Outline

- **Modeling network traffic**

  Background

  From "black-box" to physical-based models

- **Inference for scaling phenomena II: Heavy-tailed distributions**

  CCDF method, Hill's method

- **Explaining self-similar scaling; empirical evidence**

  LANs (Ethernet): Aggregation of On/Off processes

  WANs (Internet): Aggregation of sessions/connections

- **Explaining heavy tails in the networking context**

Scaling Phenomena And Data Network Traffic

Lecture II: Self-Similarity Through High-Variability

Walter Willinger
AT&T Labs – Research
Florham Park, NJ 07932
Email: walter@research.att.com

June 28, 2000

Background
How to Choose Among Different Traffic Models?

- Traditional model selection criteria
  - Familiarity with the model class (e.g., Markov-property)
  - Visual "inspections" (e.g., traces)
  - Statistical model selection criteria (e.g., AIC)
  - Analytic tractability (e.g., performance analysis)
  - Accurate prediction of performance (e.g., queueing)
- Little concern for
  - Model parsimony
  - Physical basis within underlying networking context

Background
Self-Similar Models: Yet Another Traffic Model?

- Modeling alternatives
  - Classical stationary models
    Independent, Markovian, short-range dependent
  - Non-stationary models
    With specific non-stationarity structures
  - SRD approximations of LRD models
  - Self-similar models
    With underlying long-range dependence
  - Non-linear chaotic map models
  - etc.

Background
Model Selection: SRD approximations of LRD models

- Always possible — becomes exercise in data fitting
- Resulting models
  - No predictive power!
  - No insights into actual traffic dynamics
- Leads to highly parameterized models

  The more observations the more parameters
- Interpretation?
- "Black-box" modeling at its best ...

Background
Model Selection: SRD approximations of LRD models

- Model selection criteria from standard time series analysis (Example from J. Beran, 1994, pp. 146–147)
  - Data: 10,000 observations from and LRD process
  - Models: Auto-regressive (AR) processes of order $p > 0$
  - Select "best-fitting" model according to AIC criterion

    Pick first 100 observations: AIC picks AR(2)
    Pick first 500 observations: AIC picks AR(11)
    Pick first 2000 observations: AIC picks AR(18)
    Pick first 4000 observations: AIC picks AR(21)
    Pick first 10000 observations: AIC picks AR(25)

Background
Model Selection: The "Black-Box" Approach

- Operational or "black-box" view of self-similarity
  - Self-similarity (LRD, scaling) as mathematical property
  - Self-similar processes (e.g., FGN) as black box models
  - Black box models as network traffic models
  - *" ... the success of LRD models has a rather pyrrhic flavor; a solution of one mathematical puzzle (i.e., Hurst Effect) has been paid for by the introduction of several physical riddles ..." (V. Klemes, 1974, hydrology-literature)*
  - This operational view has been the main focus of much of the self-similar traffic modeling to date.

Background
Model Selection: (Non-)Stationary Models

- Re: Stationary SRD models
  - SRD is not consistent with measured traffic data
  - Measurements do not support the Markovian assumption
  - It is easy to distinguish between SRD and measured traffic
- Re: Non-stationary models
  - Suggested structure: "Level shifts"
  - Based on non-rigorous analysis of a single data set
  - Statistical tests exist
  - "Level shifts" are not consistent with measured data

From "Black-Box" to Physical-Based Models
Physical-Based Modeling of Self-Similar Traffic

- Self-similarity (LRD, scaling):

  Result of an explicit construction

- Construction:
  - Result of a detailed physical understanding of self-similarity in the network context
  - Heavy reliance on subject-matter expertise

- **Benefit:**

  Physically-based or structural network traffic models that take fully into account the underlying networking context.

From "Black-Box" to Physical-Based Models
Physical Understanding of Self-Similar Scaling

- Central question:
  - **Why is network traffic self-similar?**
- Physical meaning in networking context
  - Account for the hierarchical network structure
  - Capture effects of protocols and network controls
  - Incorporate network-host interactions
- Parsimonious models for network traffic
  - A by-product of the physical understanding of traffic

Inference for Scaling Phenomena II
Heavy-Tailed Distributions: Properties

- Key facts:
  - If $1 \leq \alpha < 2$, F has finite mean but infinite variance (Noah Effect or infinite variance syndrome)
  - Heavy-tailed (with infinite variance) means high variability
  - Plotting $P[X > x]$ vs. x on log-log scale shows linear relationship for large x, with slope $-\alpha$.
- Examples: Pareto, Pareto-stable, Zipf, etc.

Inference for Scaling Phenomena II
Heavy-Tailed Distributions: Definition

- Definition: Let X be a random variable with distribution function F (e.g., call holding time, packet interarrival time, burst length, traffic rate). F is said to be heavy-tailed if for some $\alpha > 0$,

$$1 - F(x) = P[X > x] \sim cx^{-\alpha}, \quad \text{as } x \to \infty,$$

where c is a finite positive constant, independent of x.

- Terminology: Heavy-tailed distributions are also called power-law or scale-invariant distributions.

Inference for Scaling Phenomena II
Inference for Heavy Tails: The Hill Method

- Methodology:
  - Select k upper order statistics $X_{n-k+i,n}, 1 \leq i \leq k$
  - Hill estimator

$$H_{n,k} = (\frac{1}{n}\sum_{i=1}^{k}\log(\frac{X_{n-k+i,n}}{X_{n-k,n}}))^{-1}$$

  - Plot $H_{n,k}$ vs. k ("dynamic" Hill plot)

Inference for Scaling Phenomena II
Inference for Heavy Tails: The CCDF Method

- Methodology:
  - Calculate $1 - \hat{F}_n(x)$, the empirical complementary cummulative distribution function (CCDF) of F.
  - For $1 \leq i \leq n$, plot $1 - \hat{F}_n(X_i)$ vs. X_i on log-log scale
  - Look for "linear asymptotic regime"
  - "Naive" estimator = slope of linear regime
- Properties of the "naive" or CCDF method:
  - Eyeballing method, very useful for large samples
  - No known distributional results for "naive" estimator

Explaining LAN Traffic Self-Similarity

Intuitive Explanation

- Aggregation of many On/Off processes

  The superposition of many i.i.d. On/Off sources becomes white noise, provided the durations of the On- and/or Off-periods have finite variance.

  The superposition of many i.i.d. On/Off sources is self-similar, provided the durations of the On- and/or Off-periods are heavy-tailed with infinite variance.

Inference for Scaling Phenomena II

Inference for Heavy Tails: The Hill Method

- Properties of the Hill-method:
  - More sophisticated eyeballing technique
  - Useful in practice to control bias and variability
  - Heuristics: choice of cutoff-value k
  - More enhanced versions of Hill's method are known
  - $H_{n,k}$ is consistent and asymptotically normal, but only under conditions that are unverifiable in practice
  - Reference: S. Resnick and C. Starica, 1995-1997.

Explaining LAN Traffic Self-Similarity
Mathematical Explanation (cont.)

- M i.i.d. host-host pairs:
  With associated traffic rate processes
  $W_m = (W_m(t) : t \geq 0); m = 1, 2, \ldots, M$
- Aggregate traffic $W^\star(M) = (W^\star(M, t) : t \geq 0)$:

$$W^\star(M, t) = \int_0^t (\sum_1^M W_m(u))du$$

  where $W^\star(M, t)$ denotes the total # of packets/bytes generated by all M host-host pairs in the time interval $[0, t]$

Explaining LAN Traffic Self-Similarity
Mathematical Explanation

- Single source traffic (host-host pair):
  On/Off or bursty/idle or active/inactive with i.i.d. On-periods (distribution F_1) and i.i.d. Off-periods (distribution F_2), where for $i = 1, 2$,

$$1 - F_i(x) \sim x^{-\alpha_i} L_i(x), \quad \text{as } x \to \infty.$$

- Single source traffic rate process:
  $W = (W(t) : t \geq 0)$, with $W(t) = 1/0$ depending on whether the source is On/Off at time t

Explaining LAN Traffic Self-Similarity
Mathematical Explanation (cont.)

- Remarks:
  - Based on a construction originally due to Mandelbrot
  - Result can be generalized to heterogeneous host-host pairs
  - Taking $M \to \infty$ and then $T \to \infty$ is crucial ...
  - Proof: Laplace-transforms, Tauberian-type theorems
  - Weak convergence can be established
  - Theorem provides an answer to the question:
    When is FGN a good model for network traffic?

Explaining LAN Traffic Self-Similarity
Mathematical Explanation (cont.)

- Theorem:
  In the limit as $M \to \infty$ and then $T \to \infty$, we have

$$\frac{\Delta W^\star(M, T\cdot) - TM\frac{\mu_1}{\mu_1+\mu_2}}{T^H L^{1/2}(T) M^{1/2}} \to \sigma X_H(\cdot).$$

  Here, X_H denotes fractional Gaussian noise; $\to$ means convergence of the finite-dimensional distributions; the parameters σ and the slowly varying function $L(\cdot)$ can be expressed in terms of the input data (i.e., F_1 and F_2); and

$$H = \frac{3 - \min(\alpha_1, \alpha_2)}{2}$$

Explaining LAN Traffic Self-Similarity
Empirical Validation: Some Numbers

- 1 hour of Ethernet LAN traffic measurements

  1,404,444 packets

- Filtered by source address
  - 105 (out of 121) sources sent or received packets
  - Example: Source 10 sent 26,330 packets (about 1.8%)
- Filtered by host-host or source-destination pairs
  - 748 (out of 11,025) active source-destination pairs
  - Example: For source 10, there were 25 active source-destination pairs, with 5 of them generating about 70% of all the traffic.

Explaining LAN Traffic Self-Similarity
Heavy-Tailed On/Off Processes: Empirical Evidence

- Empirical validation of the On/Off nature
  - Using the header information, extract from the aggregate Ethernet LAN traffic stream individual host-host or source-destination traces
  - Check each individual trace for consistency with On/Off behavior, i.e., extract from each trace appropriate On/Off periods and check for robustness properties of the resulting distributions
  - Check the On/Off periods of each trace for consistency with heavy tails with infinite variance

Explaining LAN Traffic Self-Similarity
By-product: Fast generation of self-similar traffic

- Available hardware:

  Massively parallel machine (e.g., MasPar MP-1216 with 16,384 processors

- Method: Mandelbrot's construction
  - Let each processor generate traffic according to an On/Off-source with infinite variance On- and/or Off-periods
  - Aggregate many sources
  - Resulting algorithm is linear in the number of observations

Explaining LAN Traffic Self-Similarity
The On/Off nature: Scientific inference

- For each source-destination trace, extract On/Off periods (e.g., using a thresholding parameter)
- Plot $P[\text{On} > x]$ and $P[\text{Off} > x]$ vs. x, on log-log scale
- Straight-line behavior suggests heavy tails, i.e., $P[\text{On} > x] \sim x^{-\alpha_1}$, or $P[\text{Off} > x] \sim x^{-\alpha_2}$, for large x.
- Obtain estimate $\hat{\alpha}_i$ of the index α_i by applying a number of different inference methods (e.g., "eyeballing", qq-plot, Hill's estimate)
- Check robustness properties of $\hat{\alpha}_i$ under different choices of the thresholding parameter

Explaining WAN Traffic Self-Similarity
Intuitive Explanation for Large-Time Scaling

- Aggregation of many sessions/connections

  The aggregate traffic generated by Poisson arrivals of sessions (e.g., TELNET, FTP, WWW) with finite variance session duration (in seconds) or finite variance session size (in bytes) exhibits SRD.

  The aggregate traffic generated by Poisson arrivals of sessions (e.g., TELNET, FTP, WWW) with infinite variance session duration (in seconds) or infinite variance session size (in bytes) exhibits LRD.

Explaining LAN Traffic Self-Similarity
References

- M.S. Taqqu, W. Willinger and R. Sherman, *Computer Communication Review 27*, pp. 5–23, 1997.
- B.B. Mandelbrot, *International Economic Review 10*, pp. 82–113, 1969.
- W. Willinger, M.S. Taqqu, R. Sherman and D.V. Wilson, *IEEE/ACM Transactions on Networking 5*, pp. 71–86, 1997.
- W. Willinger, V. Paxson and M.S. Taqqu, in: *A Practical Guide to Heavy Tails: Statistical Techniques for Analyzing Heavy Tailed Distributions*, R. Adler, R. Feldman and M.S. Taqqu (eds.), Birkhauser, Boston, 1998.

Explaining WAN Traffic Self-Similarity
Mathematical Explanation for Large-Time Scaling

- Cox's construction (cont.)
  - If $W^\star = (W^\star(k) : k \geq 1)$ denotes the total # of bytes/packets/workload input into the system up to time $k, k \geq 1$, then the autocorrelation function of $(\Delta W^\star(k) : k \geq 1)$ satisfies

$$r(k) \sim ck^{1-\alpha}, \quad \text{as } k \to \infty$$

    for some finite positive constant c. Thus, $(\Delta W^\star(k) : k \geq 1)$ is (asymptotically second-order) self-similar with self-similarity parameter $H = (3 - \alpha)/2$.

Explaining WAN Traffic Self-Similarity
Mathematical Explanation for Large-Time Scaling

- Cox's construction: (Cox, 1984)
  - Also known as birth-immigration process or $M/G/\infty$-model
  - Poisson session arrivals
  - Session duration (in seconds) or session sizes (in bytes) are heavy-tailed with infinite variance (i.e., $1 \leq \alpha < 2$)
  - "Within-session" dynamics: constant ("fluid") rate
- Kurtz's construction: (Kurtz, 1996)
  - Includes Cox's construction as a special case
  - Allows for more complex "within-session" dynamics

Explaining WAN Traffic Self-Similarity
More on WAN Traffic at the Session/Connection Level

- For www-session related quantities:
  - Extraction of HTTP connection-related information from WAN packet-level traces is feasible
  - HTTP-request durations/sizes: consistent with infinite variance syndrome
  - HTTP-request arrivals: complex structure
  - Impact of www on traffic dynamics at different layers within TCP/IP protocol hierarchy

Explaining WAN Traffic Self-Similarity
WAN traffic at the Session/Connection Level

- For telnet and ftp sessions:
  - Extract session-level information from WAN traces
  - Session arrivals are consistent with Poisson
  - Session durations/sizes are consistent with infinite variance
- For www-sessions:
  - Extraction of www session-level information from WAN packet-level traces is not possible
  - Use plausible substitute: Modem calls to ISP
  - Session arrivals: consistent with Poisson arrivals
  - Session durations/sizes: consistent with infinite variance

Explaining WAN Traffic Self-Similarity
By-product: Fast generation of self-similar traffic

- Available hardware:

  Massively parallel machine (e.g., MasPar MP-1216 with 16,384 processors

- Method: Cox's construction:
  - Each active processor corresponds to an active session and generates traffic (at a constant rate) during its lifetime which, in turn, is heavy-tailed with infinite variance.
  - Aggregate over all active processors
  - Resulting algorithm is linear in the number of observations

Explaining WAN Traffic Self-Similarity
References

- V. Paxson and S. Floyd, *IEEE/ACM Transactions on Networking 3*, pp. 226-244, 1995.
- M.E. Crovella and A. Bestavros, *IEEE/ACM Transactions on Networking 5*, pp. 835–846, 1997.
- A. Feldmann, A.C. Gilbert, W. Willinger and T.G. Kurtz, *Computer Communication Review 28*, pp. 5–29, 1998.
- W. Willinger, V. Paxson and M.S. Taqqu, 1998.

THE SELF-SIMILAR NATURE OF NETWORK TRAFFIC
THE BOTTOM LINE

- ABOUT SELF-SIMILARITY
  - Additive property
  - Mainly caused by user/session characteristics
  - Has little to do with the network
  - Is likely here to stay because ...
- HEAVY TAILS ARE NOT LIKELY TO DISAPPEAR

  CPU time of UNIX processes; file sizes on file servers; document sizes on web servers; TELNET packet interarrivals; Web session sizes or durations; TCP, FTP and HTTP connection sizes/durations; etc.