Modeling Data Networks on Small and Large Time Scales

Sidney Resnick
Cornell University
School of Operations Research and Industrial Engineering
sid@orie.cornell.edu
http://www.orie.cornell.edu/∼sid
and
Krishanu Maulik
Department of Statistical Science
Cornell University
km75@cornell.edu

Traditional queueing theory has relied on

$\rightarrow$ Lots of independence; Poisson inputs,
$\rightarrow$ Light tails.

Empirical observations on networks reveal features inconsistent with traditional assumptions:

$\rightarrow$ self-similarity (ss) & long-range dependence (LRD) of various transmission rates:

packet counts/time in LANS, WANS,
www bits/time (web downloads)

$\rightarrow$ heavy tails

file sizes,
transmission rates,
transmission durations,
Unix job processing times,
call lengths
Packet interarrival times.

Origins: Bellcore study in early 90's revealed *packet counts per unit time exhibit self similarity and long range dependence*.

See Beran et al. (1995); Garrett and Willinger (1994); Leland et al. (1993); Park and Willinger (2000); Willinger et al. (1995a); Crovella and Bestavros (1997, 1996).

Simple models help to

- Understand origins and effects of long-range dependence and self-similarity
- Understand connections between { SS & LRD } and heavy tails.

Models:

→ Superposition of on/off processes (Taqqu et al. (1997); Mikosch and Stegeman (1999); Mikosch et al. (1999); Stegeman (1998); Willinger et al. (1995b); Heath et al. (1998, 1997); Jelenković and Lazar (1998)),

→ Infinite source Poisson model, sometimes called the $M/G/\infty$ input model (Guerin et al. (1999); Heath et al. (1999); Jelenković and Lazar (1996); Jelenković and Lazar (1999); Mikosch et al. (1999); Resnick and Rootzén (2000); Resnick and van den Berg (2000)).

These models lead to the paradigm: heavy tailed file sizes cause LRD in network traffic.

The infinite node Poisson model:

Infinitely many potential users connected to single server which processes work at constant rate r. At a Poisson time point,
$\rightarrow$ some user begins transmitting work to the server at constant (ugh!) rate (=1).
$\rightarrow$ the active user selects a file size from a heavy tailed distribution.

The good news:

- Fairly flexible and simple.
- The aggregate transmission rate at time t is the number of active users at t.
- The length of the transmission is random and heavy tailed.
- The model offers a simple explanation of LRD being caused by heavy tailed file sizes.
- The model predicts traffic aggregated over users and accumulated over time $[0, T]$ is approximated by either FBM (Gaussian) or stable Lévy motion (heavy tailed).

The less good news:

- The model does not fit data all that well.

(a) Constant transmission rate assumption is clearly wrong.

(b) Not all times of transmissions are Poisson. Some are machine triggered. (Eg CNN website)

(c) There is no hope the simple model can successfully match fine time scale behavior observed below, say, 100 milliseconds which is speculated to be multifractal.

RE (b): Note the invariant: behavior associated with humans acting independently is Poisson; behavior associated with machine behavior is not.

Model Inadequacies

Inadequacy 1. Difficult to identify Poisson points.

EXAMPLE: UCB data; http sessions via modem.

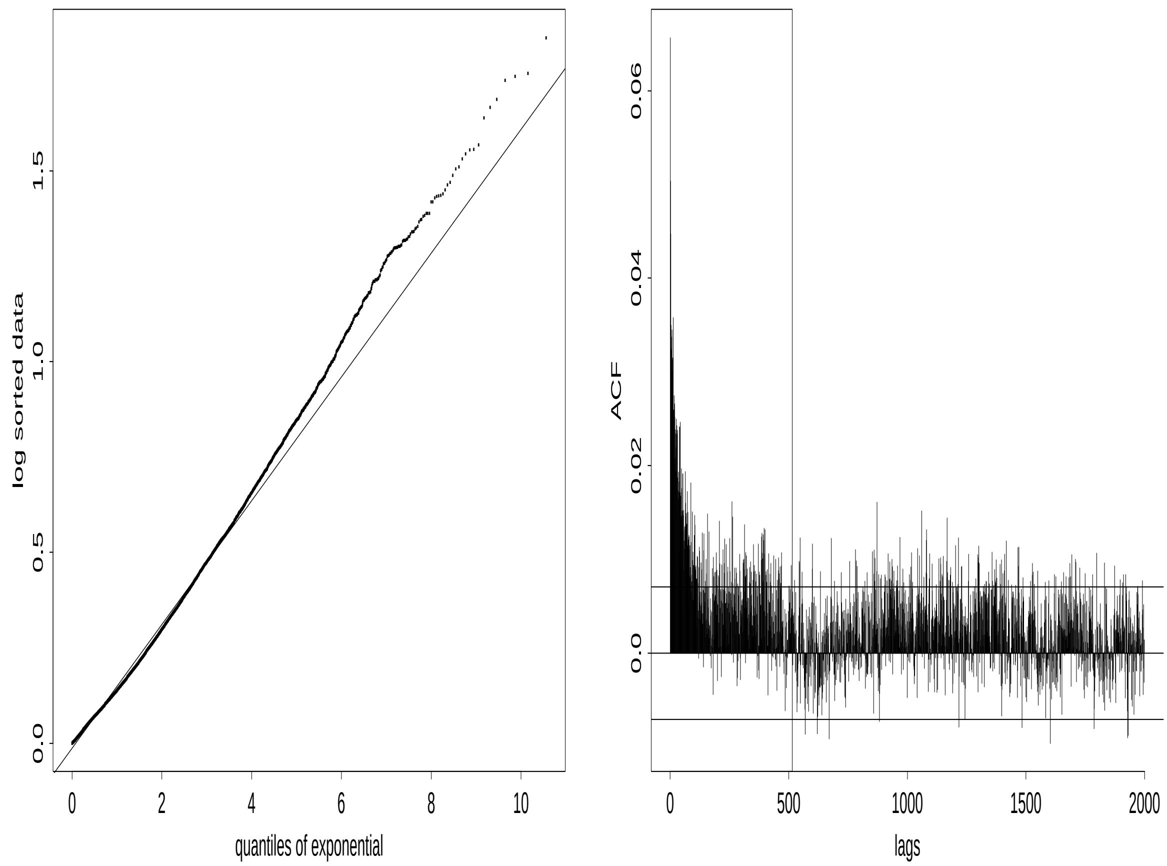

Figure: UCB inter-arrival times of requests: *left*) qqplot against exponential distribution, *rt*) autocorrelation function of interarrival times.

Inadequacy 2. Model $\Rightarrow$ traffic rates approximated by either FBM or stable Lévy motion. In practice, traffic rates not heavy tailed.

POT MLE fitting of Pareto uses McNeil software but estimates threshold sensitive.

Data set	γ (=$1/\alpha$)
simM/G/∞	$-.13 \pm .03$
BUburst 10s	$-.36 \pm .13$
BUburst 1s	$.17 \pm .03$
UCB 10s	$.05 \pm .18$
UCB syn 10s	$-.60 \pm .14$
Munich lo TX	$.09 \pm .04$
Munich lo RX	$-.03 \pm .05$
Munich hi .1s	$.17 \pm .12$
Munich hi .01s	$.10 \pm .03$
Ericsson	$-.47 \pm .09$
Eri syn 1s	$-.31 \pm .12$

Table: Point estimates for traffic rate tail $\pm$ "standard deviation".

Inadequacy 3. Model requires heavy tailed durations to be iid. Often heavy tailed data not iid and dependence models for heavy tailed data notoriously difficult.

Coping strategies:

Q&D method 1: check if sample correlation function

$$\hat{\rho}(h) = \frac{\sum_{i=1}^{n-h}(X_i - \bar{X})(X_{i+h} - \bar{X})}{\sum_{i=1}^{n}(X_i - \bar{X})^2}, \quad h = 1, 2, \ldots,$$

is **close** to identically 0. How to put meaning to phrase **close** to 0? If

(a) If finite variances, Bartlett's formula provides asymptotic normal theory.

(b) If heavy tailed, Davis and Resnick formula provides asymptotic distributions for $\hat{\rho}(h)$.

Q&D method 2: If data heavy tailed, take a function of the data (say the log) to get lighter tail and test. (But this may obscure the importance of large values.)

Q&D method 3: Subset method. Split data into (say) 2 subsets. Plot acf of each half separately. If iid, pics should look same.

EXAMPLE: *Silence*: 1026 times between transmission of packets at a terminal.

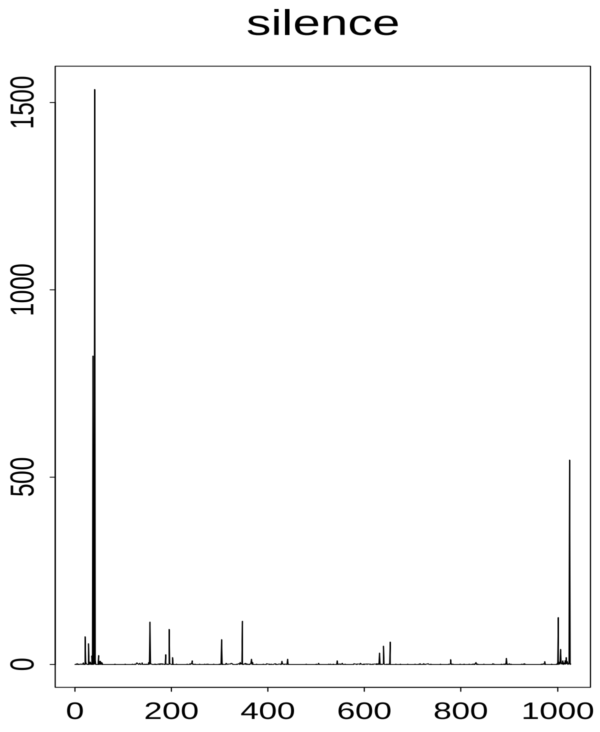

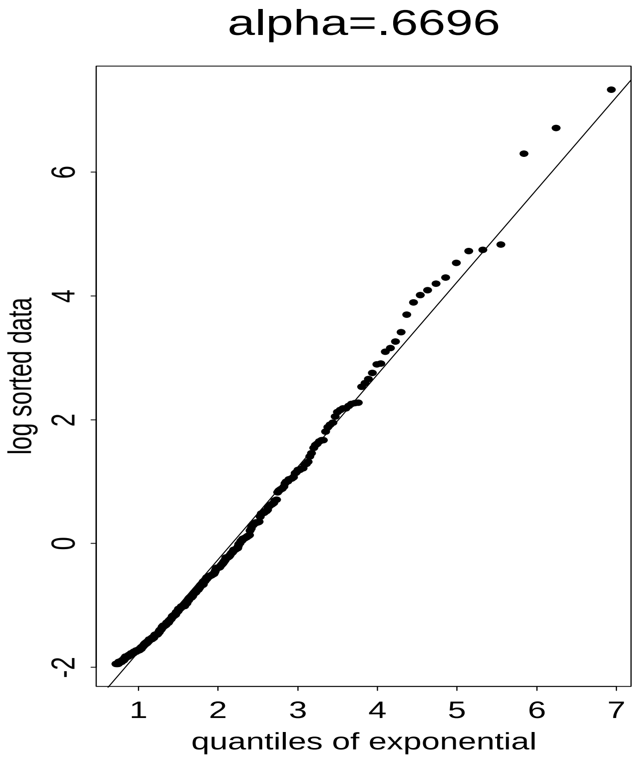

Try to fit black box time series model to the data. Best available techniques suggest AR(9)

$$X_n = \sum_{i=1}^{9} \phi_i X_{n-i} + Z_n,$$
$$n = 9, \ldots, 1029 - 9, \{Z_n\} \text{ iid.}$$

Estimate the coefficients and a goodness of fit test for the AR model is:

Q: Are the residuals iid?

A: Nope.

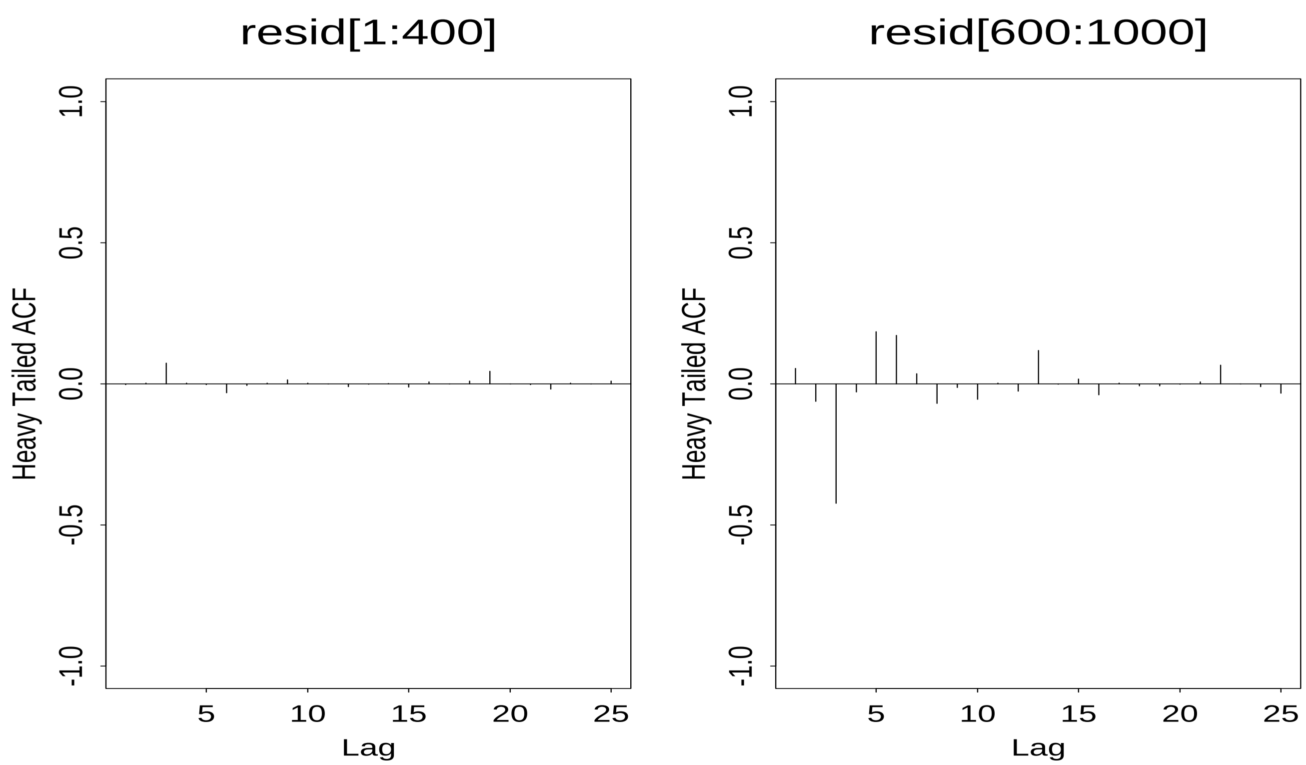

Inadequacy 4. The traffic data is often not **stationary**. For example, there are diurnal cycles.

Q&D Coping Method: Take a slab of the copius data which looks stationary.

- Rule of thumb: don't take more than 4 hours.
- Should we try to model the cycles?

EXAMPLE: Traffic rates resulting from HTTP requests to a large server. 24 hours of data (approx).

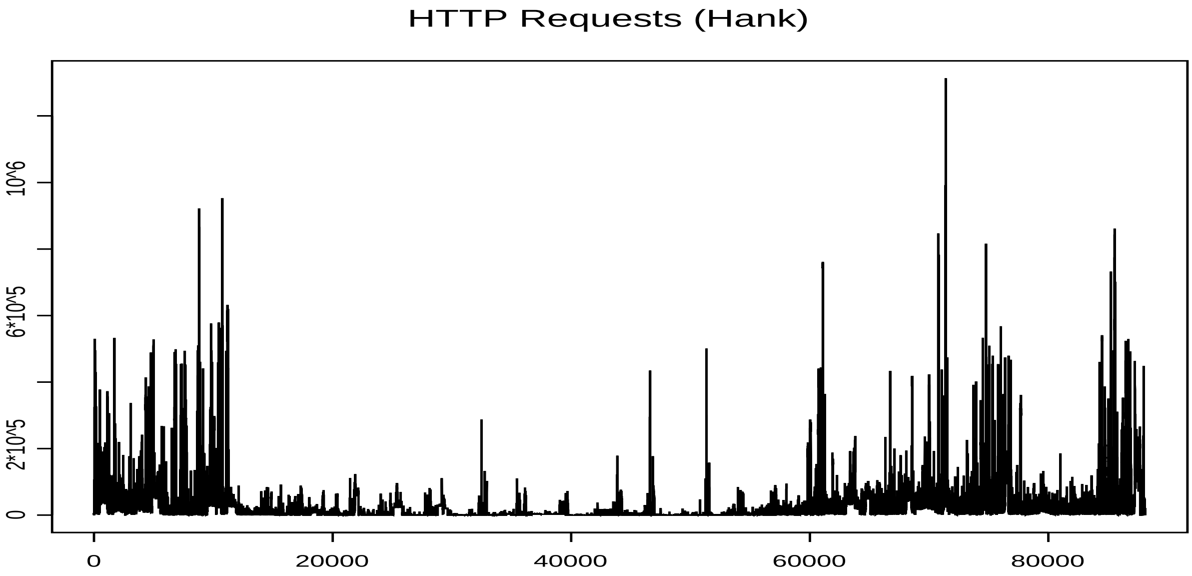

More General Model Specification

1. The k-th transmission begins at Γ_k. $\{\Gamma_k\}$ is a sequence strictly increasing to ∞.

2. The size of the file transmitted is $J_k > 0$.

3. A transmission proceeds according to a transmission schedule $A_k(\cdot)$.

   $A_k(t)$ =the amount of data transmitted in t time units after the kth transmission has begun.

   Properties of $A_k(\cdot)$:

   (i) $A_k(\cdot)$ is a non-decreasing càdlàg process
   (ii) $A_k(t) = 0$, $t \leq 0$,
   (iii) $A_k(t) \uparrow \infty$ as $t \uparrow \infty$.

The process of interest is cumulative traffic over $[0, t]$

$$X(t) = \sum_{k=1}^{\infty} A_k(t - \Gamma_k) \wedge J_k.$$

Small Time Scale Behavior

For small time scales, we need additional assumptions:

4. $\{A_k\}$ are identically distributed with stationary increments.

5. The processes $A_k(\cdot)$ are random multifractals; the multifractal spectrum of $A_k(\cdot)$ is not degenerated to a single point (ie, there is real multifractal behavior.

6. The multifractal spectrum of $A_k(\cdot)$ restricted to any (non-random) interval is non-random.

Multifractal Background

Hölder Exponent: The *Hölder exponent based on exponential growth rate* of the function x at t is defined as

$$h_x(t) := \liminf_{\epsilon \downarrow 0} \frac{\log \sup_{u:|u-t|\le\epsilon} |x(u) - x(t)|}{\log \epsilon}.$$

Rough: In ϵ-nbd of t:

$$|x(u) - x(t)| \approx \epsilon^{h_x(t)}.$$

Properties:

→ $h_{x+y}(t) \geq h_x(t) \wedge h_y(t)$ and equality holds if either

- ◦ $h_x(t) \neq h_y(t)$ OR
- ◦ $x(\cdot)$, $y(\cdot)$ monotone.

→ If x, y are non-decreasing:

$$h_x(t) = \liminf_{\epsilon \downarrow 0} \frac{\log(x(t+\epsilon) - x(t-\epsilon))}{\log \epsilon}.$$

Multifractal Spectrum: *The multifractal spectrum of the function x for the Hölder exponent based on exponential growth rate* is

$$d_x(a) = \dim(\{t > 0 : h_x(t) = a\}), \quad a \in [0, \infty),$$

where $\dim(\Lambda)$ is the Hausdorff dimension of Λ.

Relating multifractality of transmission schedules to cumulative traffic:

Theorem: If the assumptions (1)–(6) hold, then with probability 1, $d_X = d_{A_1}$.

Possible explanation for multifractality in aggregate traffic: intermittancy in individual transmissions caused by blocking and congestion.

Proof. Sample path analysis.

Example: $A_k(t)$ is an increasing Lévy process.

Large Time Scale Behavior

Additional distributional assumptions:

7. $\{\Gamma_k\}$ is homogeneous Poisson process with intensity parameter λ.

8. $\{(A_k, J_k) : k \geq 1\}$ are iid independent of $\{\Gamma_k\}$.

   Note:

   $$\begin{aligned} L_k &= \inf\{t : A_k(t) \geq J_k\} = A_k^{\leftarrow}(J_k) \\ &= \text{length of } k\text{-th transmission} \end{aligned}$$

   and

   $$\bar{F}_L(x) = P[L_1 > x] = P[A_1(x) < J_1].$$

9. Assumption about the joint distribution of $(A_1(\cdot), J_1)$:

   $\exists$ regularly varying function $\sigma \in RV_H$,
   $\exists$ proper random process χ with stationary increments, taking values in $\mathbb{D}[0,\infty)$,
   and for each fixed $\epsilon > 0$,

$$\frac{1}{\bar{F}_J(\sigma(T))} P\left[\frac{J_1}{\sigma(T)} > \epsilon, \frac{A_1(T\cdot)}{\sigma(T)} \in \cdot\right] \to_w \epsilon^{-\alpha_J} P[\chi \in \cdot]$$

   on $\mathbb{D}[0,\infty)$.

10. Technical: For all $\gamma > 0$, assume

$$\lim_{\epsilon \downarrow 0} \limsup_{T\to\infty} \frac{1}{\bar{F}_J(\sigma(T))} P\left[\frac{J_1}{\sigma(T)} \le \epsilon, \frac{L_1}{T} > \gamma\right] = 0.$$

11. Technical: Assume

$$E\left[\chi(1)^{-\alpha_J}\right] < \infty.$$

Consequences:

1. χ is H-self-similar. (Elaborate methods of Lamperti, Vervaat, Durrett & Resnick.)

2. χ càdlàg, ss, non-degenerate implies $H > 0$.

3. χ is non-decreasing, so $H \geq 1$.

4. If $H = 1$, then $\chi(t) = t\chi(1)$ is linear and no interesting structure so exclude.

5. The distribution tail of J_1 is $RV_{-\alpha_J}$:

$$\bar{F}_J(x) = 1 - F_J(x) \sim x^{-\alpha_J} L(x).$$

6. Cumulative traffic on large time scales looks like stable Lévy motion. More precisely:

Theorem. Suppose (1)–(3) and (7)–(11) hold and define

$$Y_T(t) = \frac{X(Tt) - \lambda TtE(J_1)}{b_J(T)}$$

where

$$b_J(T) = \left(\frac{1}{1 - F_J}\right)^{\leftarrow}(T),$$

is the usual quantile function. Then we have

$$Y_T \to_{\text{fidi}} Z_{\alpha_J},$$

where Z_α is mean 0, skewness 1, α-stable Lévy motion.

Sufficient conditions:

(i) J_1 and A_1 are independent.
(ii) J_1 has a tail of index α_J.
(iii) A_1 is itself a proper H-ss process.
(iv) $E[A_1(1)^{-\rho}] < \infty$ for some $\rho > \alpha_J$.

If A_1 is a $\frac{1}{H}$-stable Lévy motion, (iii) & (iv) ok.

How to Prove:

I. The point process

$$M = \sum_{k=1}^{\infty} \epsilon_{(\Gamma_k, A_k, J_k)},$$

is Poisson with mean measure

$$\lambda d\gamma \times P[A_1 \in da, J_1 \in dj]$$

on $(0, \infty) \times \mathbb{D}_{\uparrow} \times (0, \infty)$.

II. $X(T)$ is a function of M. Write

$$X(T) = X_1(T) + X_2(T),$$

where

$X_1(T)$ =accumulation from transmissions
starting and ending before T,
$X_1(T)$ =accumulation from transmissions
starting before T but ending after T.

$X_2(T)$ is negligible.

How to Prove (continued):

III. $X_1(T)$ is a functional of M restricted to a finite region and hence can write

$$X_1(T) =^d \sum_{k=1}^{P(T)} W_k^{(T)}$$

where $\{W_k^{(T)}\}$ iid indep of $P(T)$ which is PO.

IV. Regular variation type condition holds for tail of W_1^T

$$\lim_{T\to\infty} TP\left[W_1^{(T)} > b_J(T)w\right] = w^{-\alpha_J}.$$

V. Point process method to get that the Poissonized sum of W's is asymptotically stable.

*References

J. Beran, R. Sherman, W. Willinger and M. Taqqu (1995): Long range dependence in Variable-bit-rate video. *IEEE Transactions on Communications* 43:1566–1579.

M. Crovella and A. Bestavros (1996): Self-similarity in world wide web traffic: evidence and possible causes. *Performance Evaluation Review* 24:160–169.

M. Crovella and A. Bestavros (1997): Self-similarity in world wide web traffic: evidence and possible causes. *IEEE/ACM Transactions on Networking* 5(6):835–846.

M. Garrett and W. Willinger (1994): Analysis, modeling and generation of self similar VBR video traffic. In *ACM SigComm*. London.

C. Guerin, H. Nyberg, O. Perrin, S. Resnick, H. Rootzen and C. Starica (1999): Empirical testing of the infinite source Poisson data traffic model. Technical Report 1257, School of ORIE, Cornell University, Ithaca NY 14853; available at www.orie.cornell.edu.

D. Heath, S. Resnick and G. Samorodnitsky (1997): Patterns of buffer overflow in a class of queues with long memory in the input stream. *Ann. Appl. Probab.* 7:1021–1057.

D. Heath, S. Resnick and G. Samorodnitsky (1998): Heavy tails and long range dependence in on/off processes and associated fluid models. *Math. Oper. Res.* 23:145–165.

D. Heath, S. Resnick and G. Samorodnitsky (1999): How system performance is affected by the interplay of averages in a fluid queue with long range dependence induced by heavy tails. *Ann. Appl. Probab.* 9:352–375.

P. Jelenković and A. Lazar (1998): Subexponential asymptotics of a Markov-modulated random walk with queueing applications. *J. Appl. Probab.* 35:325–347.

P. Jelenković and A. Lazar (1999): Asymptotic Results for Multiplexing Subexponential On-Off Processes. *Advances in Applied Probability* 31:394–421.

P. R. Jelenković and A. A. Lazar (1996): A network multiplexer with multiple time scale and subexponential arrivals. In *Stochastic Networks*, G. Paul, K. Sigman and D. D. Yao, editors, volume 117 of *Lecture Notes in Statistics*. Springer, New York, pp. 215–235.

W. Leland, M. Taqqu, W. Willinger and D. Wilson (1993): On the self-similar nature of Ethernet traffic. *ACM/SIGCOMM Computer Communications Review* :183–193.

T. Mikosch, S. Resnick, A. Stegeman and H. Rootzén (1999): Is network traffic approximated by stable Lévy motion or fractional brownian motion? Technical report 1247, available at www.orie.cornell.edu/trlist/trlist.html, to appear: Ann. Applied Probability.

T. Mikosch and A. Stegeman (1999): The interplay between heavy tails and rates in self-similar network traffic. Technical Report. University of Groningen, Department of Mathematics, Available via www.cs.rug.nl/ eke/iwi/preprints/.

K. Park and W. Willinger (2000): Self-similar network traffic: An overview. In *Self-Similar Network Traffic and Performance Evaluation*, K. Park and W. Willinger, editors. Wiley-Interscience, New York.

S. Resnick and H. Rootzén (2000): Self-similar communication models and very heavy tails. *Ann. Applied Probability* 10:753–778.

S. Resnick and E. van den Berg (2000): Weak Convergence of high-speed network traffic models. *Journal of Applied Probability* 37:575–597.

A. Stegeman (1998): *Modeling Traffic in High-Speed Networks by ON/OFF Models*. Master's thesis, University of Groningen, Department of Mathematics.

M. Taqqu, W. Willinger and R. Sherman (1997): Proof of a fundamental result in self-similar traffic modeling. *Computer Communications Review* 27:5–23.

W. Willinger, M. Taqqu, M. Leland and D. Wilson (1995a): Self–similarity in high–speed packet traffic: analysis and modelling of ethernet traffic measurements. *Statistical Science* 10:67–85.

W. Willinger, M. Taqqu, M. Leland and D. Wilson (1995b): Self–similarity through high variability: statistical analysis of ethernet

LAN traffic at the source level. *Computer Communications Review* 25:100–113. Proceedings of the ACM/SIGCOMM'95, Cambridge, MA.