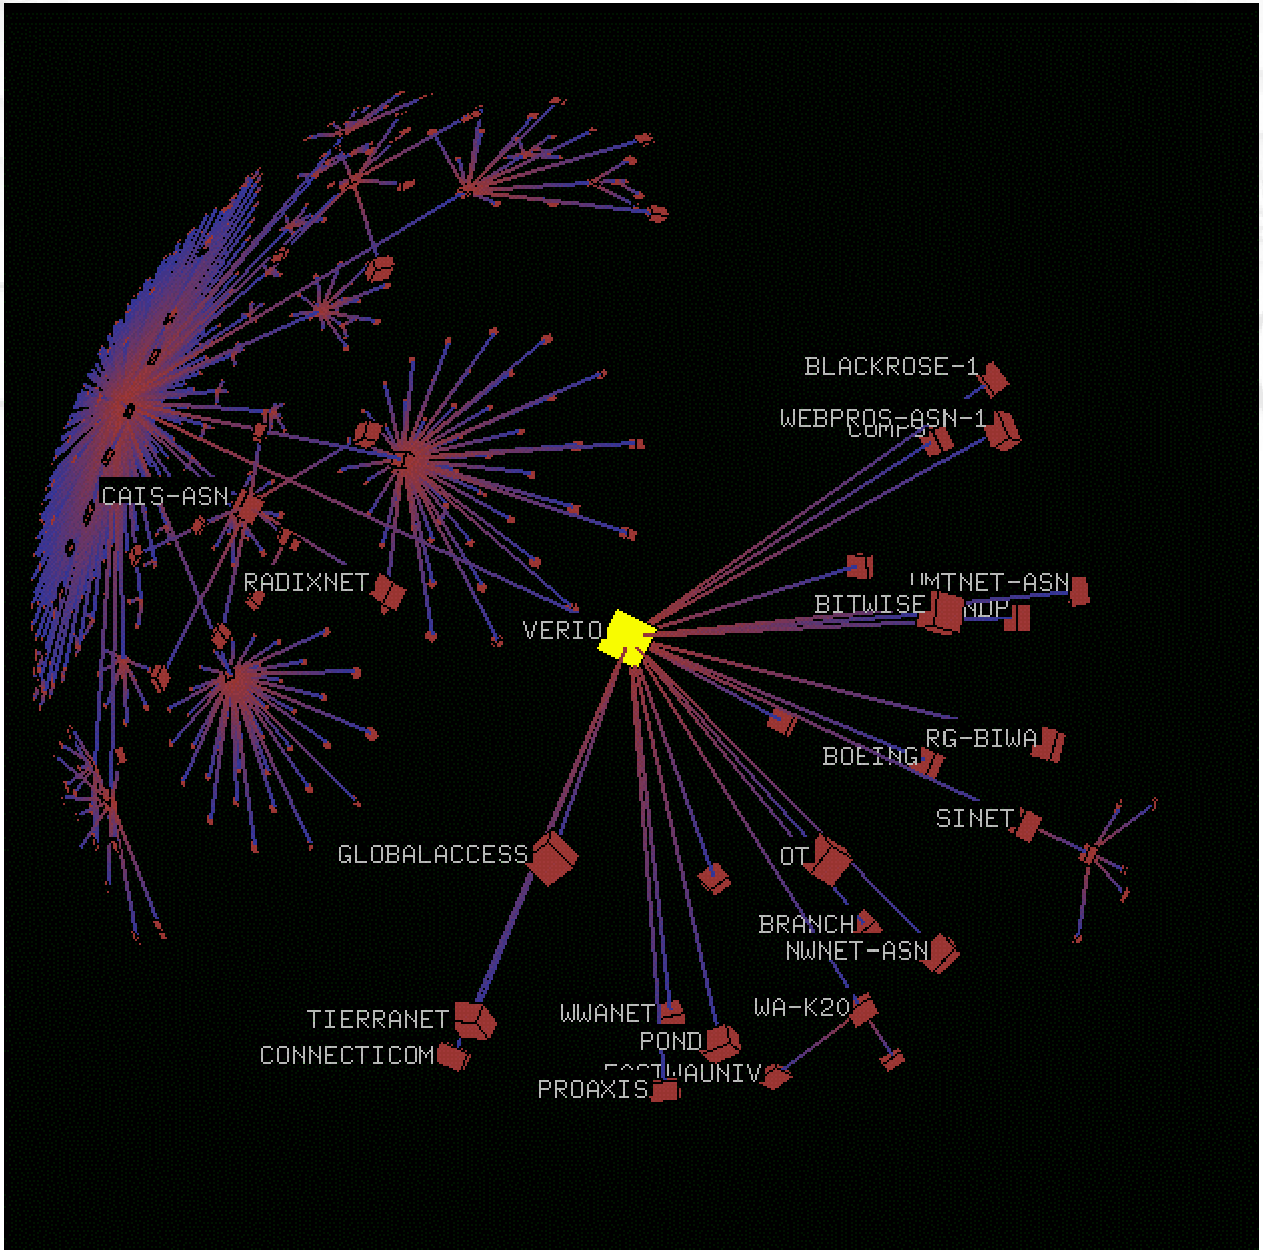

Internet is packet switched

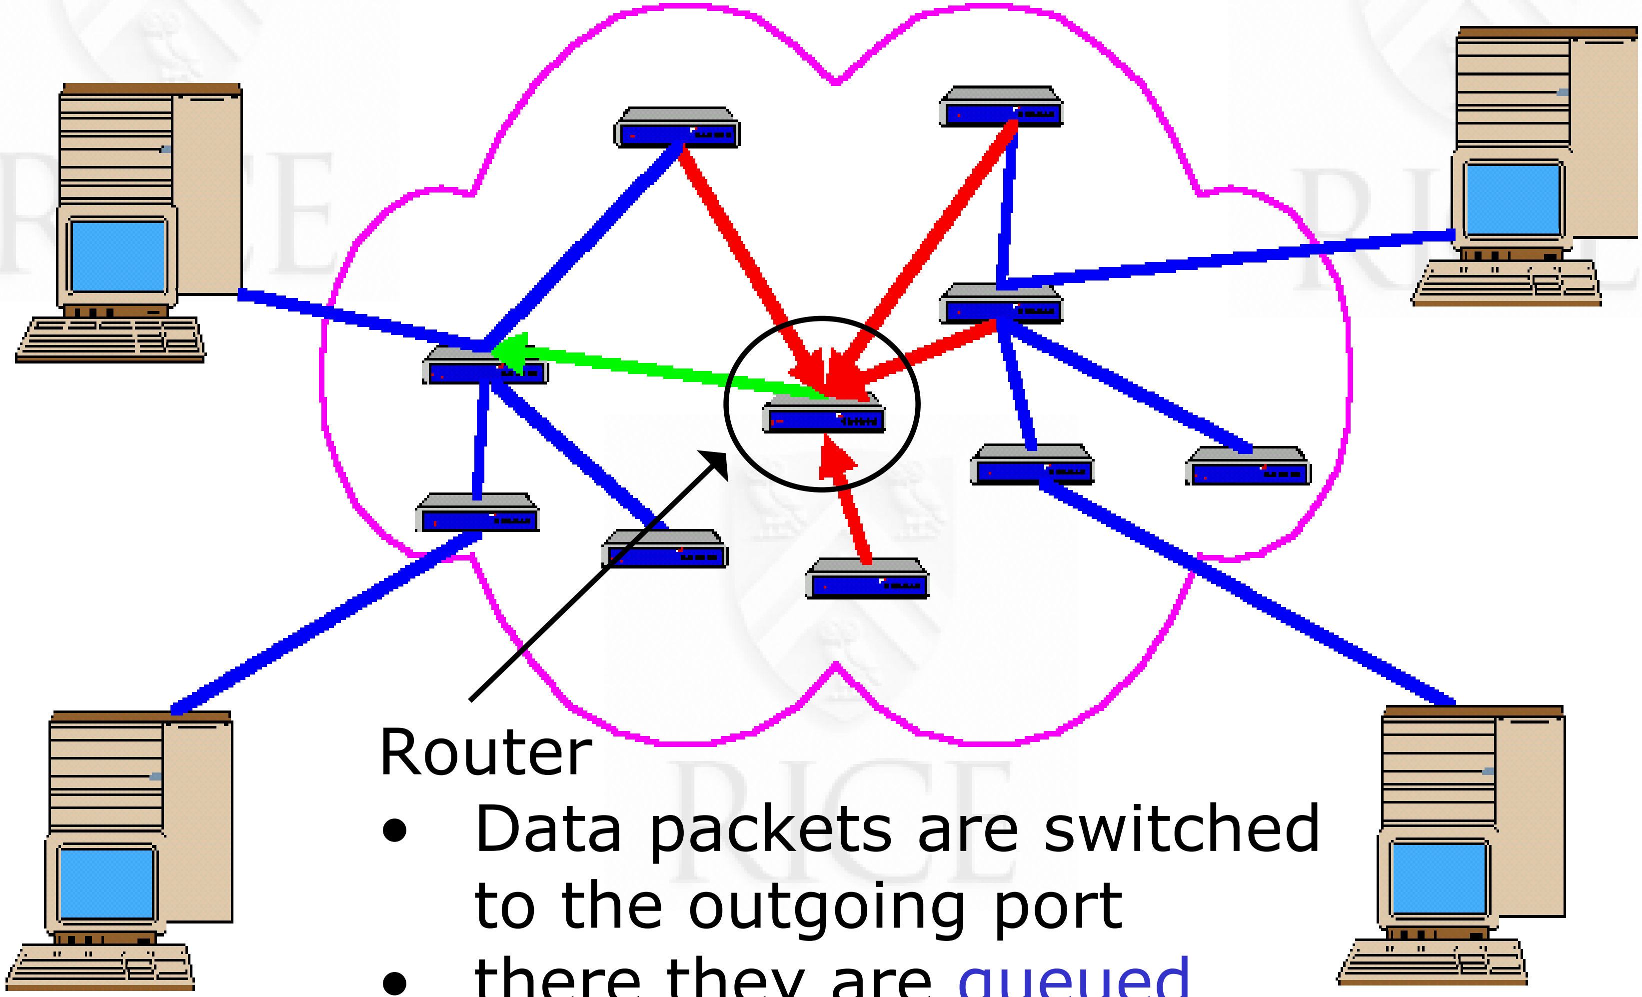

Router

- Data packets are switched to the outgoing port
- there they are queued
- …or dropped.

Measured Data

- Time series (A_k, Z_k) collected at gateway of LAN
  - k= number of data packet
  - A_k= arrival time of packet
  - B_k= size of packet

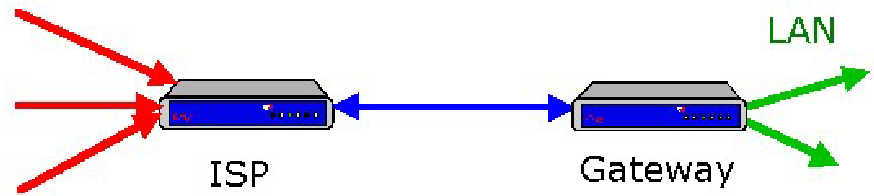

- Working data:

Bytes arriving in intervals of m time units

$$X_n^{(m)} = \sum_{mn+1<A_k<m(n+1)} B_k$$

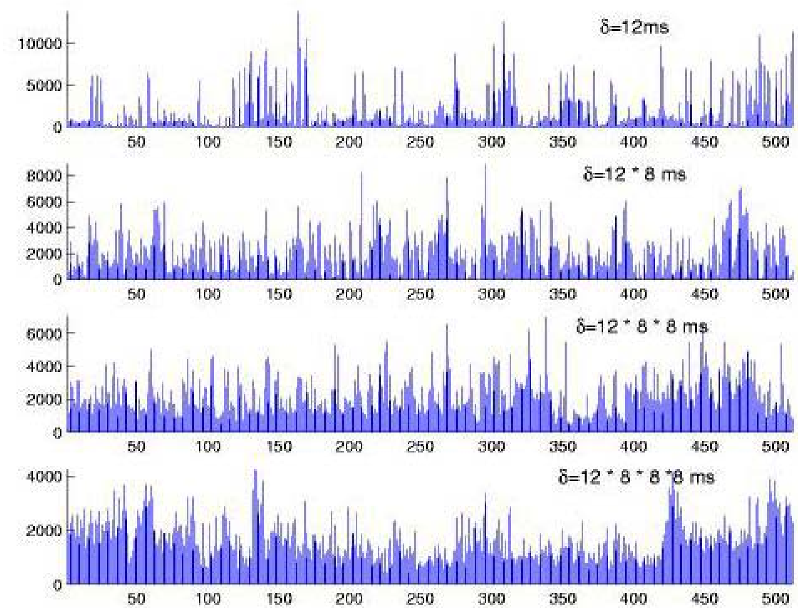

- Observation/Motivation:

High variability of $X_n^{(m)}$ for large m

degrading performance (loss, delay)

Network traffic is LRD
(Long Range Dependent)

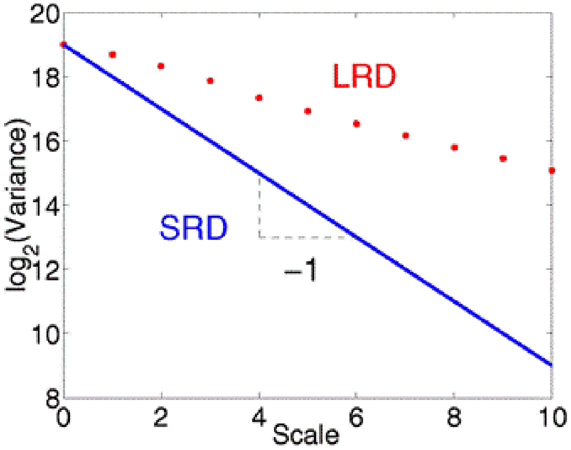

Bellcore '89

Network traffic is LRD

- Failure of Poisson modeling (Paxson Floyd 1995)
- Discovery of LRD -- Mathematical model

Source i: $$X_i(t) = \begin{cases} 1 & \text{if time } t \text{ is an } \textit{on} \text{ interval,} \\ 0 & \text{if time } t \text{ is in } \textit{off} \text{ interval.} \end{cases}$$

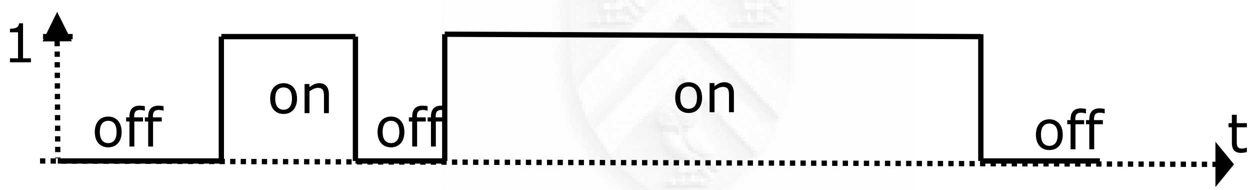

Heavy tailed on- and off-durations, e.g.,

$$1 - F_{on}(x) \sim \ell_{on} x^{-\alpha_{on}}, \; 1 < \alpha_{on} < 2$$

(Willinger Taqqu Wilson Leland '93, Renewal reward: Mandelbrot 60's)

Limiting behavior

Multiplexed Traffic: $$W^{(m)}(t) = \sum_{i=1}^{m} X_i(t)$$

$$\frac{W^{(m)}(t) - \mathsf{E}W^{(m)}}{m^{1/2}} \stackrel{fdd}{\longrightarrow} G(t)$$

$$\frac{1}{T^H}\int_0^{Tt} G(u)du \stackrel{fdd}{\longrightarrow} \sigma B_H(t)$$

$$H = \frac{3 - \min(\alpha_{on}, \alpha_{off})}{2}$$

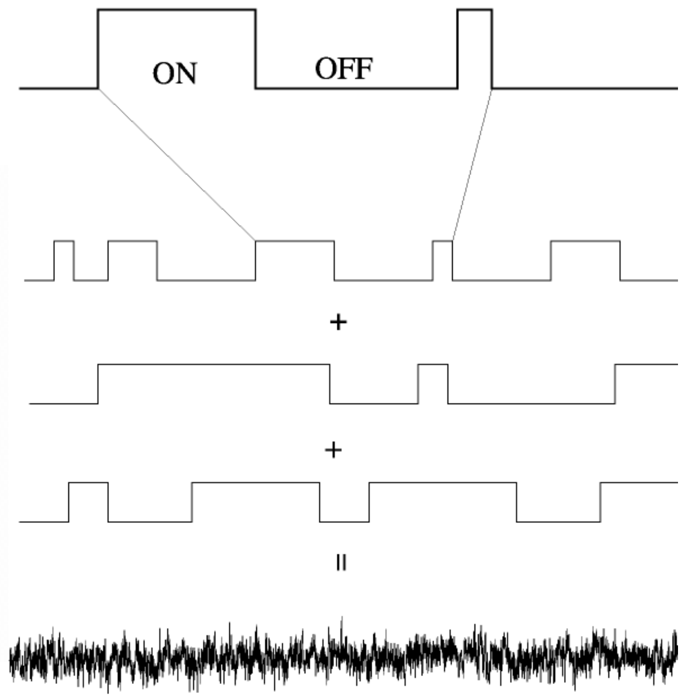

(Lamperti 1962)
(Taqqu Levy 1986)

Physical model

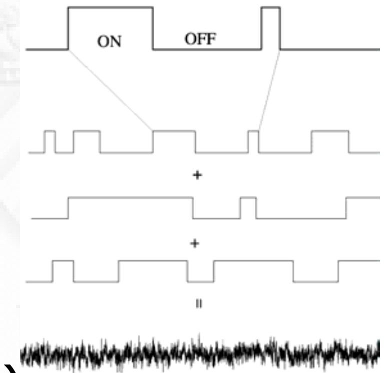

- On-off model: physically appealing
- Verification
  - Heavy tailed file sizes on web (Crovella '96)
  - Client behavior
- Implications
  - Using self-similarity: (Norros '94)
    To reduce overflow increase bandwidth, not buffer
  - Using large deviations: (Duffield '95)
    Queue-tail is Weibull $P[Q > b] \simeq \exp(-b^{2H})$
  - Predictive control
- Huge success

Voice: Heavy tales are rare
Data: Heavy tails are there

Understanding Large Scale

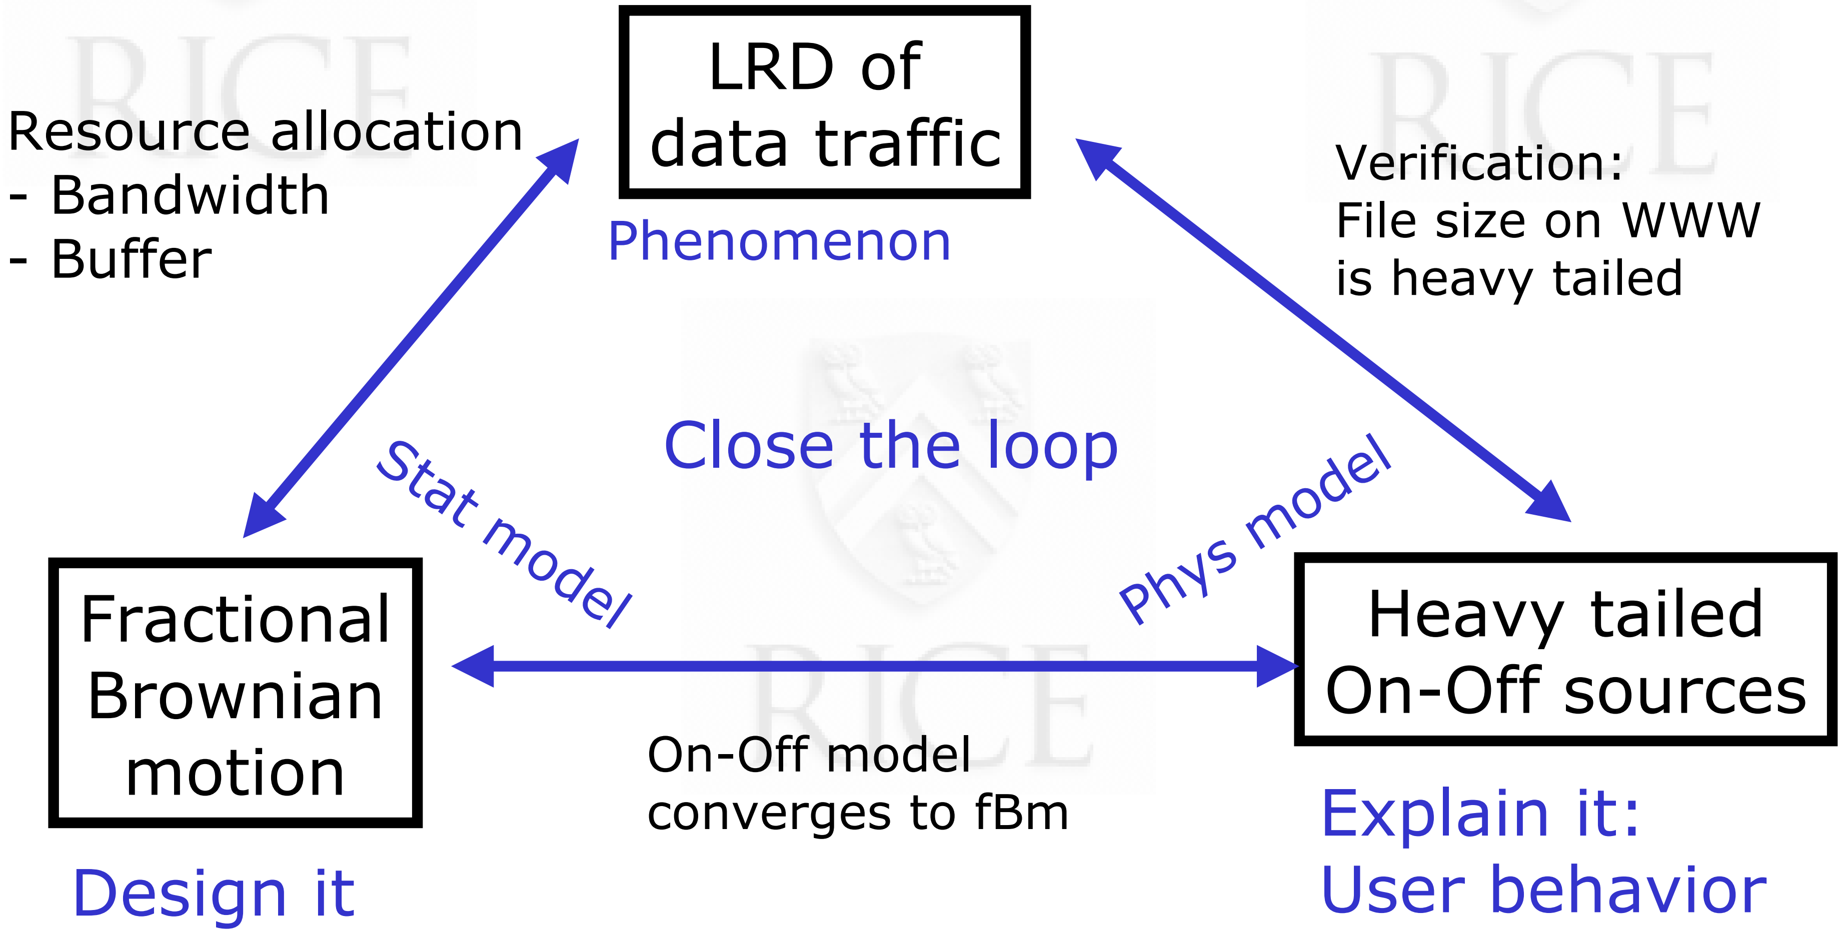

Small scale behavior

Below Round Trip Time

- fBm is realistic only at large scales
  - Explains how to design
- …but is unrealistic at the small time scales relevant for
  - Control,
  - Performance (jitter, delay)
  - Service level agreements
  - Quality of Service (QoS)

Network Traffic is Multifractal

- Visually striking
- Scaling of impressive quality
  (Levy Vehel & RR ’96, Norros & Mannersalo ’97, Willinger et al ‘98)
- Statistical models:
  - Binomial cascades
    (Crouse & RR ’98, Willinger et al ‘98)
  - On-off cascades (Baras)
  - Stationary Products
    (Norros Mannersalo RR ‘99)
  - Products of Cylindrical Pulses
    (Barral Mandelbrot ‘99)
  - IDC
    (Bacry-Muzy ‘02, Chainais Abry RR)

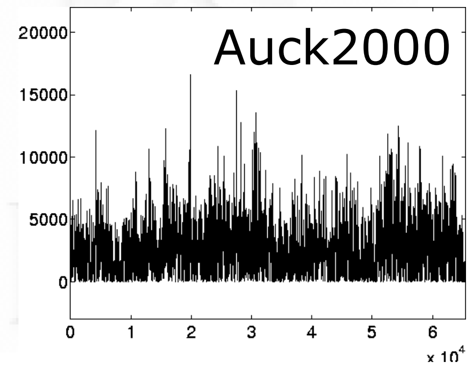

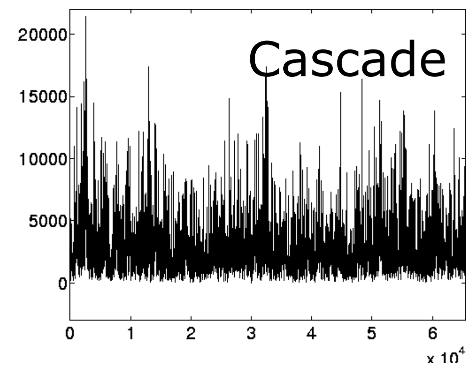

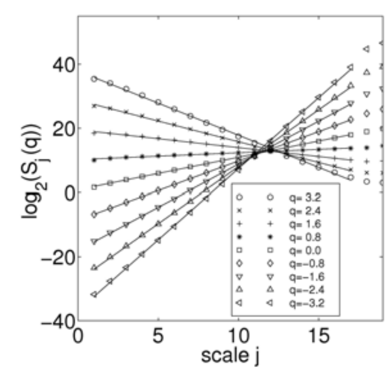

Multiscale statistical modeling

Why multiplicative?

Multiscale Modeling

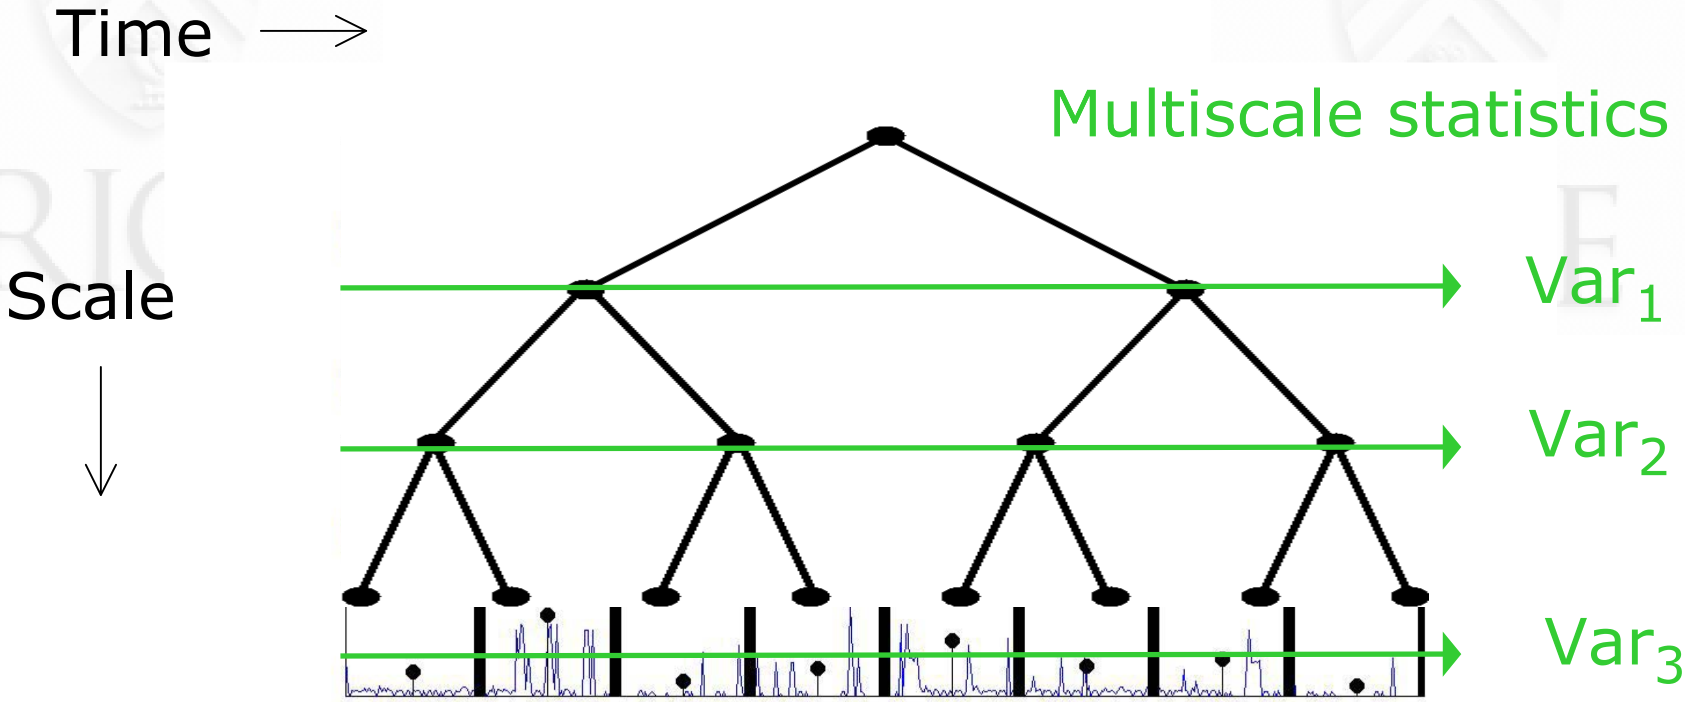

Analysis: flow up the tree by adding

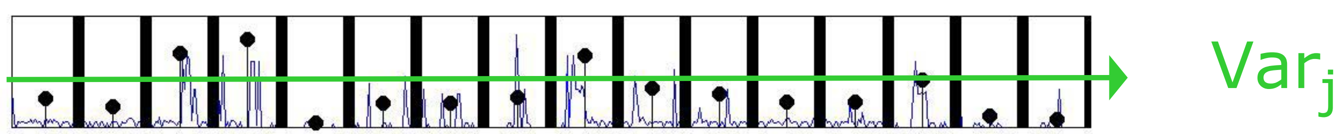

Start at bottom with trace itself

Multiscale Modeling

Time →

Start at top with total arrival

Scale ↓

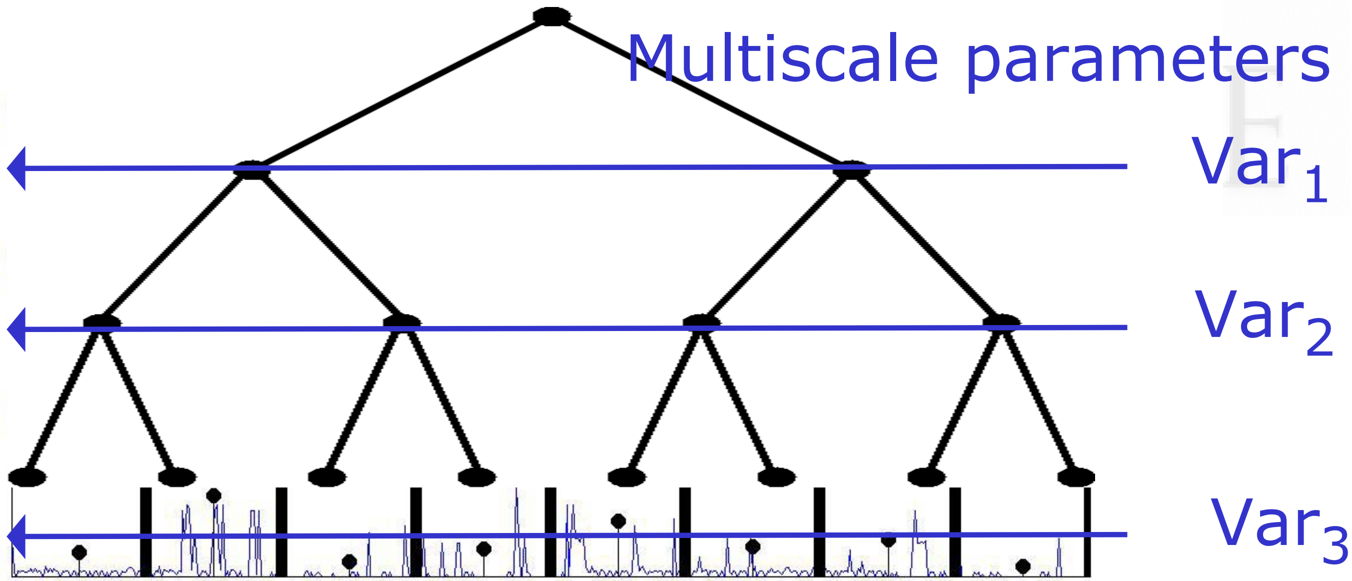

Synthesis: flow down via innovations

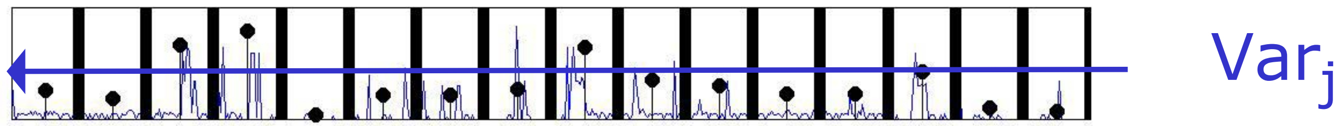

Signal: bottom nodes

Additive innovations: Linear Processes

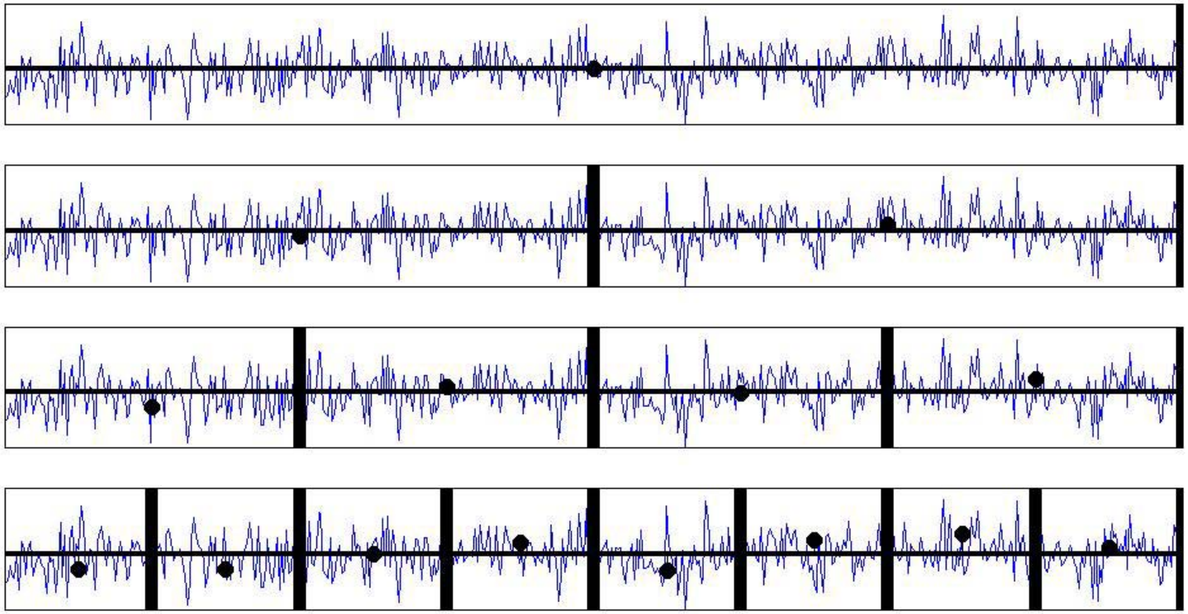

Match variances on all dyadic scales
CLT: asymptotically Gaussian

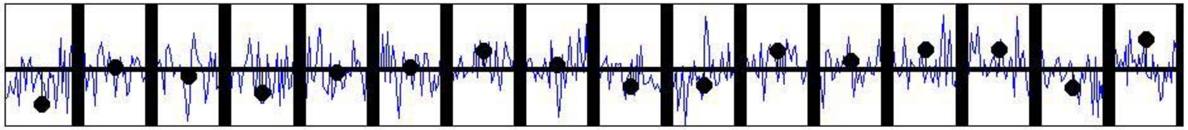

Additive Innovations $W_{jk} \sim \mathcal{N}(0, \sigma^2 2^{-j(2H-1)})$: Model for $B_H(t)$

Multiplicative Wavelet-Model: MWM

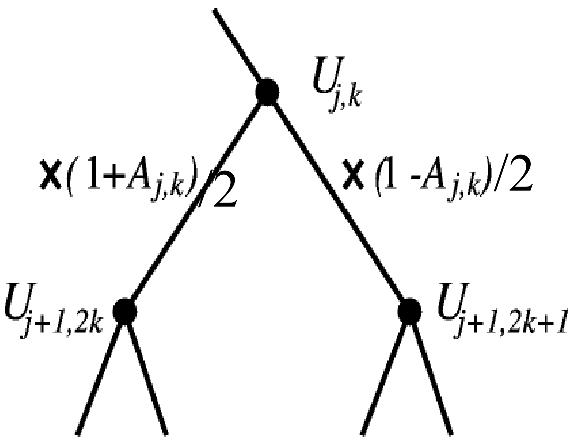

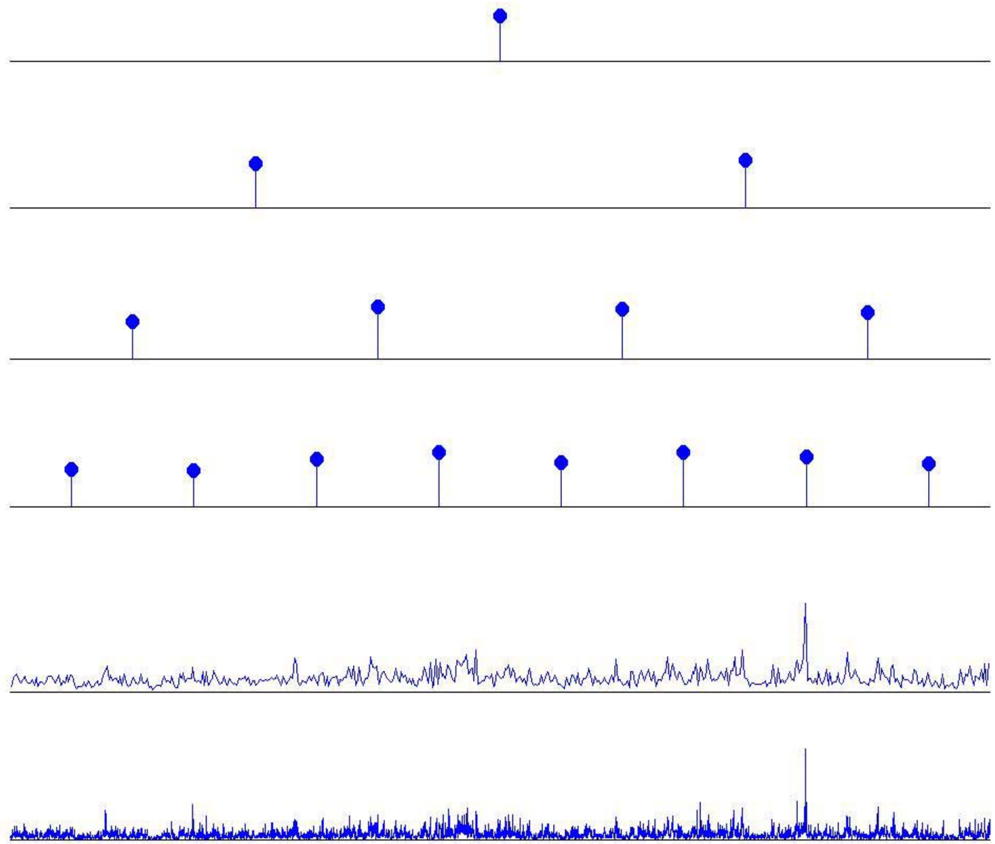

Multiplicative Innovations $(1 \pm A_{jk})/2 \sim \mathcal{B}eta\,(1/2, \sigma_j)$: Variance-Match

Statistical relevance
of cascades in networking

scale | Real trace (Auck00) | Cascade | Gaussian (fGn)

4ms

16ms

64ms

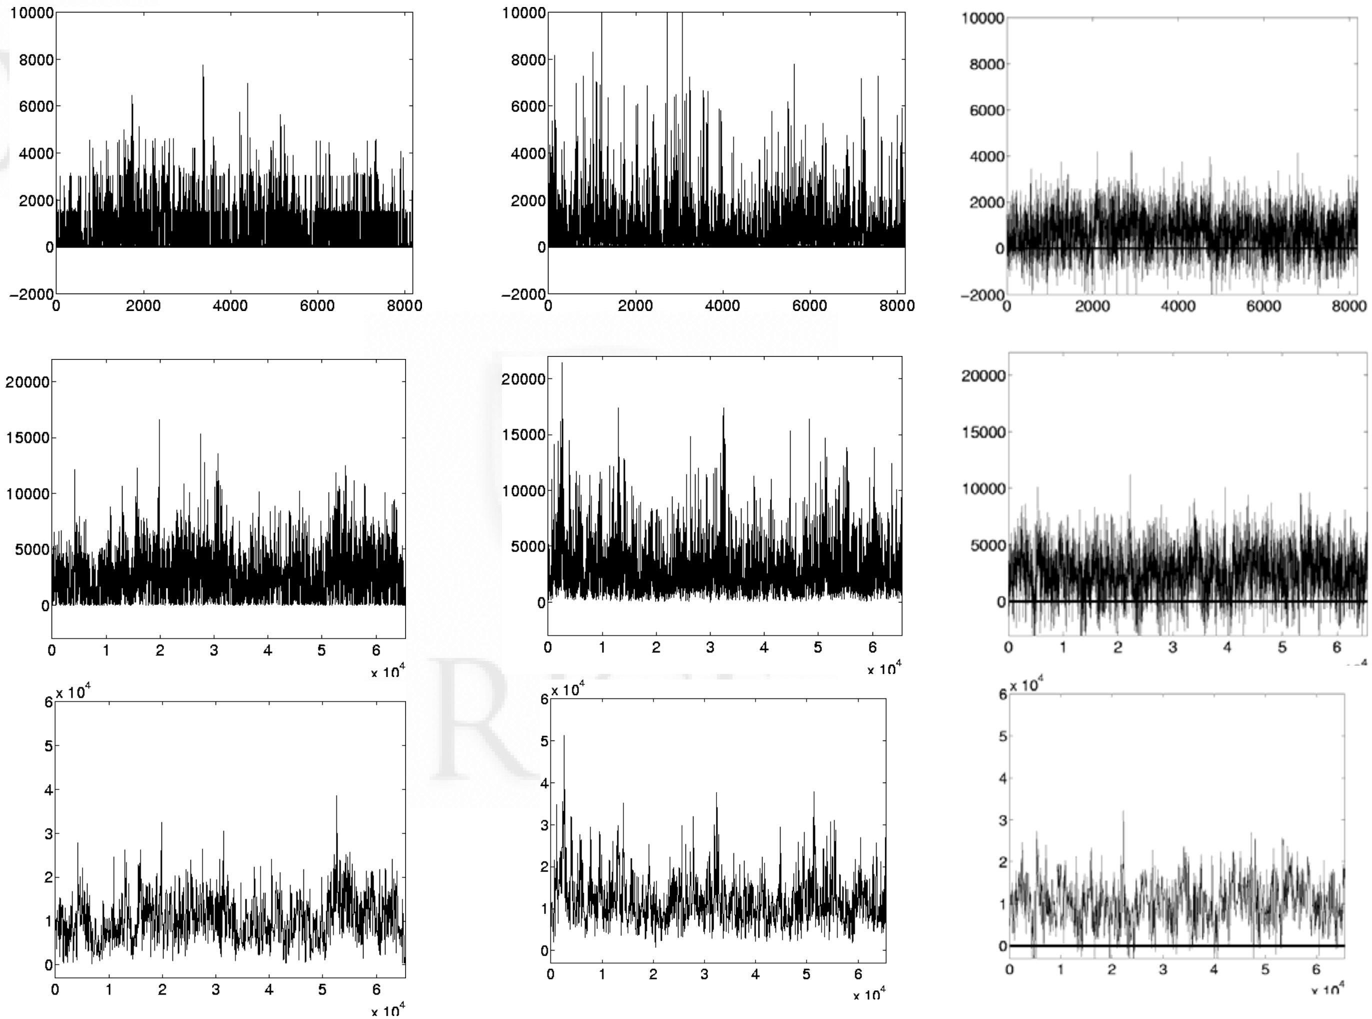

Equal variance on all scales

Statistical relevance

of cascades in networking

scale | Real trace (Auck00) | Cascade | Gaussian (fGn)

4ms

16ms

64ms

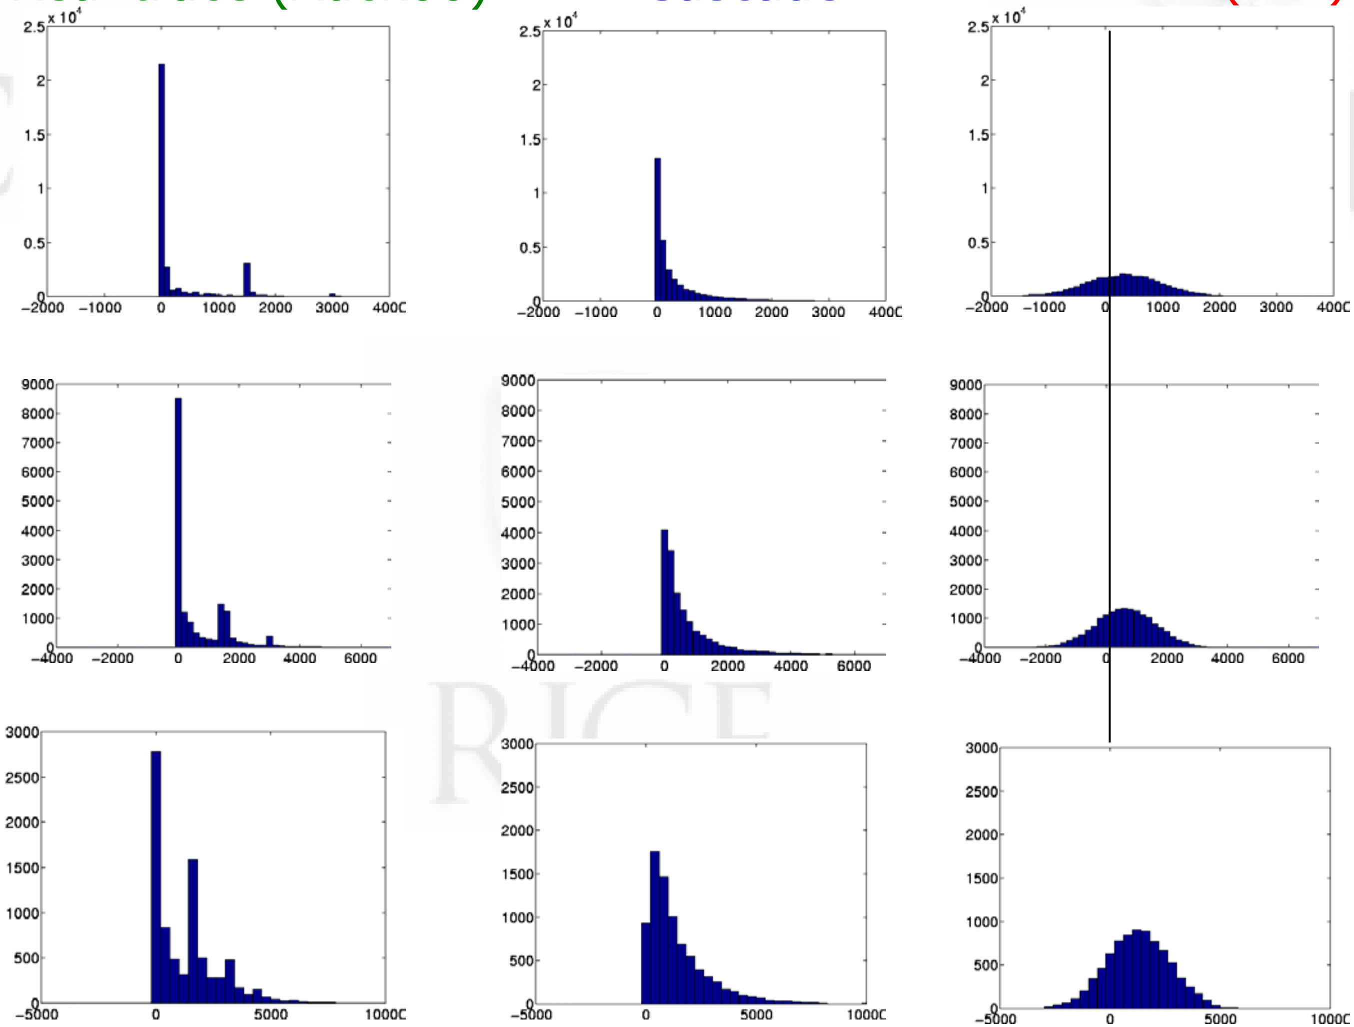

Equal variance on all scales

Own it:

Queuing
Traffic inference

Queuing 101

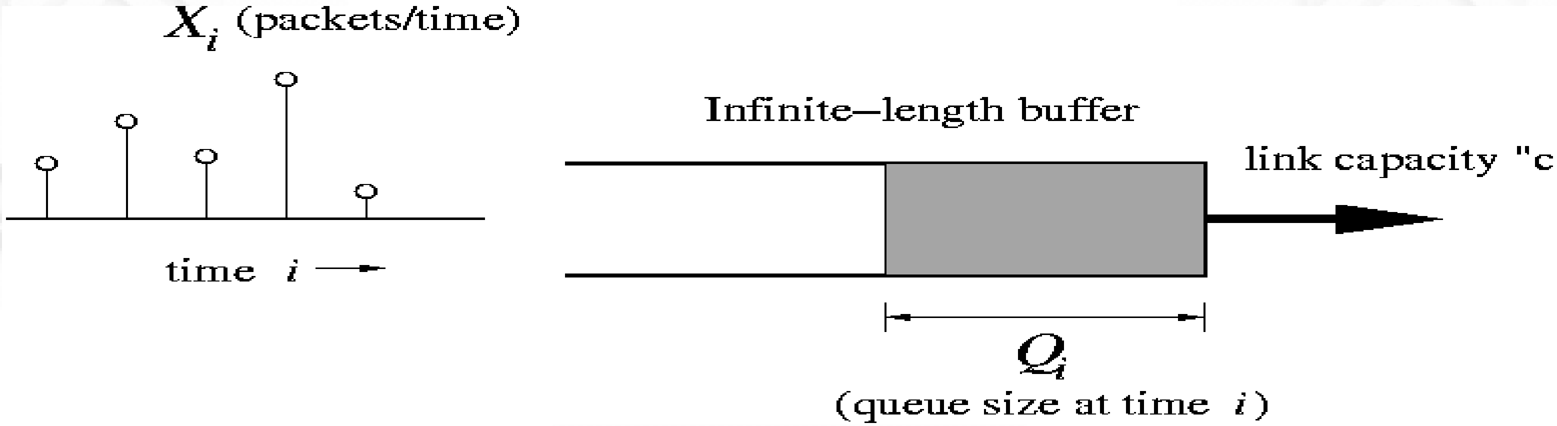

- Lindley's recursion

Q_i: queue size at time i

X_i: traffic arriving at time bin i

$$
\begin{aligned}
Q_0 &= \max(Q_{-1} + X_0 - c, 0) \\
&= \max(\max(Q_{-2} + X_{-1} - c, 0) + X_0 - c, 0) \\
&= \max(Q_{-2} + X_{-1} + X_0 - 2c, X_0 - c, 0)
\end{aligned}
$$

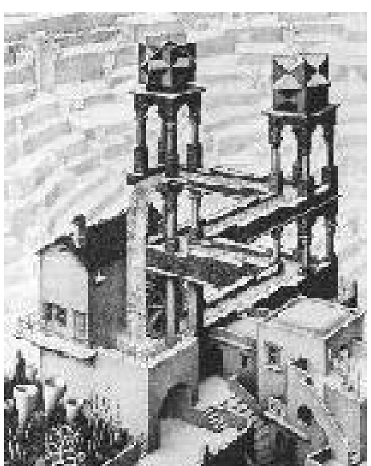

$$Q_0 = \max_n [X_1 + \dots X_n - nc]$$

MultiScale Queuing approach

For tree models of traffic:

- Innovations W_i between dyadic times are independent (pass to log for multiplicative model)

Lemma: $E_i := \{W_0 + \ldots + W_{i-1} < b_i\}$, W_i *are independent, otherwise arbitrary. Then*

$$P[E_i | E_{i-1}, \cdots, E_0] \geq P[E_i].$$

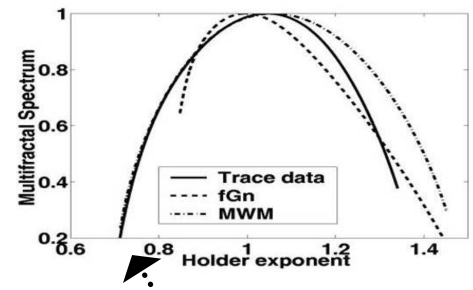

Multifractal analysis, (large deviations): If mfr spectrum has positive zeros then MSQ converges exponentially fast

Dyadic queue tails

$$\begin{aligned} P[Q_D > b] &= 1 - P[Q_D < b] = 1 - P[\cap_{i=0}^{n} E_i] \\ &= 1 - P[E_0] \prod_{i=1}^{n} P[E_i | E_{i-1}, \cdots, E_0] \\ &\leq 1 - \prod_{i=0}^{n} P[E_i] =: \mathrm{MSQ}_n(b). \qquad (1) \end{aligned}$$

Queuing analysis

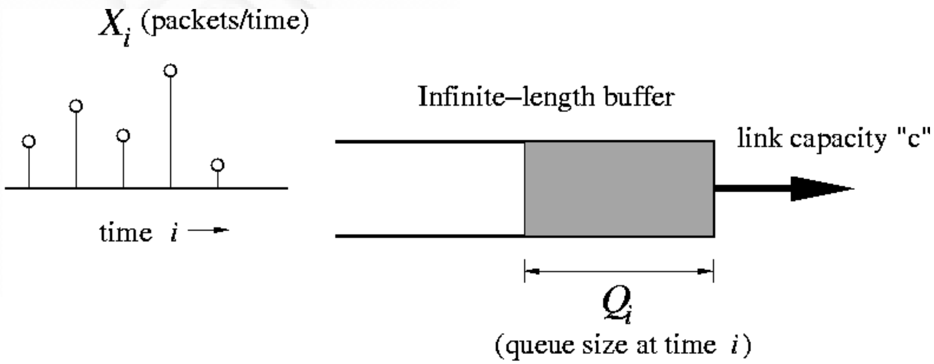

Q-tail: P[Q>b]

- Tree structure allows for analytical queuing formula
- Multiplicative model superior to additive
- Importance of multiscale tails (not only second order statistics)

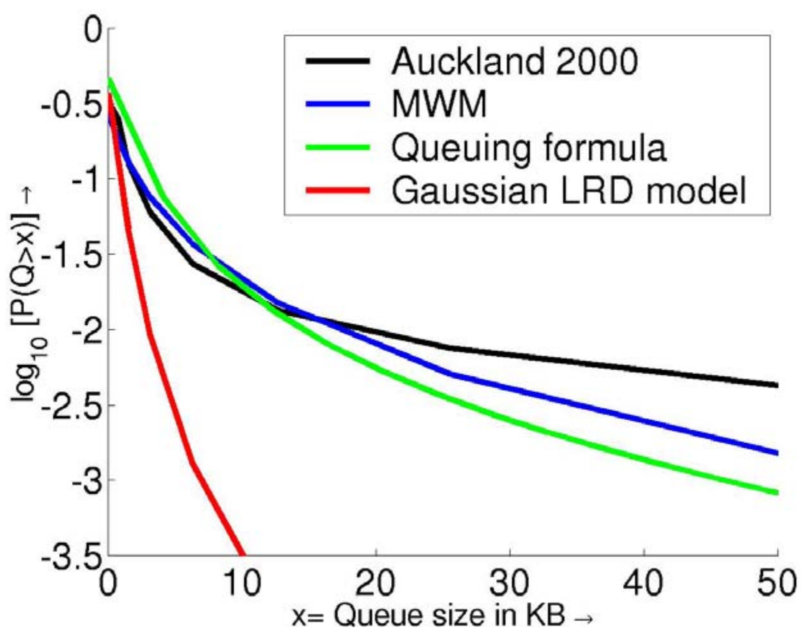

InfoComm2000, Transactions on Networking

Efficient Probing: *Packet Chirps*

- Tree inspires geometric chirp probe
- MLE estimates of cross-traffic at multiple scales

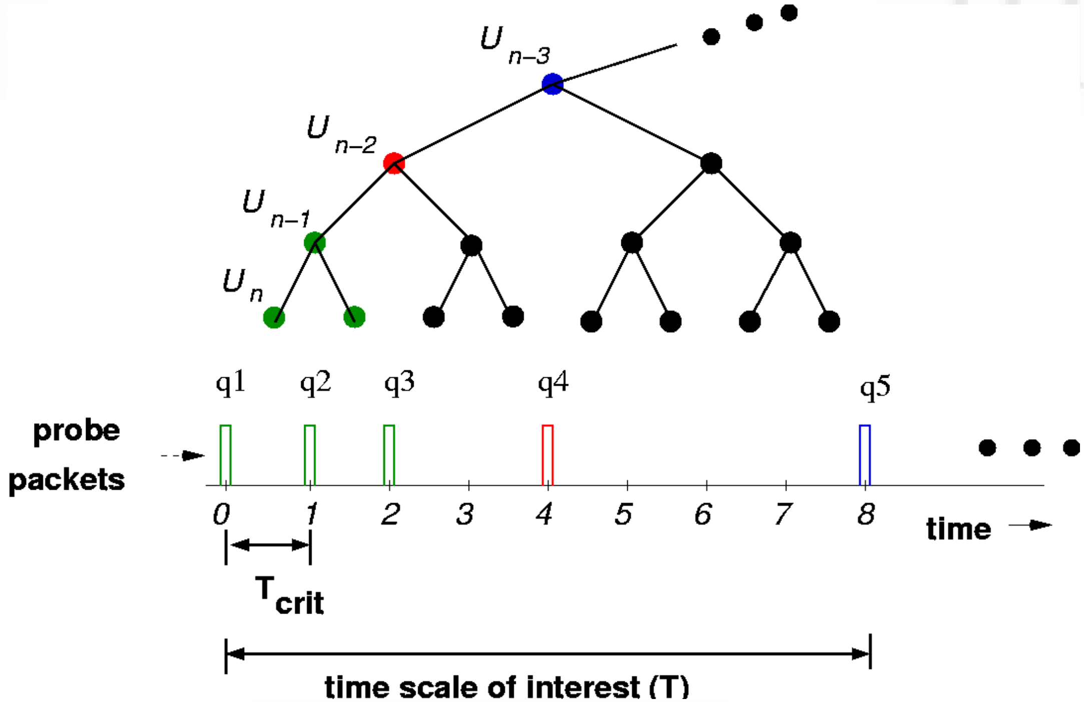

- Chirp is practical (ITC2000, Best paper PAM2003)
- Uniform is optimal spacing (Lehmann symposium 2004)

Statistical scaling at fine resolution

Beyond powerlaws

Real world data

- is stationary
- can deviate from powerlaws: traffic
- has no preference for dyadic scales

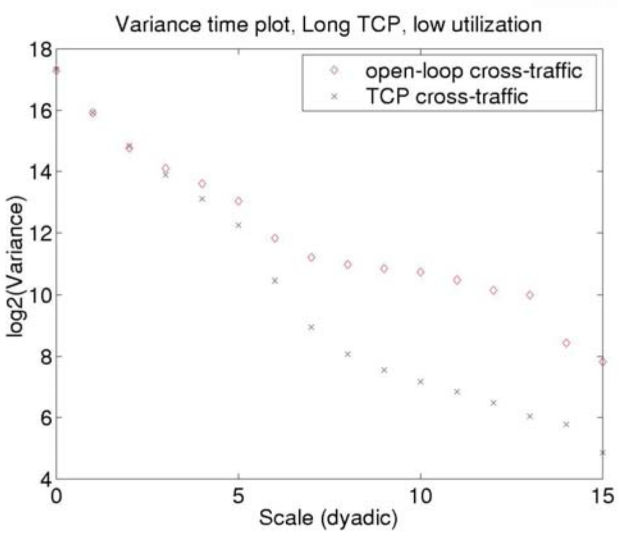

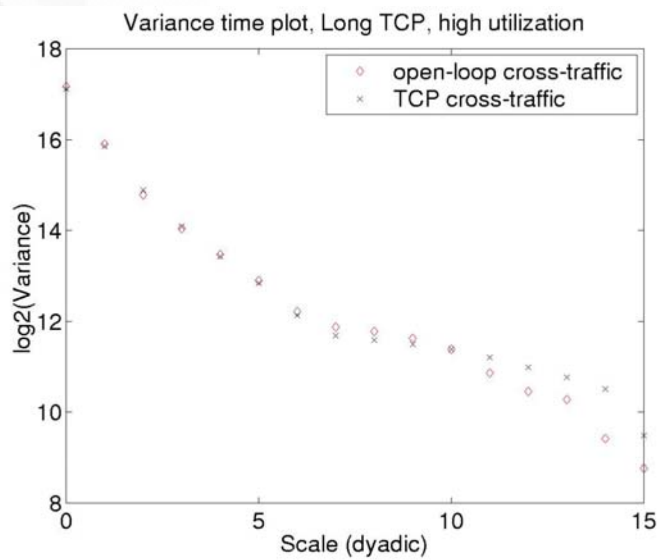

Beyond Self-similarity

- Self-similarity revisited:
  - $B(at) =^d C(a)\, B(t)$ B: process, C: scale function
  - $B(abt) =^d C(a)C(b)\, B(t)$
  - $C(a)C(b)=C(ab) \rightarrow C(a) = a^H$
  - $E[|\, B(a^n)\, |^q] = c(q)\, (a^{qH})^n$
  - linear in q (mono-fractal)

Beyond Self-similarity

- Self-similarity revisited:
  - $B(at) =^d C(a) B(t)$ B: process, C: scale function
  - $B(abt) =^d C(a)C(b) B(t)$
  - $C(a)C(b)=C(ab) \rightarrow C(a) = a^H$
  - $E[| B(a^n) |^q] = c(q) (a^{qH})^n$
  - linear in q (mono-fractal)
- More flexible rescaling "Ansatz":
  - $C=C(a,t)$? : non-stationary increments
  - C=independent r.v. for every re-scaling :
  - $X(a\ldots at)= X(a^n t) = C_1(a)\ldots C_n(a) X(t)$: multiplicative
  - $E[| X(a^n) |^q] = c(q) E[| C(a) |^q]^n$
  - non-linear in q (multifractal)

Beyond power-laws

- Self-similarity revisited:
  - $B(at) =^d C(a) B(t)$ B: process, C: scale function
  - $E[| B(a^n) |^q] = c(q) (a^{qH})^n$ (mono-fractal)
- More flexible rescaling "Ansatz":
  - $X(a \ldots at) = X(a^n t) = C_1(a) \ldots C_n(a) X(t)$: multiplicative
  - $E[| X(a^n) |^q] = c(q) E[| C(a) |^q]^n$ (multifractal)
- Infinitely divisible scaling
  - $E[| X(a^n) |^q] = \exp(\nu(a^n) \xi(q))$
  - Multifractal for $\nu(x) = \log_a x$; mono-fractal for $\xi(q)=qH$
  - Interpretation:
    - $E[| X(a^n) |^q]$ is the Laplace trafo of $\log[X(a^n)]$
    - The distribution of $\log[X(a^n)]$ is an $\nu(a^n)$-fold convolution
    - Thus $X(a^n)$ itself is in distribution the $\nu(a^n)$-fold product of a unit multiplier W with = $E[| W |^q] = \exp(\xi(q))$
  - First model: Product of Cylindrical Pulses (Mandelbrot Barral)

Binomial Measure

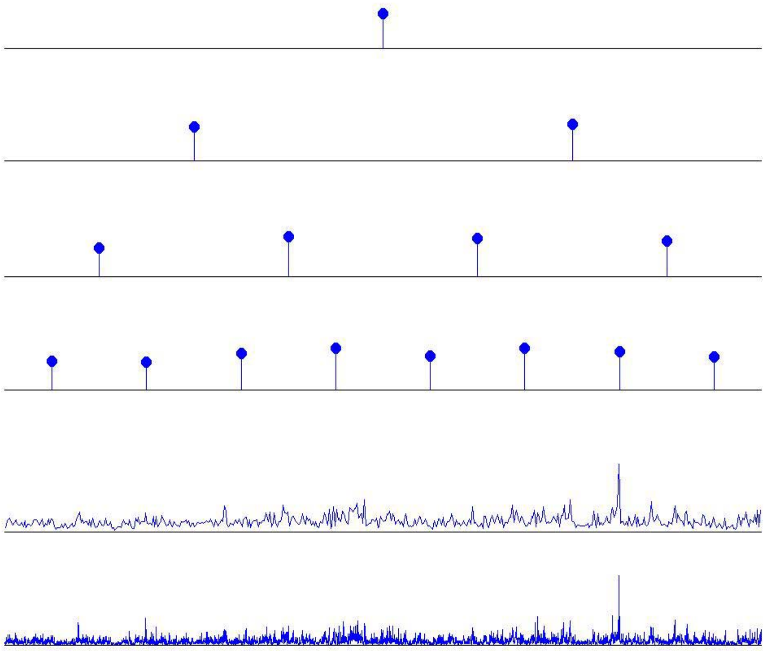

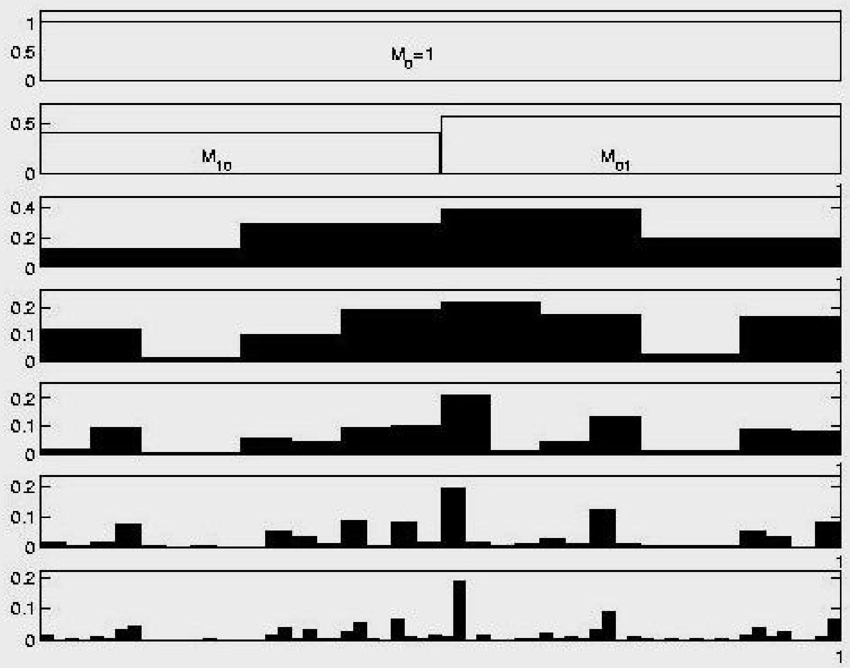

As a tree

As a cascade of multiplicative pulses

Geometry of Binomial Pulses

- Time-Scale plane: codes shape of pulses
  - **Position** (T=center)
  - **Size** (R=length)

Pulses:

$P_i(t) = W_i$ if $|t-t_i|<r_i/2$

$\quad\quad\quad\;\; 1$ else

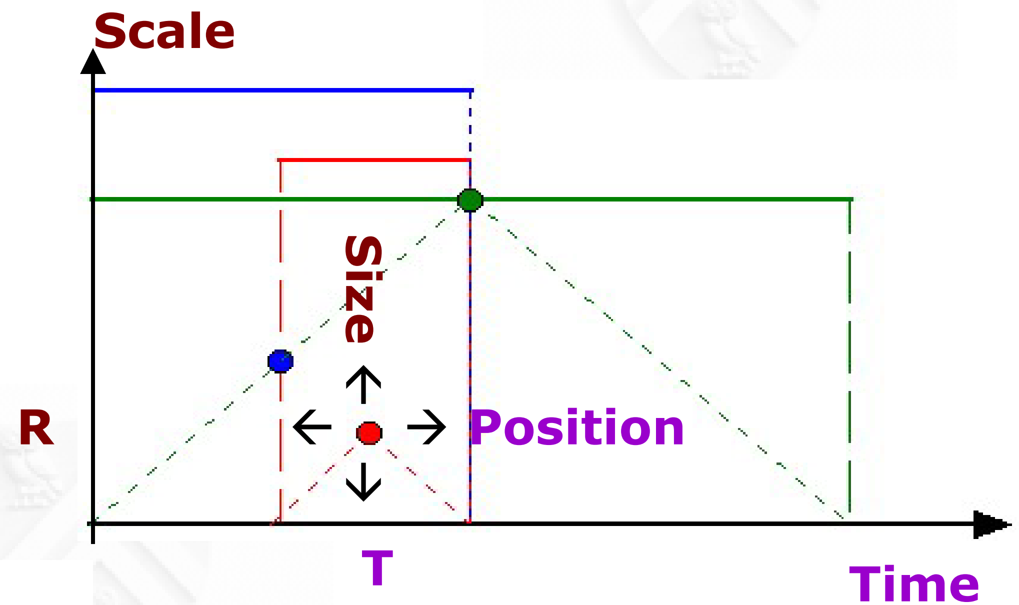

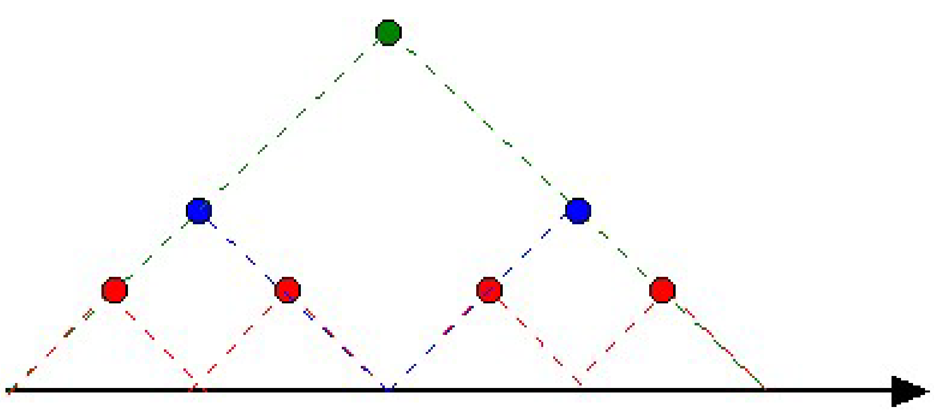

←

For Binomial: Strict dyadic geometry

Stationary geometry

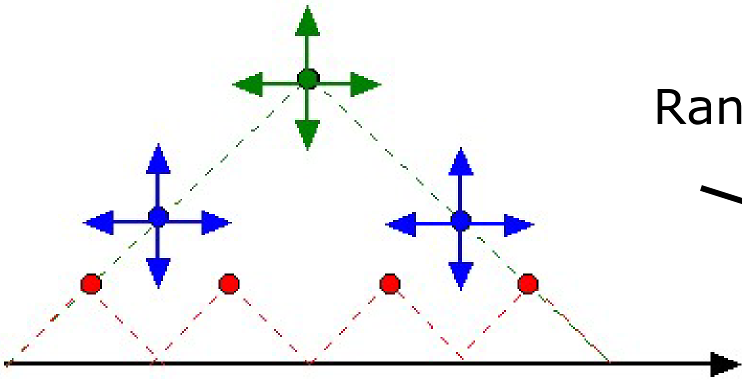

Randomize Positions
and Sizes

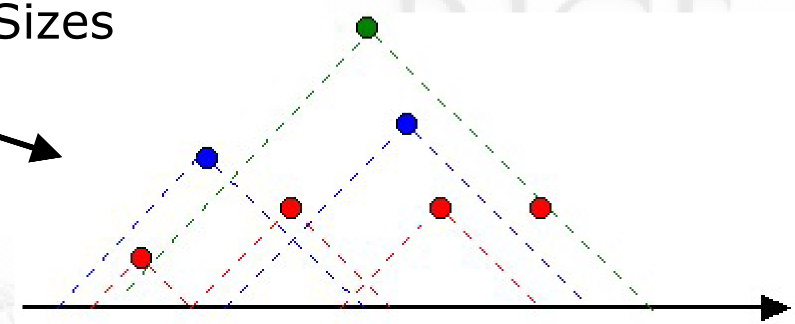

Large Scale Pulses

Medium Scale Pulses

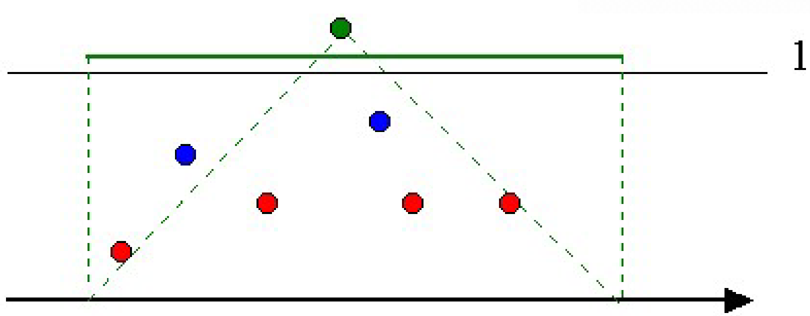

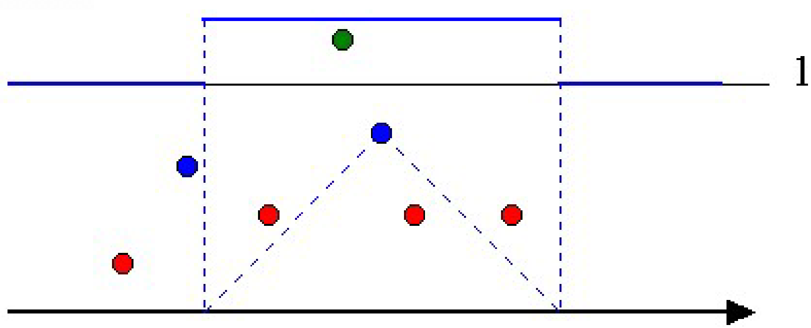

Compound Poisson Cascade

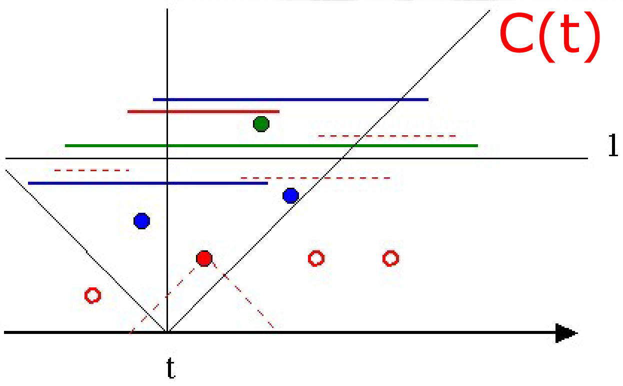

Poisson points with control measure m

Cone of influence at t
$C(t)=\{(t_i,r_i) : |t-t_i|<r_i/2 \}$

$Q(t) = \Pi \{W_i : |t-t_i|<r_i/2 \}$

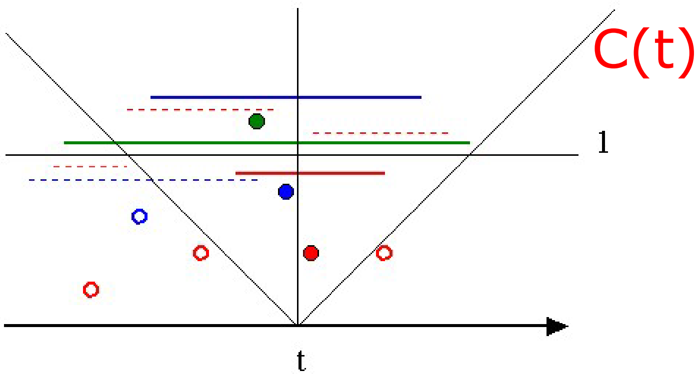

Multifractal formalism
(self-similar case) Barral Mandelbrot
Multifractal Scaling
(non-powerlaws) (with Abry & Chainais)

$T(q) = \exp[m(C(t,r))(1-E[W^q])]$

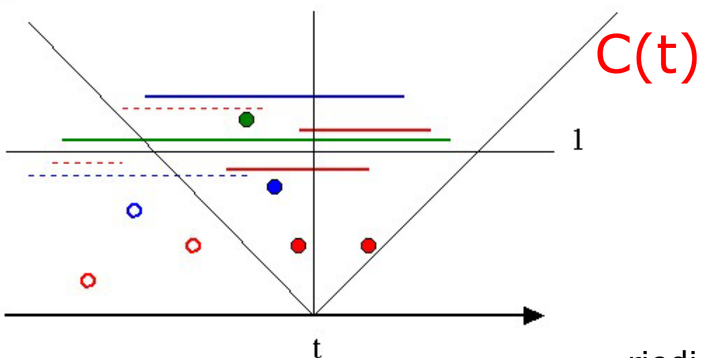

→powerlaw only if $m(C(t,r)) = -\log(r)$
→More general Infinitely Divisible Laws