Intelligent control of data networks: effect of network traffic

Ljiljana Trajković

Communication Networks Laboratory
http://www.ensc.sfu.ca/research/cnl
School of Engineering Science
Simon Fraser University

communication networks laboratory

Who we are: the university

- Simon Fraser University
- Located in Vancouver, British Columbia
- School of Engineering Science
- SFU: 18,000 students
- Faculty of Applied Science
- ENSC: ~400 undergraduate and ~60 graduate students, with ~24 faculty
- Communication Networks Laboratory

Who we are: our laboratory

- Postdoctoral fellows:
  - Luc A. Andriantiatsaholiniaina, Judy Liu
  - Slobodan Petrovic, Fei Xue
- Ph. D. student:
  - Inas Khalifa
- M. A. Sc. students (engineering):
  - Milan Nikolic, Grace Hui, Kenny Shao
  - Nazy Alborz, Maryam Keyvani, Velibor Markovski, Michael Jiang

Who we are: our laboratory

- M. A. Sc. students (computing science):
  - Johnson Chen, Leo Chen, Tony Feng, James Song
- M. Eng. students:
  - Jeffery Lau, Barry Moss, Thomas Pang, Duncan Sharp
  - Zelimir Lucic, Ricky Ng, Hubert Pun, Danny Yip
- B. A. Sc. students:
  - Amir Jodari, Eric Keung, Savio lau
  - Bruce Chan, David Ciampi, Arash Haydari-Khabbaz, Robert Trost, Randeep Ghakal

http://www.ensc.sfu.ca/research/cnl

Students involved in the traffic project

- Amir Jodari (undergraduate)
- Bruce Chan (B. A. Sc.)
- Želimir Lučić (M. Eng.)
- Nazy Alborz (M.A.Sc.)
- Velibor Markovski (M. A. Sc.)
- Giorgio Bombelli (directed studies)
- Dr. Fei Xue (postdoctoral fellow)

Road map

- Modeling of Internet traffic
- Trace-driven simulations
- Simulation scenarios and results
- Analytical tools for traffic characterization:
  - wavelet analysis
- Current and future work

Introduction: traffic patterns

- Complex traffic patterns from multiplexed data, voice, image, and video patterns.
- Traditional traffic models fail to capture essential characteristics of these traffic patterns.
- Presence of the traffic invariants has been detected in traffic traces.
- Traffic exhibits long-range dependent (self-similar, fractal, chaotic) behavior.

Traffic collection

- Active measurements: sending probes and measuring network response.
- Passive measurements: non-intrusive, simply observing packet loss and delay.
- Tools:
  - traceroute
  - tracedump
- Traffic measurement cannot be replaced with simple traffic modeling.

Self-similar traffic patterns

Genuine MPEG traffic trace

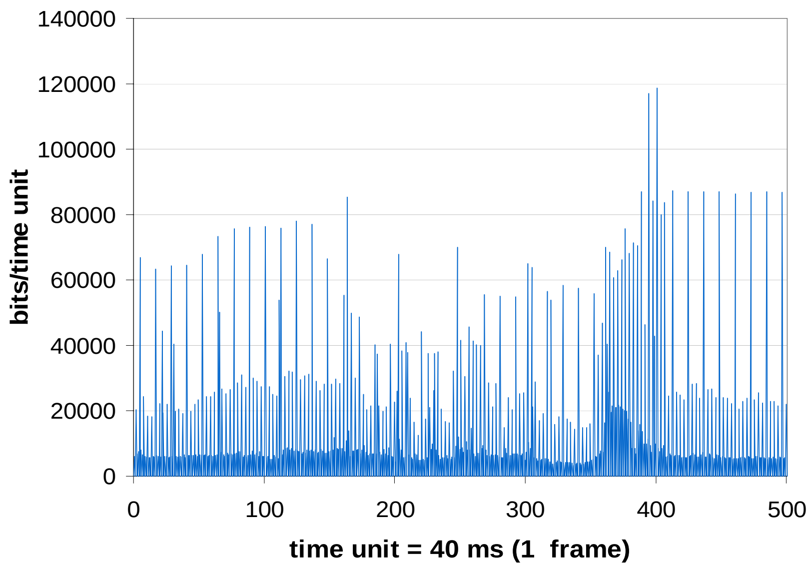

Poisson model

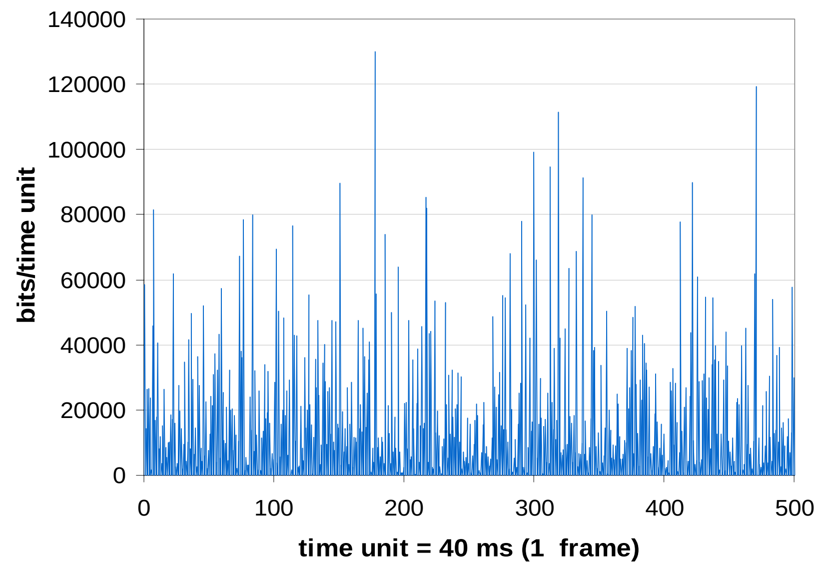

The two traces have identical mean.

Influence of time-scales

- Genuine MPEG traffic trace:

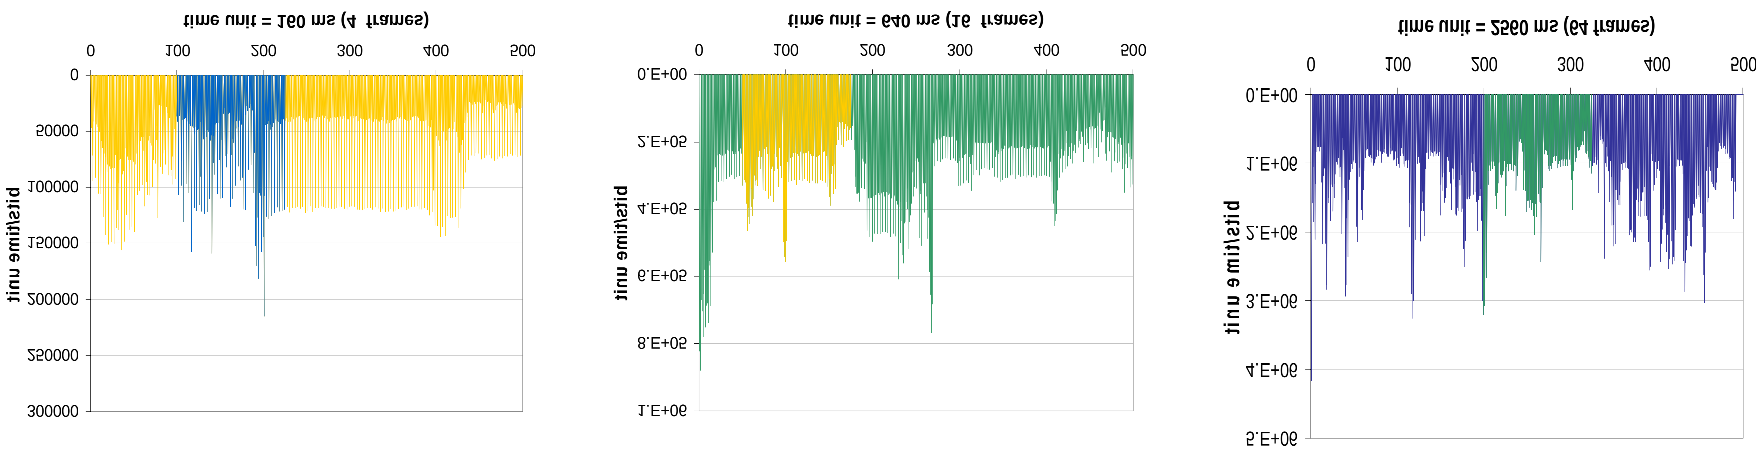

W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Trans. Networking,* vol. 2, pp. 1 – 15, 1994.

Influence of time-scales

- Synthetically generated Poisson model:

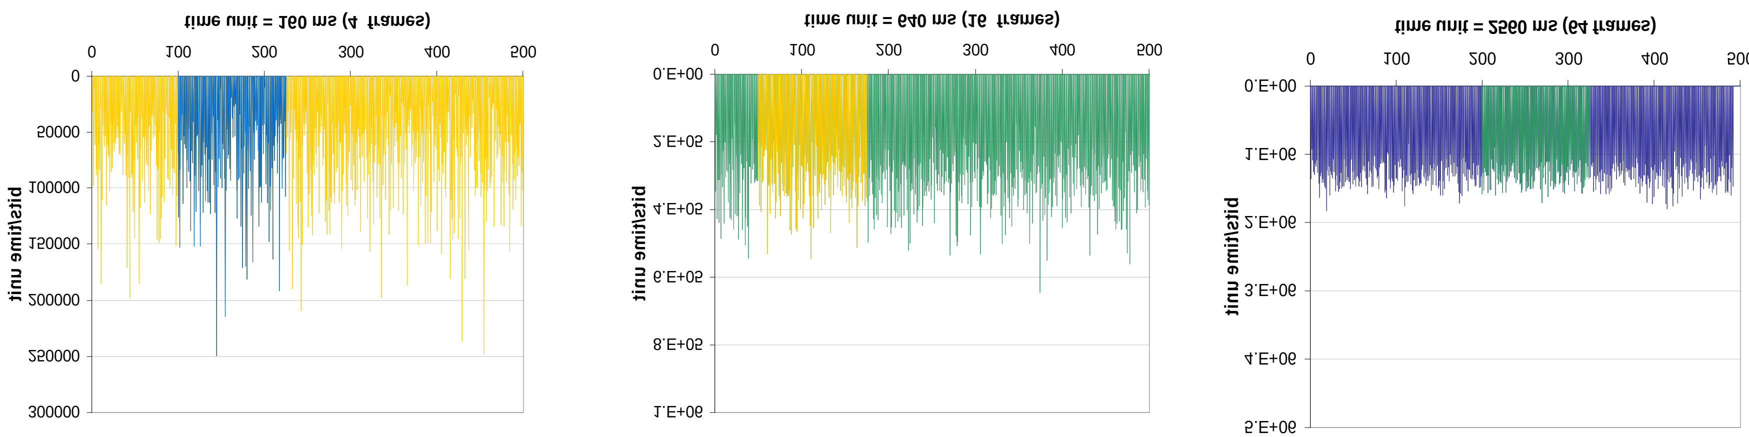

W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Trans. Networking,* vol. 2, pp. 1 – 15, 1994.

Heavy-tailed distribution of packet loss

- Pareto distribution can be used to model the distribution of packet loss episode lengths:

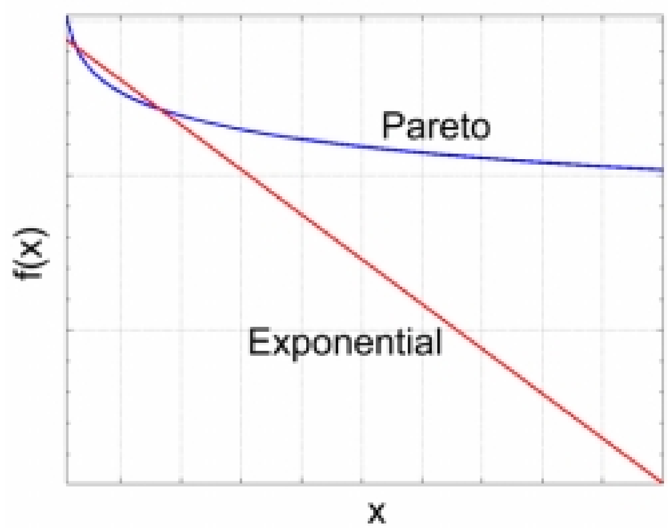

LRD and Hurst parameter

- Self-similar process:
  - long-range dependent (LRD)
  - fractal
  - with heavy-tailed distributions
  - Hurst parameter: $0.5 < H < 1$
- Network traffic often exhibits self-similarity.

LRD and Hurst parameter

Long-tailed distributions:

- the variance of a sample mean $X^{(m)}$ decreases (slowly) as:

$$\mathrm{var}\left(X^{(m)}\right) = \sigma^2 m^{-\beta}$$

- the power spectral density of a self-similar process is:

$$\Gamma(\nu) = c \cdot |\nu|^{-\alpha}$$

Road map

- Modeling of Internet traffic
- Trace-driven simulations
- Simulation scenarios and results
- Wavelet analysis of traffic and packet loss
- Current and future work

Introduction: simulations

- Computer simulation and empirical techniques play an important role in understanding networks' behavior.
- We use genuine traffic traces to simulate loss in packet networks such as Internet.
- Simulation results indicate that underlying transport protocols and time-scales are essential for understanding loss behavior in packet networks.

Traffic patterns: "Star Wars" trace

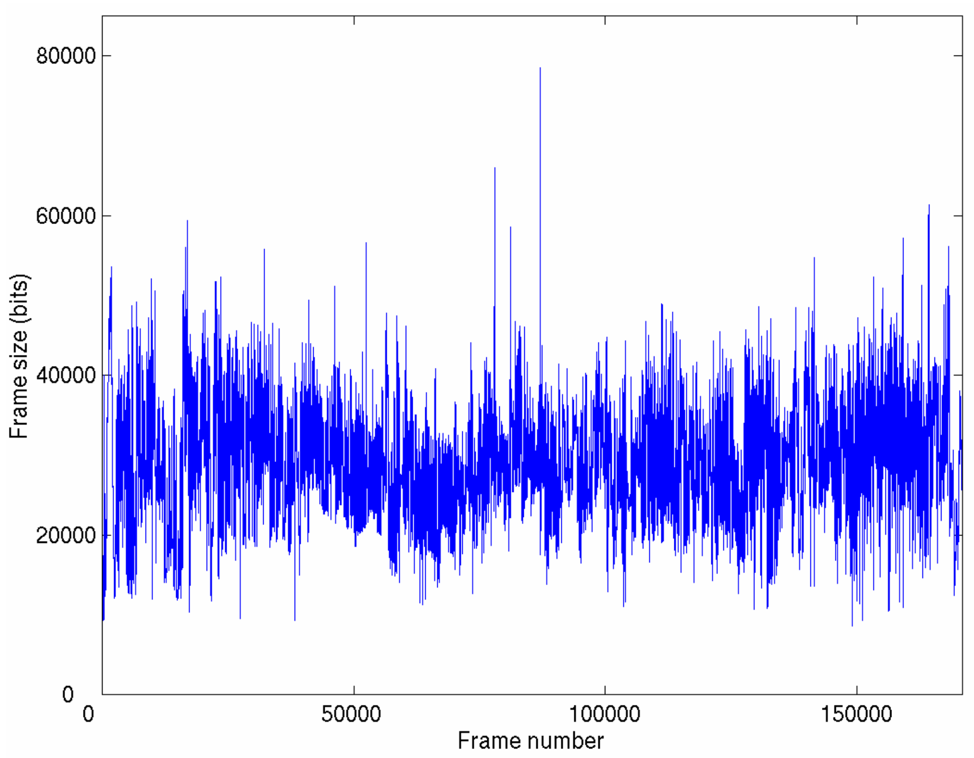

- 170,000 frames (2 hours)
- 24 frames a second
- each source starts at a random point within the trace
- at the end, the trace wraps around to its beginning

Video traces

Traces last 30 minutes.

- Two types of traces:
  - movies: action scenes
  - still videos: parking camera

Sources:

http://nero.informatik.uni-wuerzburg.de/MPEG

http://www-tkn.ee.tu

http://www.berlin.de/research/trace/trace.html

Requirements for data transmission

For multimedia applications:

- packet loss
- packet delay
- delay jitter

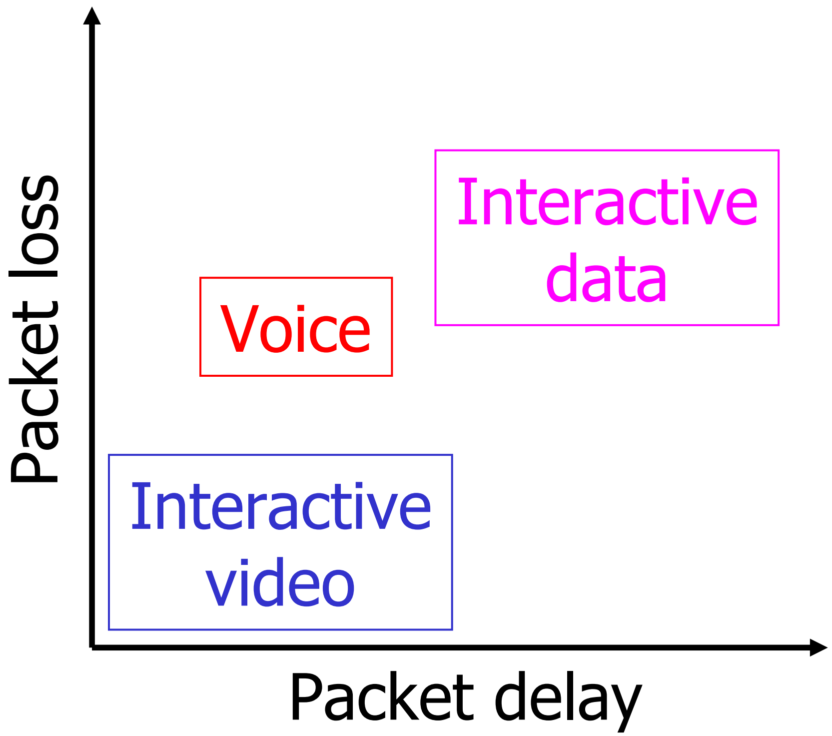

Quality of Service parameters

- Packet loss occurs in the routers, on the links, at the end hosts:
  - packet loss probability
  - packet loss behavior
- Packet delay:
  - real-time applications
- Packet delay jitter (cell delay variation):
  - video on demand applications

Road map

- Modeling of Internet traffic
- Trace-driven simulations
- Simulation scenarios and results
- Wavelet analysis of traffic and packet loss
- Current and future work

ns-2 network simulator

- Collaborative project among USC, Xerox PARC, LBL, and UCB (http://www.isi.edu/nsnam/ns)
- Discrete event network simulator
- Open source code
- Provides support for various:
  - network protocols
  - topologies
  - traffic generators
  - queue management and packet scheduling techniques

OPNET simulator

- Commercial tool (http://www.opnet.com)
- Discrete event network simulator
- Built-in models and contributed models
- Provides support for various:
  - network elements
  - topologies
  - traffic generators

- OPNET university program: free 6 month license

Simple simulation scenario

- *n* video sources, one router, and one sink (destination)
- DropTail queue with buffer size set according to delay requirements
- Trace-driven simulation using genuine video traffic traces (*Star Wars* and *Talk show*)
- Three subscenarios:
  - all sources use User Datagram Protocol
  - all sources use Transmission Control Protocol
  - mixed UDP/TCP traffic

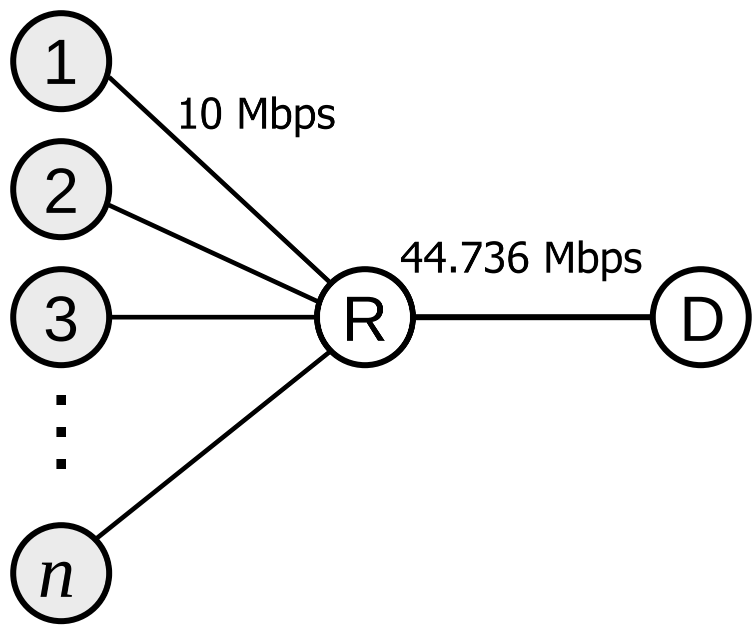

Buffer size: 46 or 200 packets
Packet size: 552 bytes

Complex simulation scenario

- Four transit routers and 200 end hosts
- Mix of network traffic (Web, FTP, and video)

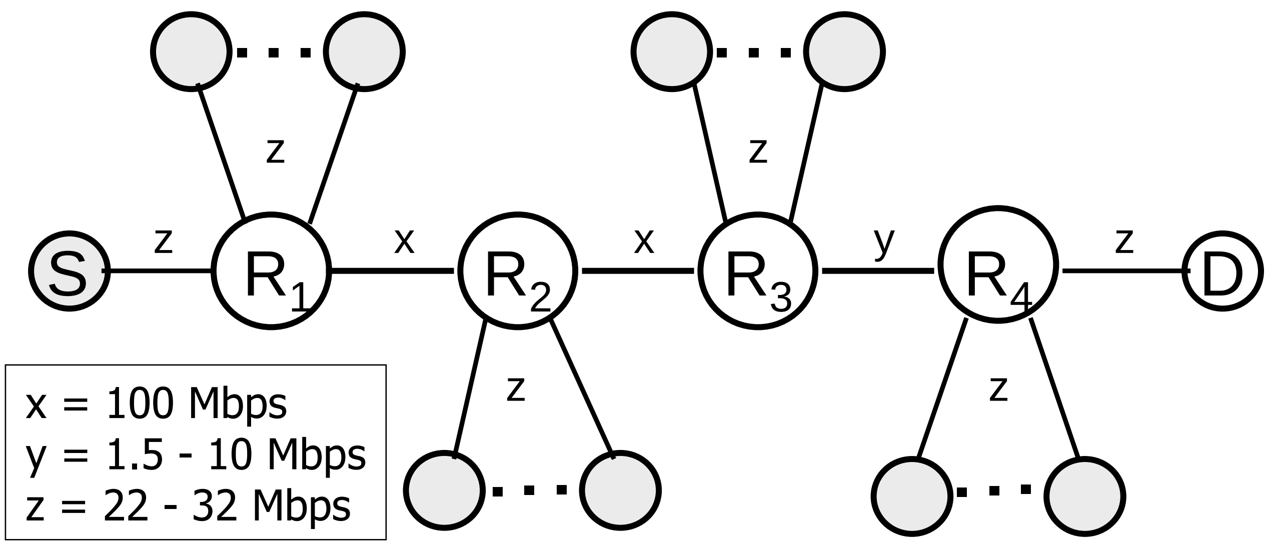

Packet loss rates: mixed UDP and TCP sources

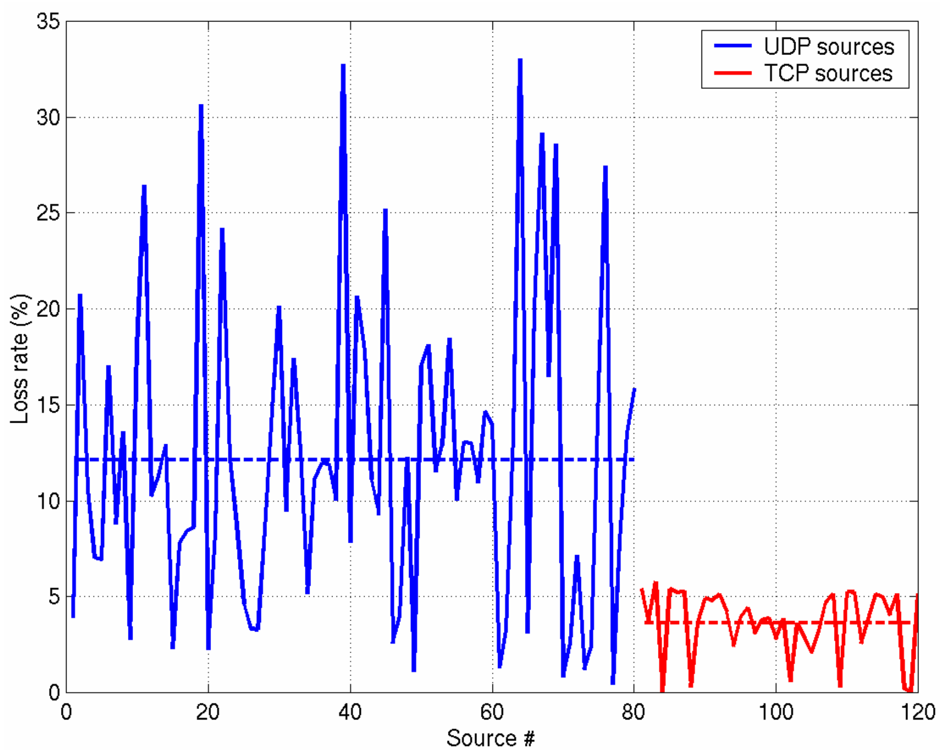

- Packet loss rate for the UDP sources is much larger than the packet loss rate for the TCP sources

Simulation run with 80 UDP and 40 TCP sources, and buffer size of 50 Kbytes (9.2 msec).

MPEG traces

Trace	Mean bit rate (Mbps)	Hurst parameter
Silence of the Lambs	0.18	0.89
Terminator 2	0.27	0.89
MTV	0.49	0.89
Simpsons	0.46	0.89
Talk show 1	0.36	0.89
Jurassic park 1	0.33	0.88
Mr. Bean	0.44	0.85
News	0.38	0.79
Star Wars	0.36	0.74
Talk show 2	0.49	0.73

Simulation results: packet loss

Trace-driven simulations with “Star Wars” trace indicate:

- Increased link utilization causes lengthier loss episodes.
- Periods of lower congestion are characterized with more frequent loss episodes of length one (single packet loss episodes) and with wider loss episode distances.
- Similar behavior was observed with other traffic traces: “Talk show” from the University of Wűrzburg archive.

Simulation results: TCP

TCP transfers:

- Lower packet loss rates than in the case of UDP transfers due to the congestion control mechanisms in TCP transfers.
- Short packet loss episodes: loss episodes of length one contribute over 90% of the overall loss.

Active queuing mechanisms

- FIFO/DropTail
- Random Early Drop (RED)
- Fair Queuing (FQ)
- Stochastic Fair Queuing (SFQ)
- Deficit Round Robin (DRR)

FIFO/DropTail queuing scheme

- Focus of our studies:
  - aggregate packet loss
  - per-flow packet loss
  - packet delay

Aggregate packet loss: buffer size

- Contribution of loss episodes of various lengths (100 sources)

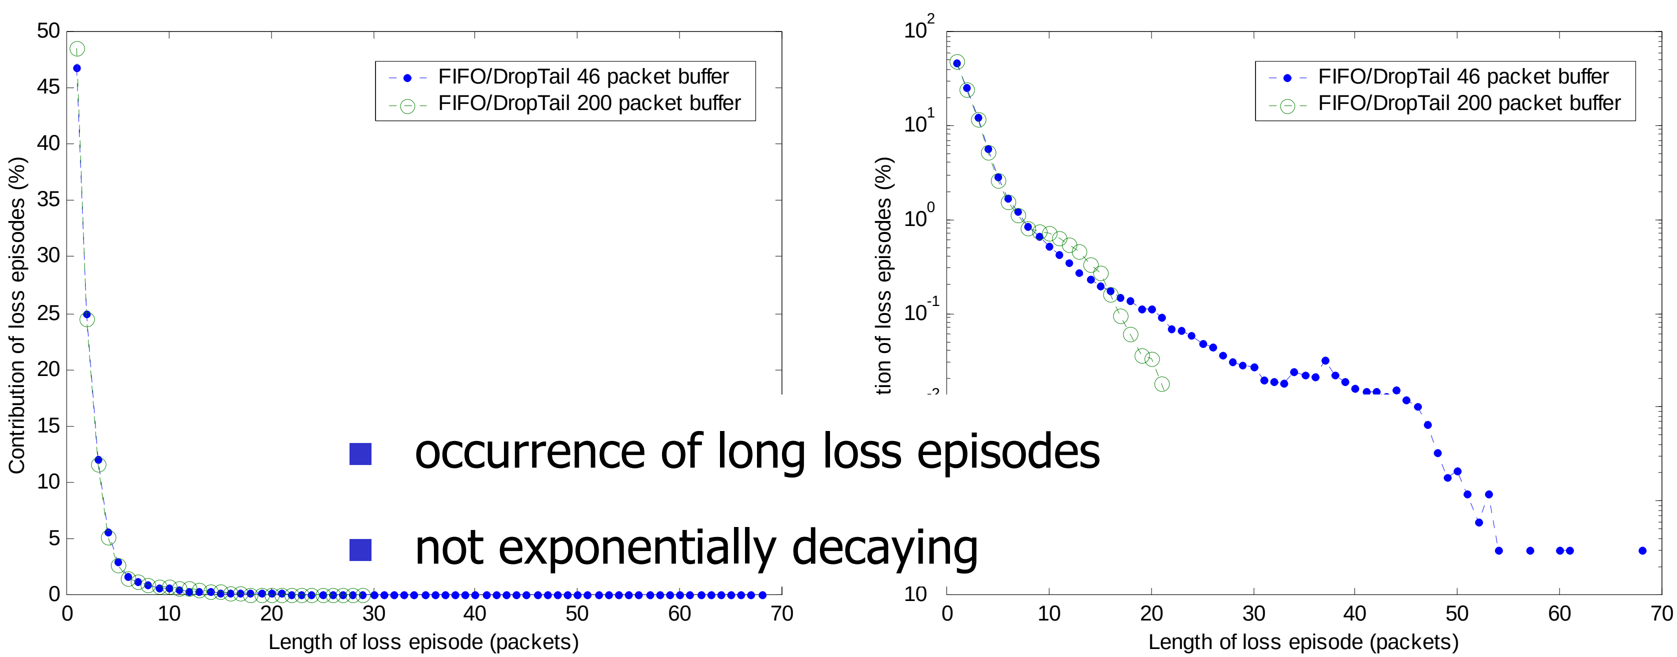

- occurrence of long loss episodes
- not exponentially decaying

Aggregate packet loss: traffic load

- Contribution of loss episodes (46 packet buffer)
- 40 to 100 traffic sources (33% to 82% traffic load)

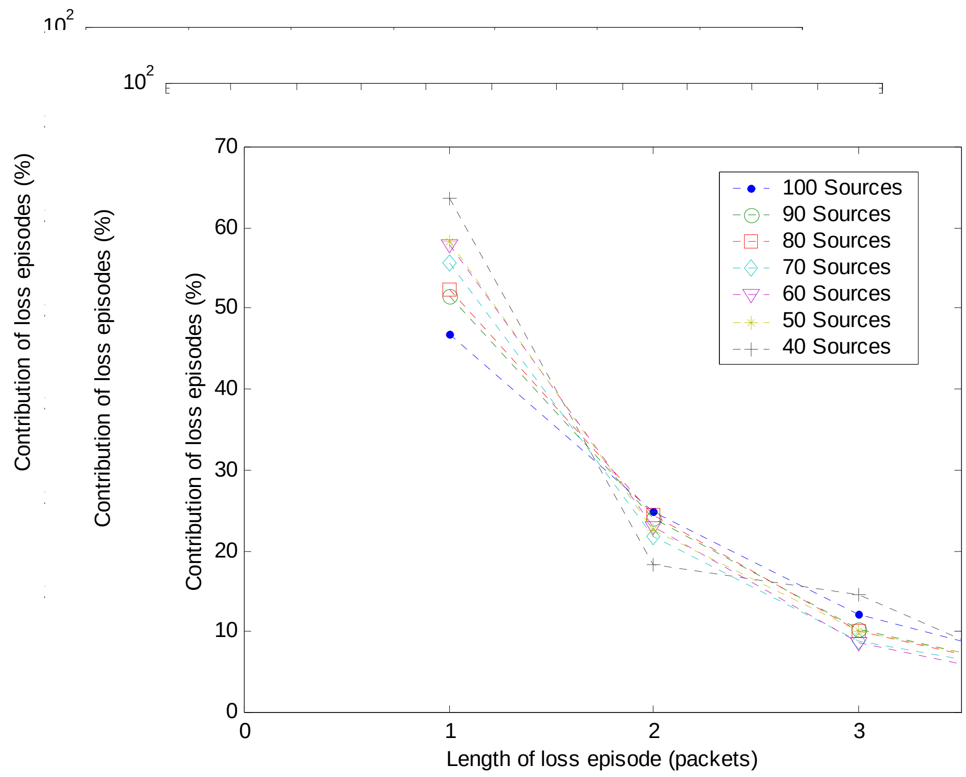

Per-flow packet loss: buffer size

- Contribution of loss episodes of various lengths, averaged over all flows (100 sources)

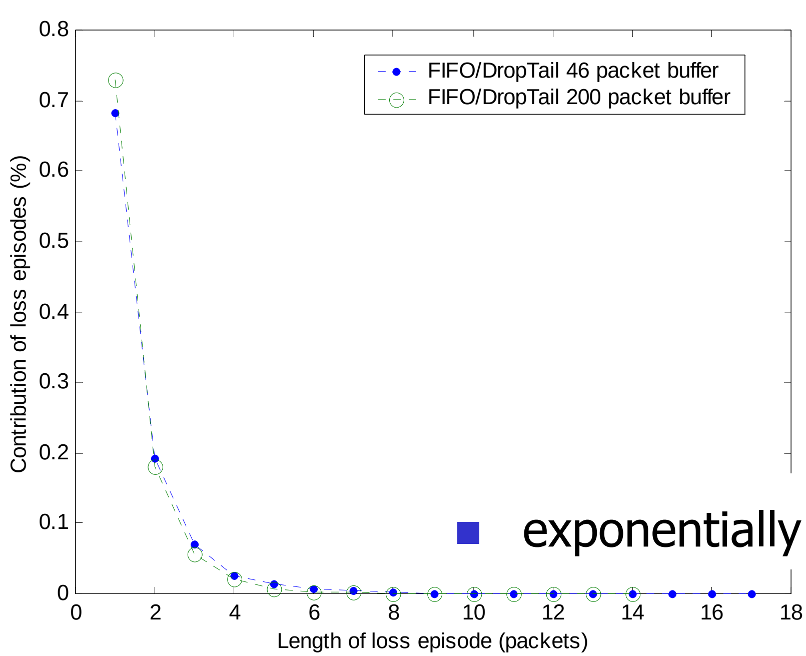

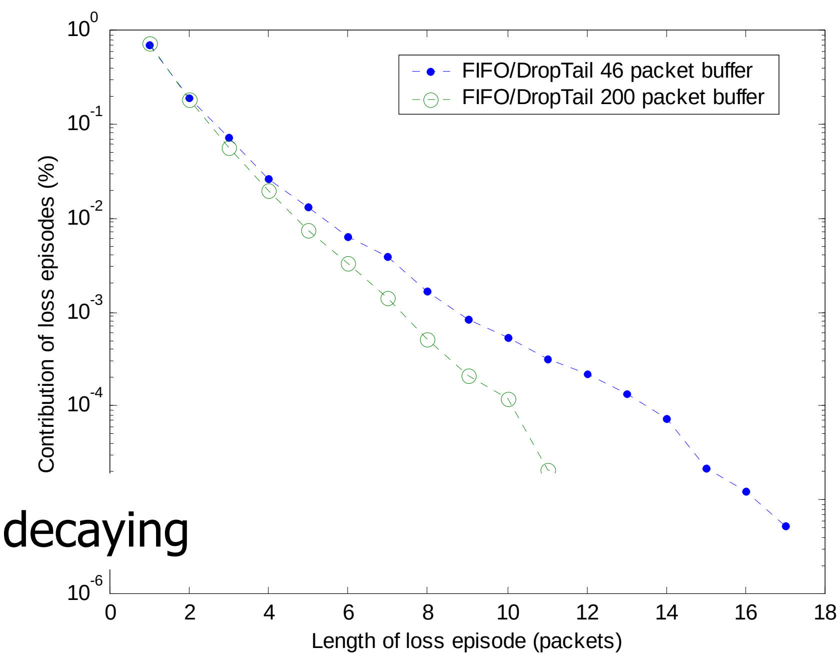

- exponentially decaying

Simulation results: other schemes

- Random Early Drop (RED)
- Fair Queuing (FQ)
- Stochastic Fair Queuing (SFQ)
- Deficit Round Robin (DRR)

Random Early Drop (RED)

- Two thresholds: min_{th} and max_{th}
- Monitors buffer size, drops packet as congestion builds:
  - between min_{th} and max_{th}: drop incoming packets
  - above max_{th}: drop packets from inside the buffer
- Congestion avoidance mechanism
- Originally designed for TCP [Floyd and Jacobson, 1993]

RED simulation results

- 100 sources, 10 MPEG traces, buffer sizes of 46 and 200 packets
- Comparison of RED to FIFO/DropTail in terms of:
  - per-flow packet loss
  - per-flow load, throughput, and loss

Per-flow packet loss: buffer size

- Contribution of loss episodes of various lengths, averaged over all flows (100 sources)

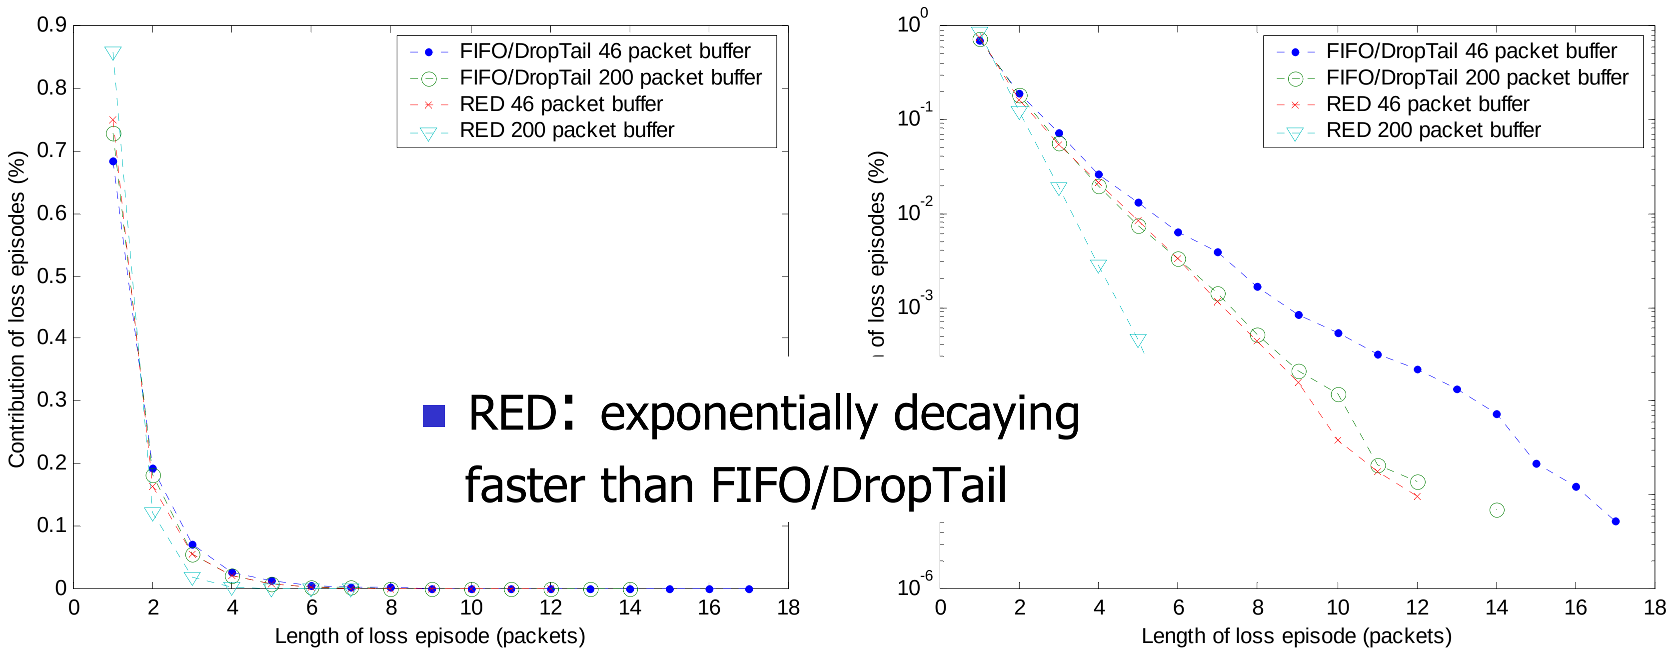

- RED: exponentially decaying faster than FIFO/DropTail

Per-flow load, throughput, and loss

- Traffic load from flow i:

$$load_i = \frac{total\ number\ of\ packets\ arrived\ from\ source\ i}{total\ number\ of\ packets\ arrived}$$

- Throughput from flow i:

$$TP_i = \frac{total\ number\ of\ packets\ from\ source\ i\ that\ are\ delivered}{total\ number\ of\ packets\ delivered}$$

- Loss from flow i:

$$loss_i = \frac{total\ number\ of\ packets\ from\ source\ i\ that\ are\ lost}{total\ number\ of\ lost\ packets}$$

Per-flow load, throughput, and loss

46 packet buffer

FIFO/DropTail

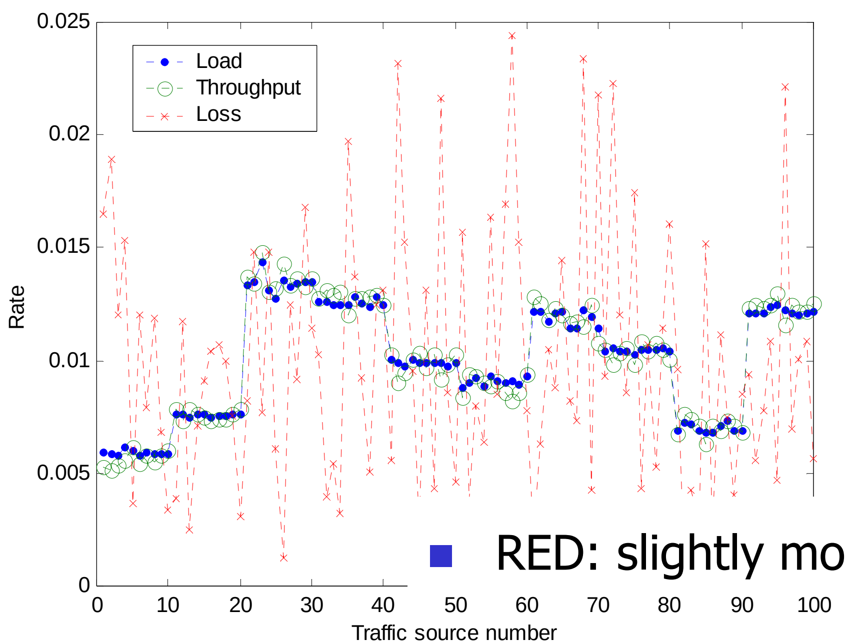

RED

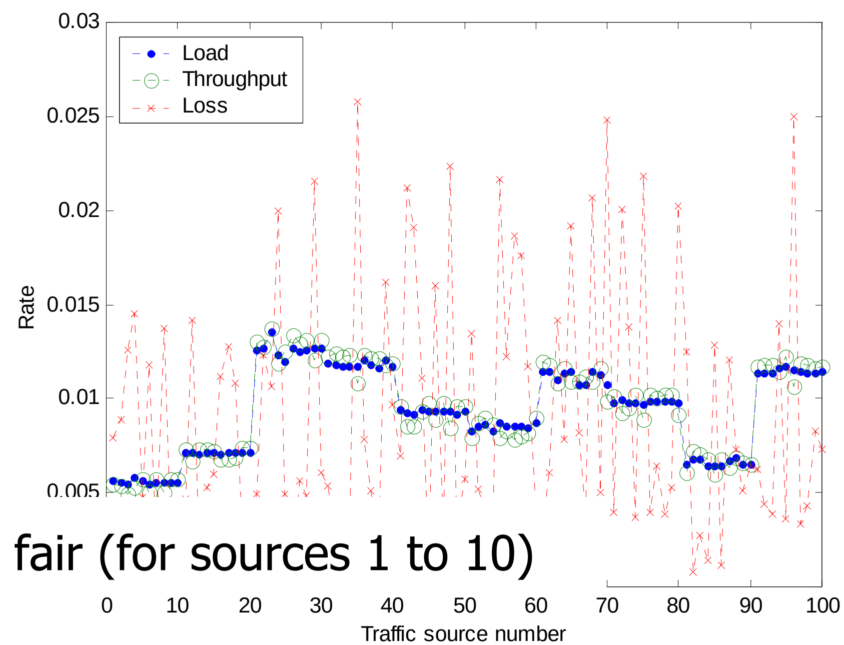

- RED: slightly more fair (for sources 1 to 10)

Simulation details

- Comparison between FIFO/DropTail and RED
- Comparison of FQ, SFQ, and DRR
  - FQ, SFQ, DRR employ per-flow queuing
  - ns-2 employs packet-queues for FQ and SFQ
  - FQ and DRR are designed to fairly serve traffic flows with variable packet sizes

Simulation results

- RED and FIFO/DropTail comparison
  - RED: better loss pattern and delay distribution
- FQ, SFQ, and DRR comparison
  - SFQ: best loss pattern and delay distribution
  - DRR: best fairness

Current work: queue management

- Impact of various queue management policies on packet loss and delay patterns:
  - FIFO/DropTail
  - Random Early Drop (RED)
  - Fair Queuing (FQ)
  - Stochastic Fair Queuing (SFQ)
  - Deficit Round Robin (DRR)
  - Class Based Queuing

Road map

- Modeling of Internet traffic
- Trace-driven simulations
- Simulation scenarios and results
- Wavelet analysis of traffic and packet loss
- Current and future work

Wavelet analysis

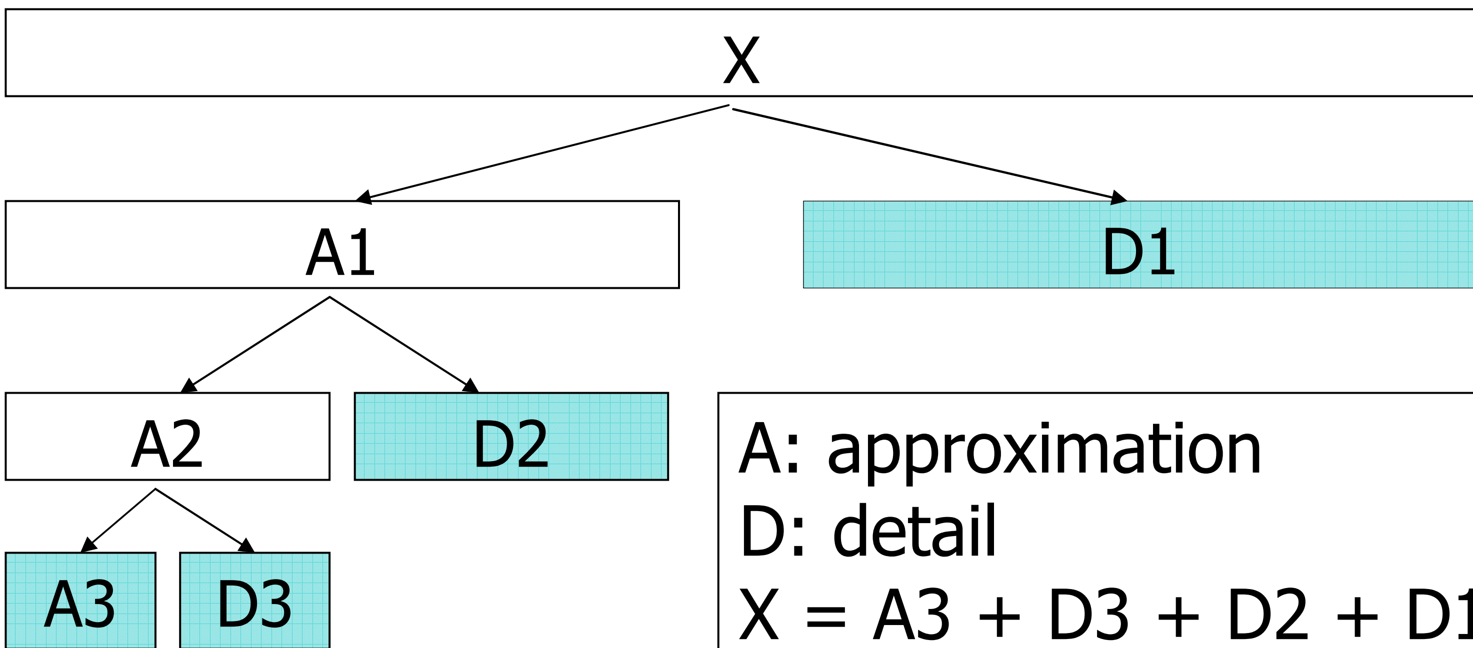

A: approximation
D: detail
X = A3 + D3 + D2 + D1

$$x(t) = \text{approx}_J(t) + \sum_{j=1}^{J} \text{detail}_j(t)$$

Wavelets

- Daubechies wavelets
- Scaling function:

$$\varphi(x) = \sum_k a_k^0 \cdot \varphi(mx - k)$$

- Wavelet function:

$$\psi^r = \sum_k a_k^r \cdot \varphi(mx - k)$$

r = 1, …, m-1

- m – rank, a - coefficients (Daubechies wavelets)

Daubechies wavelet (rank 2)

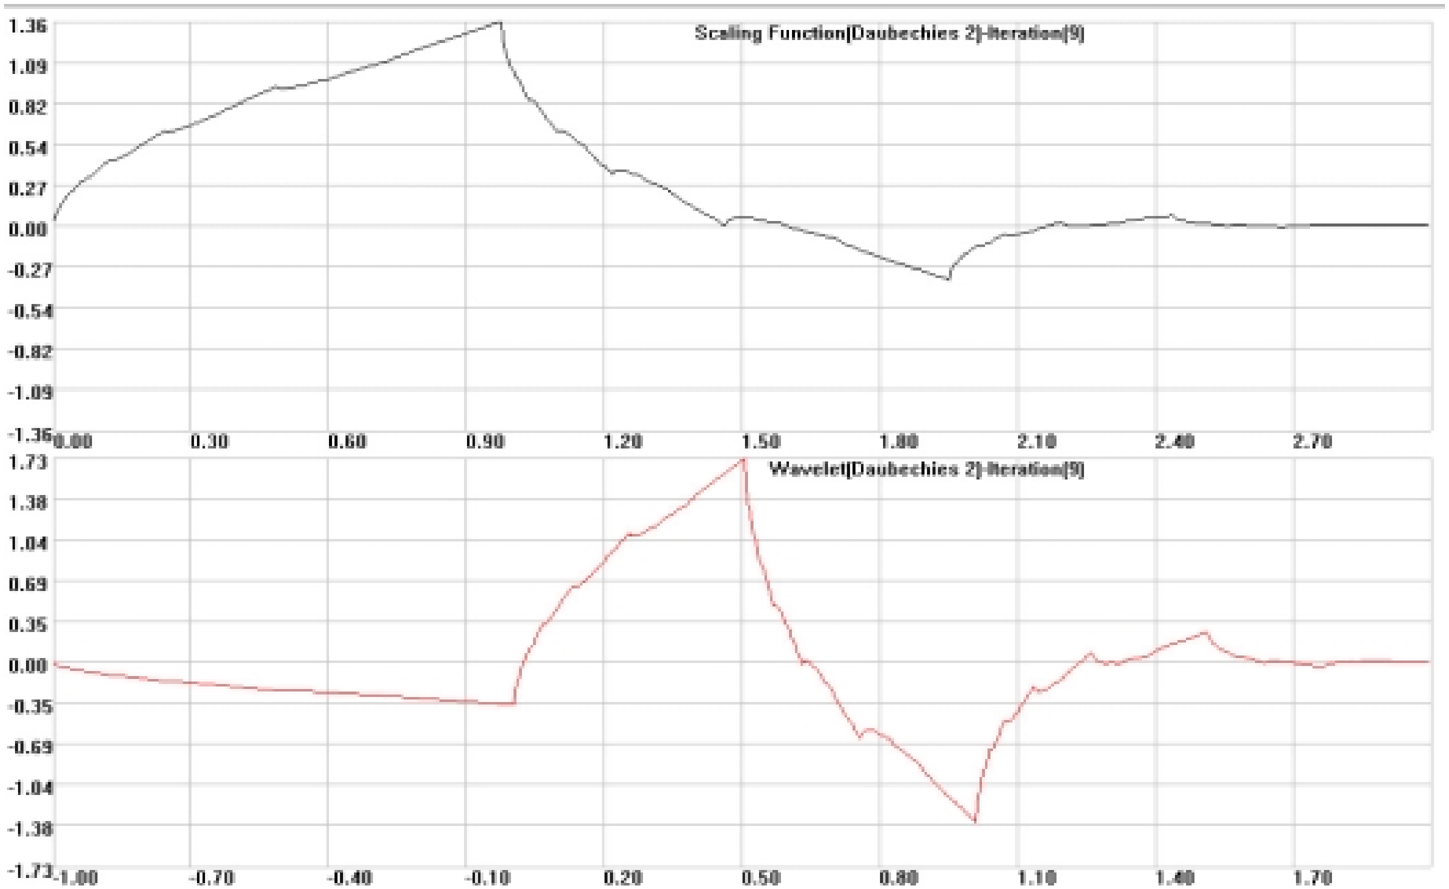

Daubechies wavelet (rank 9)

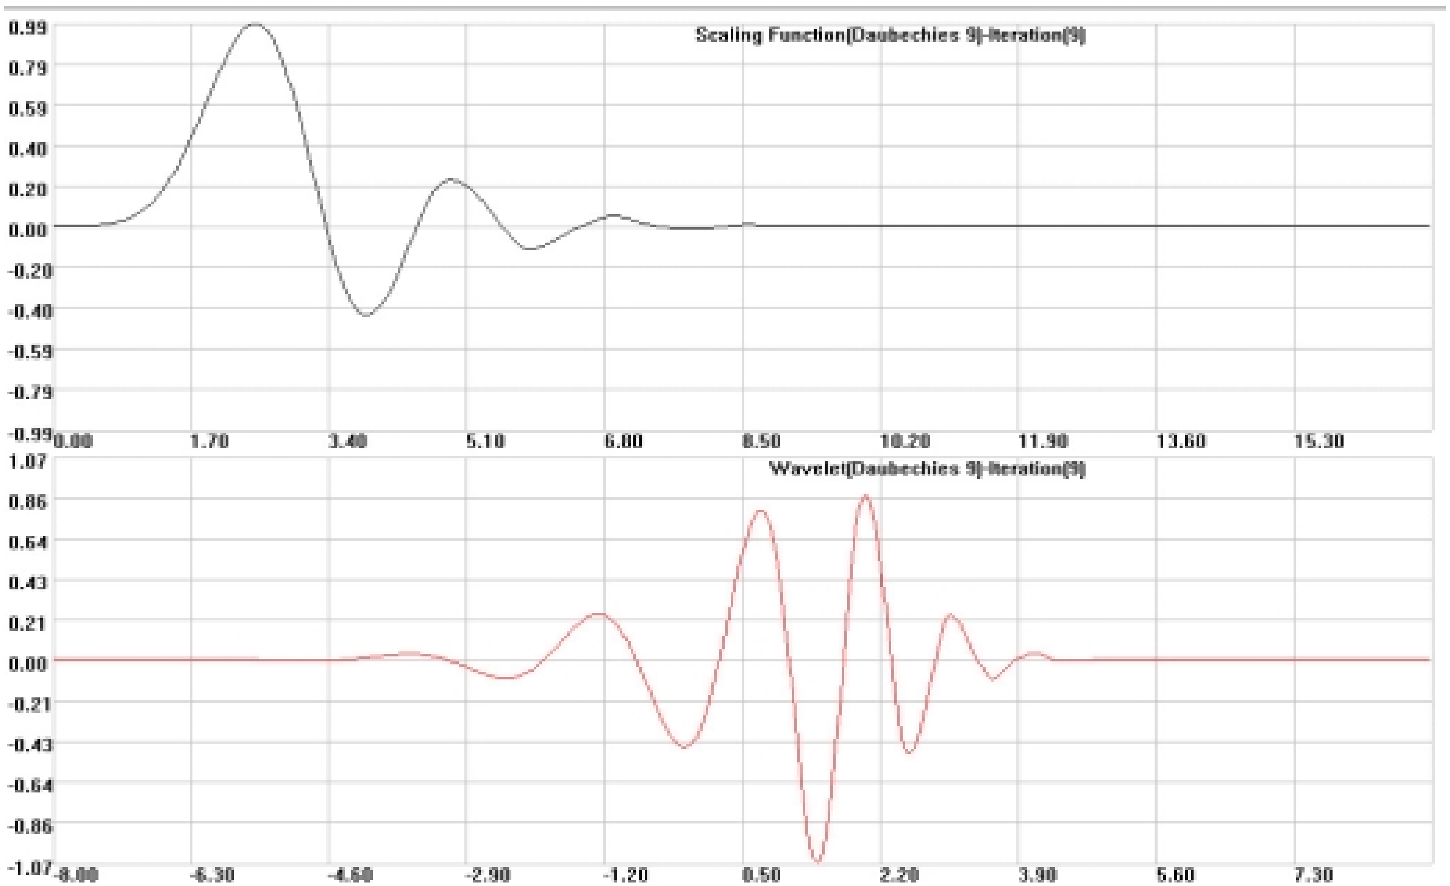

Monofractal estimator

- Monofractal property of a statistical process is homogenous: it does not depend on time scales.
- Monofractal wavelet estimator relies on second order statistics in characterizing a statistical process.

Wavelet-based Hurst parameter estimator

$$\Gamma(2^{-j} v_0) = \frac{1}{n_j} \sum_k \left| d(j,k) \right|^2$$

Γ: spectrum of time series
$d(j,k)$: detail coefficients at level j
n_j: number of coefficients at level j

P. Abry and D. Veitch, "Wavelet analysis of long-range-dependent traffic," *IEEE Trans. Information Theory*, vol. 44, pp. 2-15, 1998.

Abry-Veitch estimator

- Linear relationship between $\log_2(\Gamma)$ and j indicates LRD.
- $\log_2(\Gamma) = (2H - 1)j + c$, H: estimated Hurst parameter
- Scale level j:
  - time-scale = 2^j
  - large j implies coarser time-scales

Wavelet analysis of Ethernet traffic traces from Bellcore

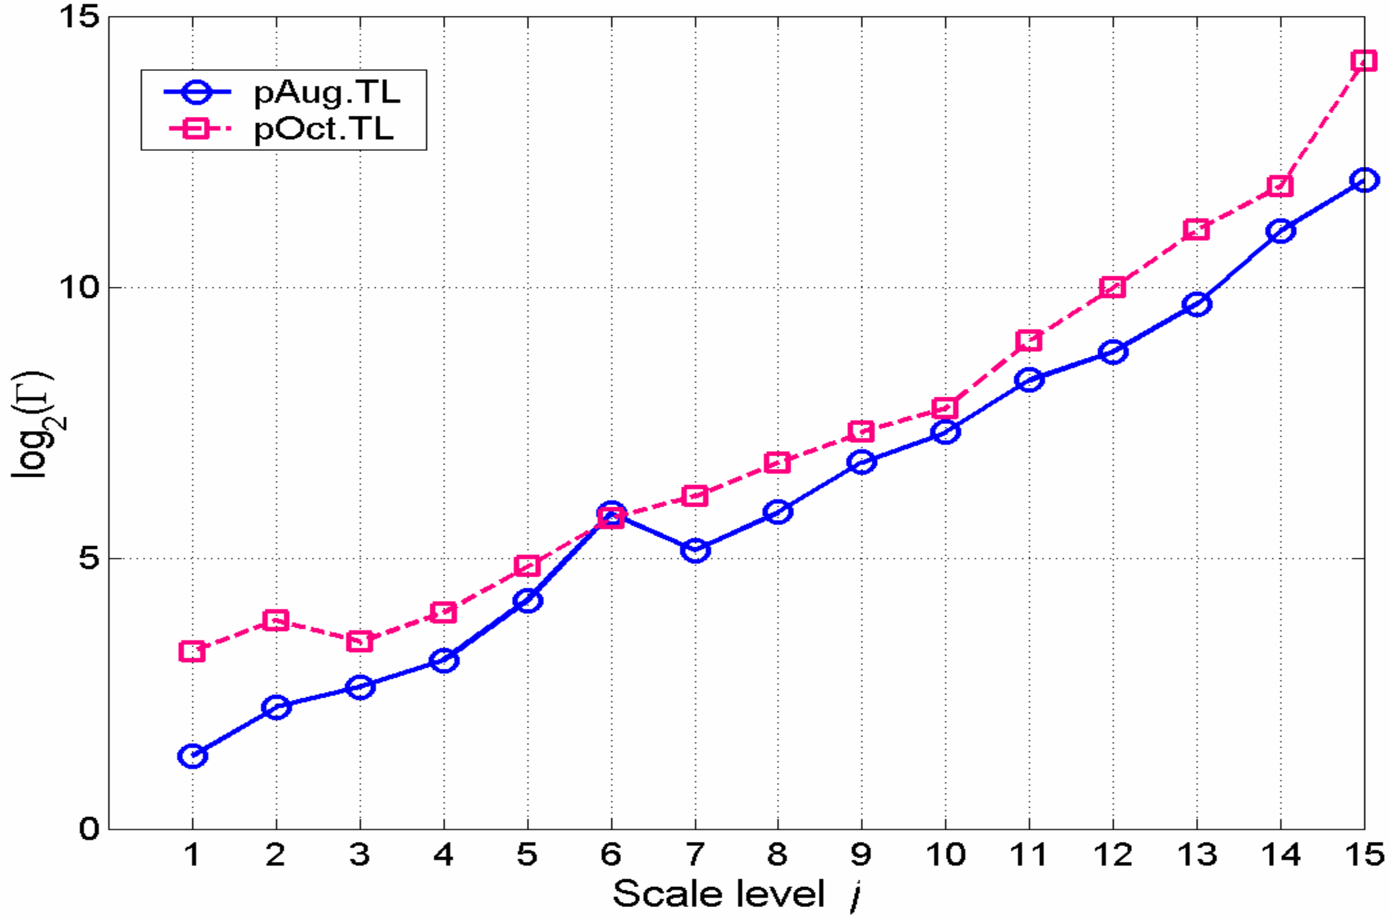

Wavelet analysis of the "Star Wars" traffic trace

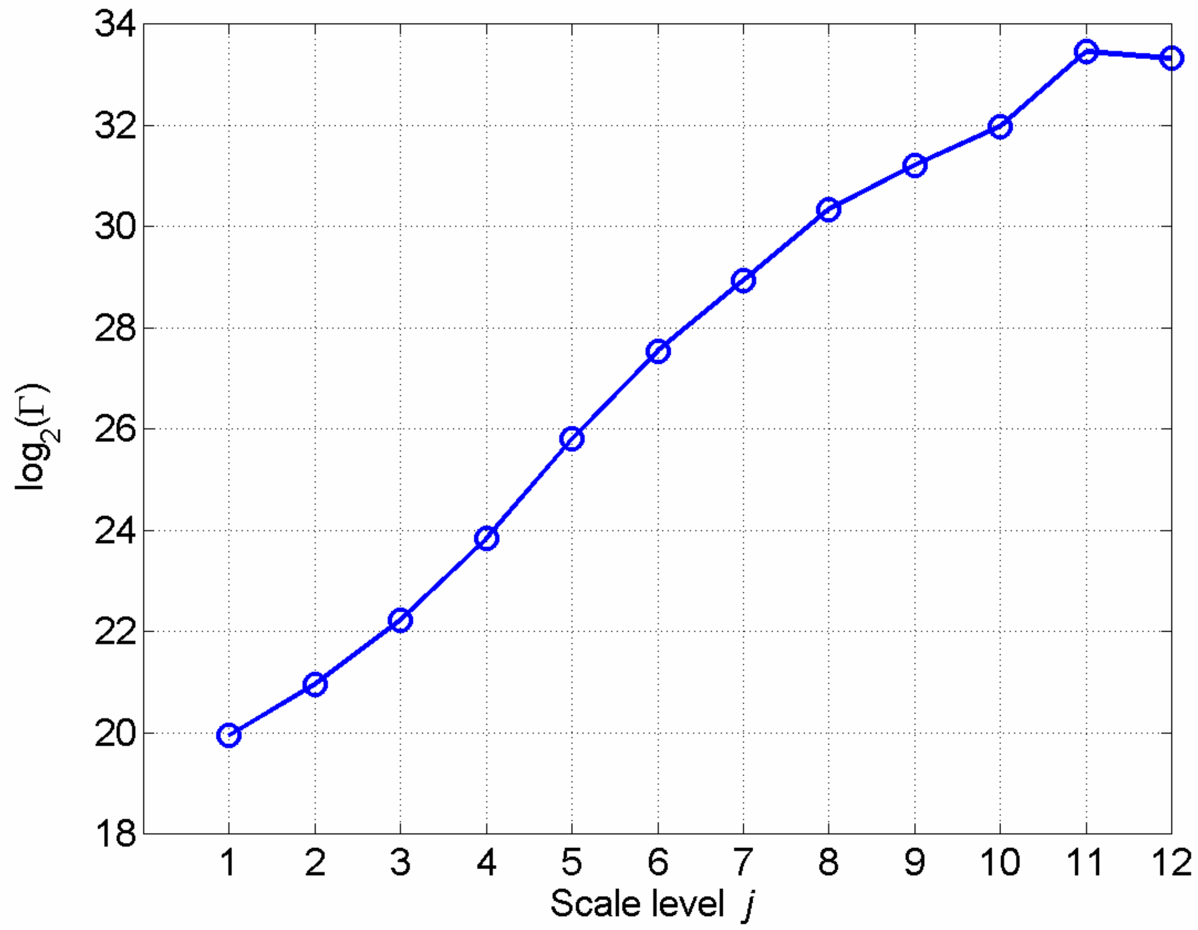

Wavelet analysis of the traffic trace after packetization

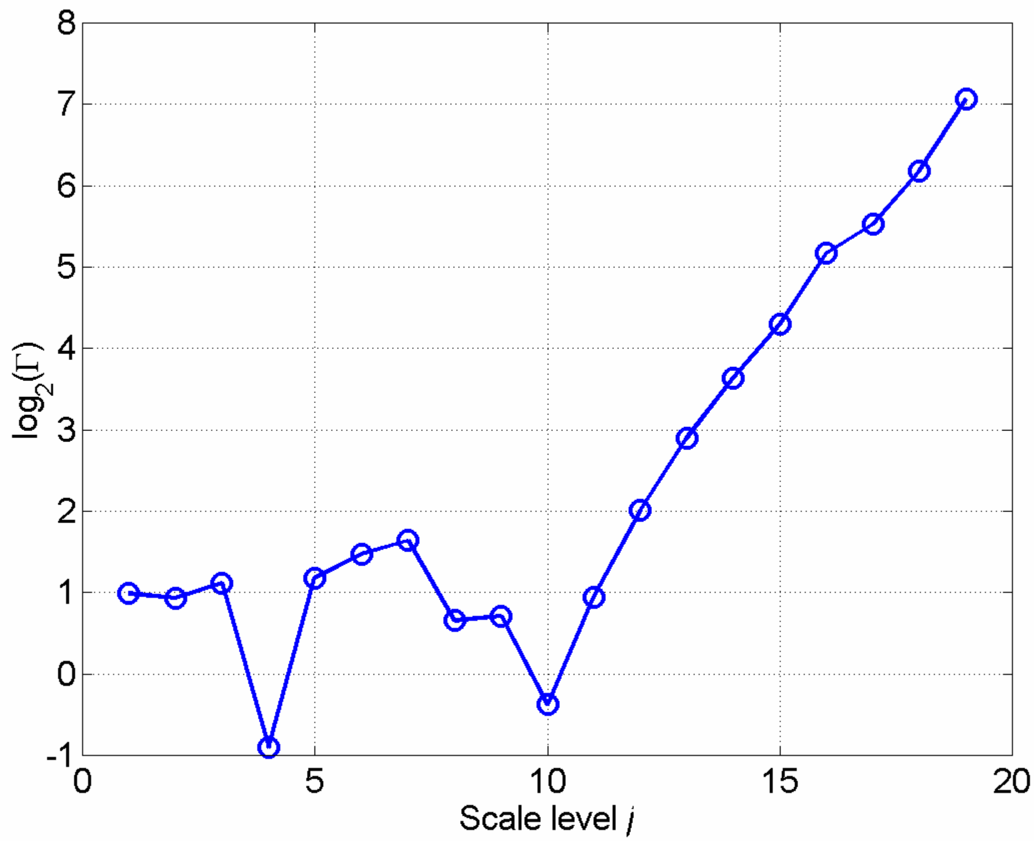

Graph exhibits a linear behavior for $j \in [10, 19]$.

Aggregate traffic arriving to router R

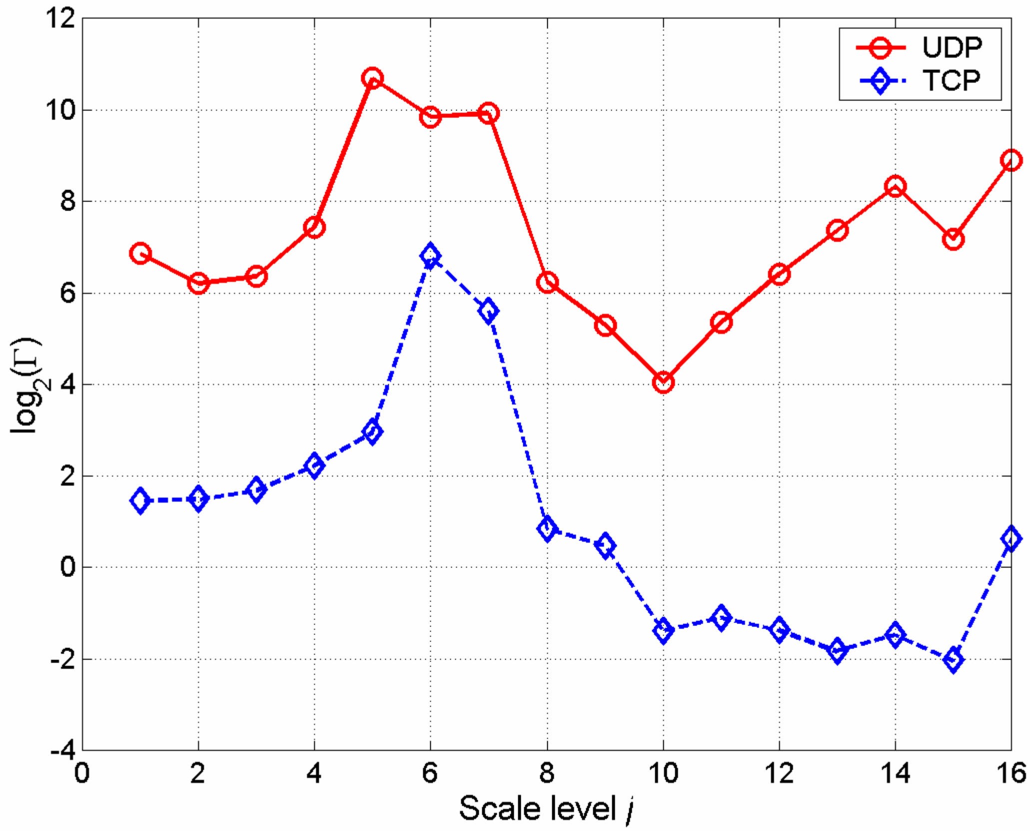

Graph exhibits a linear behavior for $j \in [10, 14]$.

Characteristics of the “Star Wars” traffic trace

- “Star Wars” traffic is of medium-burstiness.
- Estimated Hurst parameter is H = 0.84 ~ 0.97:
  - R/S analysis
  - variance-time plots
- The coexistence of both:
  - LRD: long-range dependent
  - SRD: short-range dependent

  components is due to the scene changes and the coding algorithms.

Wavelet analysis of packet loss

- We use wavelet-based analysis to detect both the presence and the location of LRD.
- Loss rate process:
  each sample represents the number of lost packets over 1 msec interval.

ns-2 simulation results: packet loss

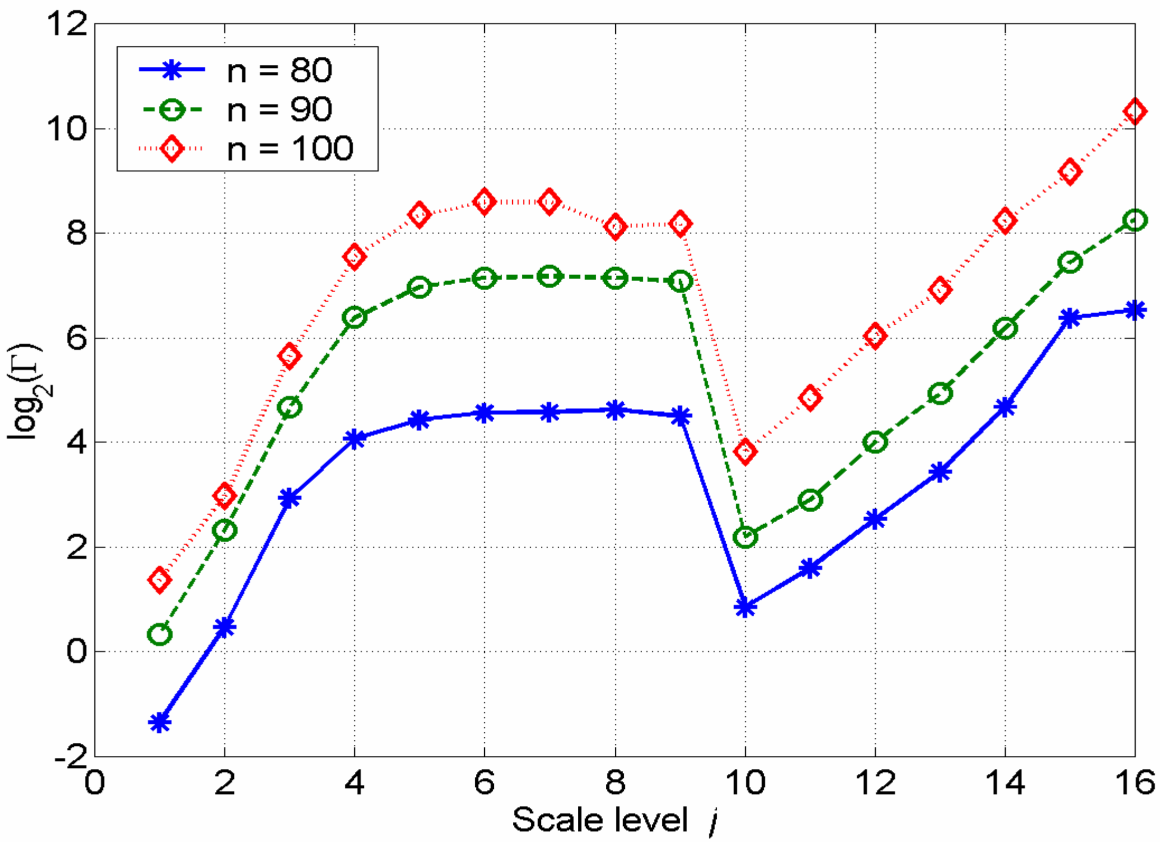

Buffer size B = 100 Kbytes.

Analysis of packet loss on multiple time-scales

- Similarity between traffic and packet loss patterns:

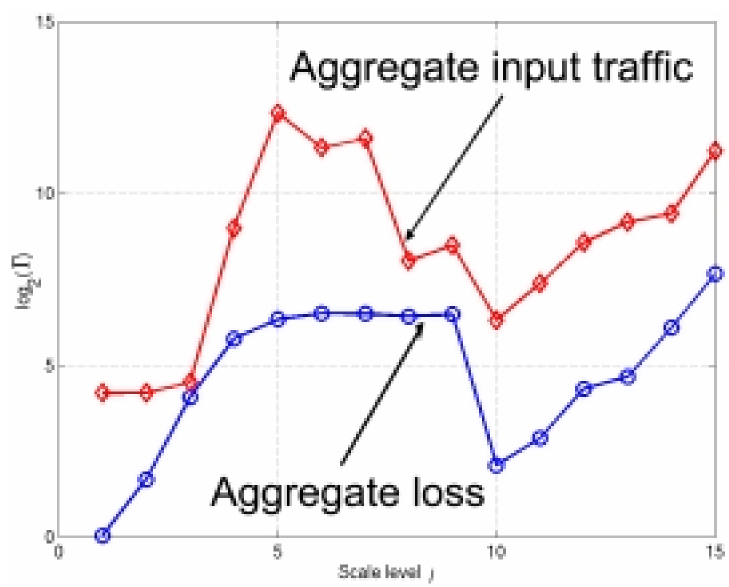

UDP scenario

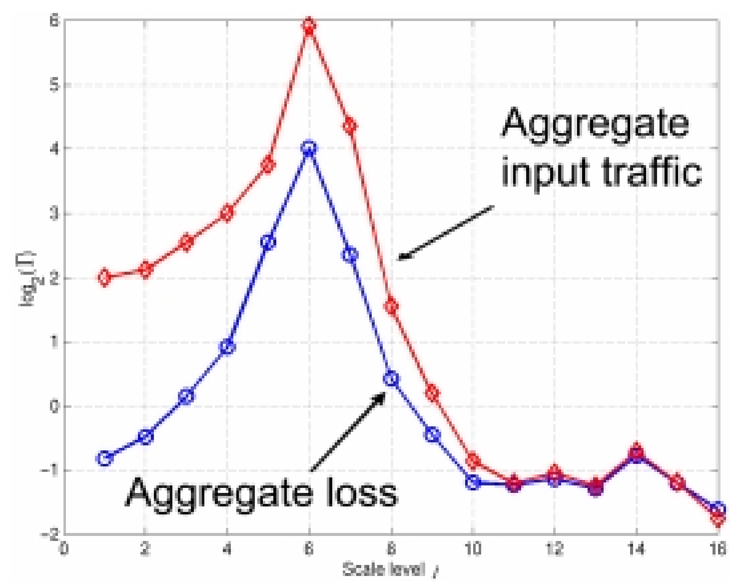

TCP scenario

Packet loss and time-scales

- Packet loss behavior and loss properties vary over different time-scales.
- Linear relationship between $\log_2(\Gamma)$ and j is evident for the coarser scales, i.e., beyond the break-point $j = 10$ (equivalent to 2^{10} msec).

Video traces: sources

- Traffic traces from Wurzburg:
  - MPEG-1 coded movies
- Traffic traces from Berlin:
  - MPEG-4 and H263 coded movies and stills sequences

Video traces: coders

- Three coders: MPEG-1, MPEG-4, H263
- MPEG: Motion Picture Experts Group
- H263: ITU recommendation H263
- All three coders use:
  - discrete cosine transformation
  - predictive encoding
- MPEG-1 and H263 coders are similar and both are frame based
- MPEG-4 is object based and it is more efficient

Video coders: examples

- Wavelet analysis may indicate a difference between coders.
- We experimented with video traces coded with both MPEG-4 and H263.
- Example: 30-minute trace from a cartoon named "Futurama".

Wavelet analysis: Futurama (MPEG4)

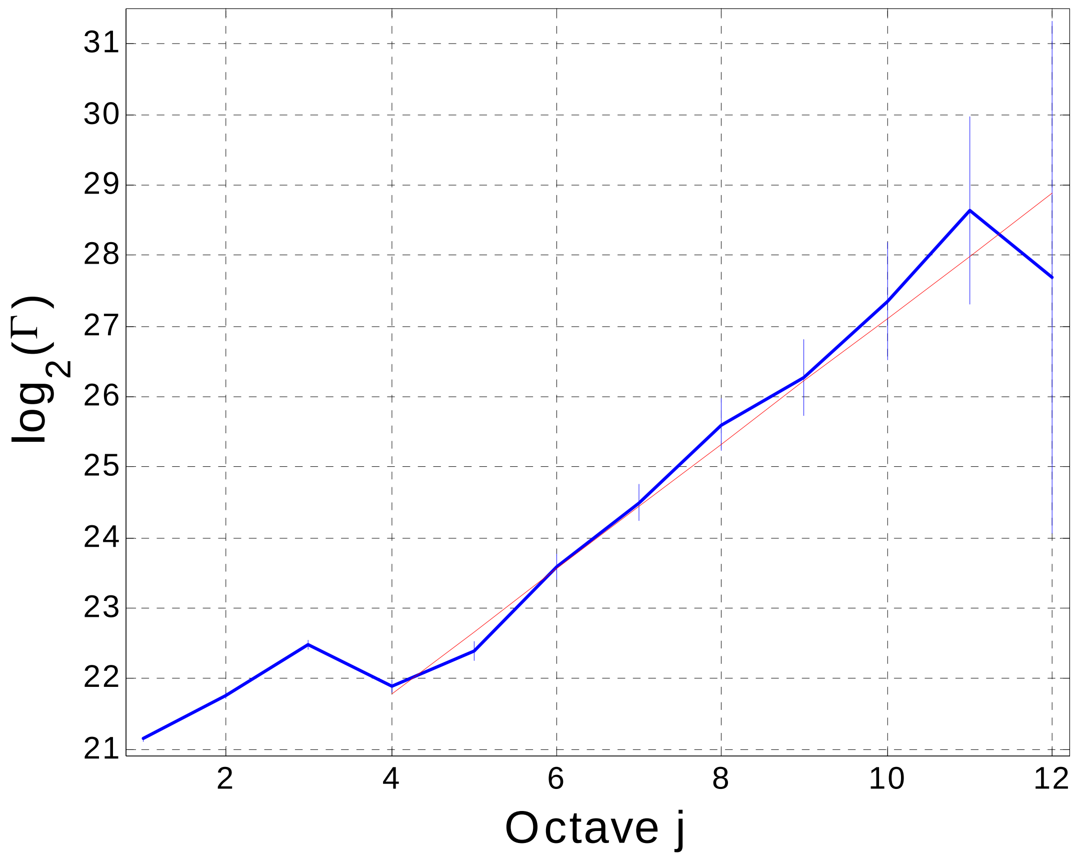

Wavelet analysis: Futurama (H.263)

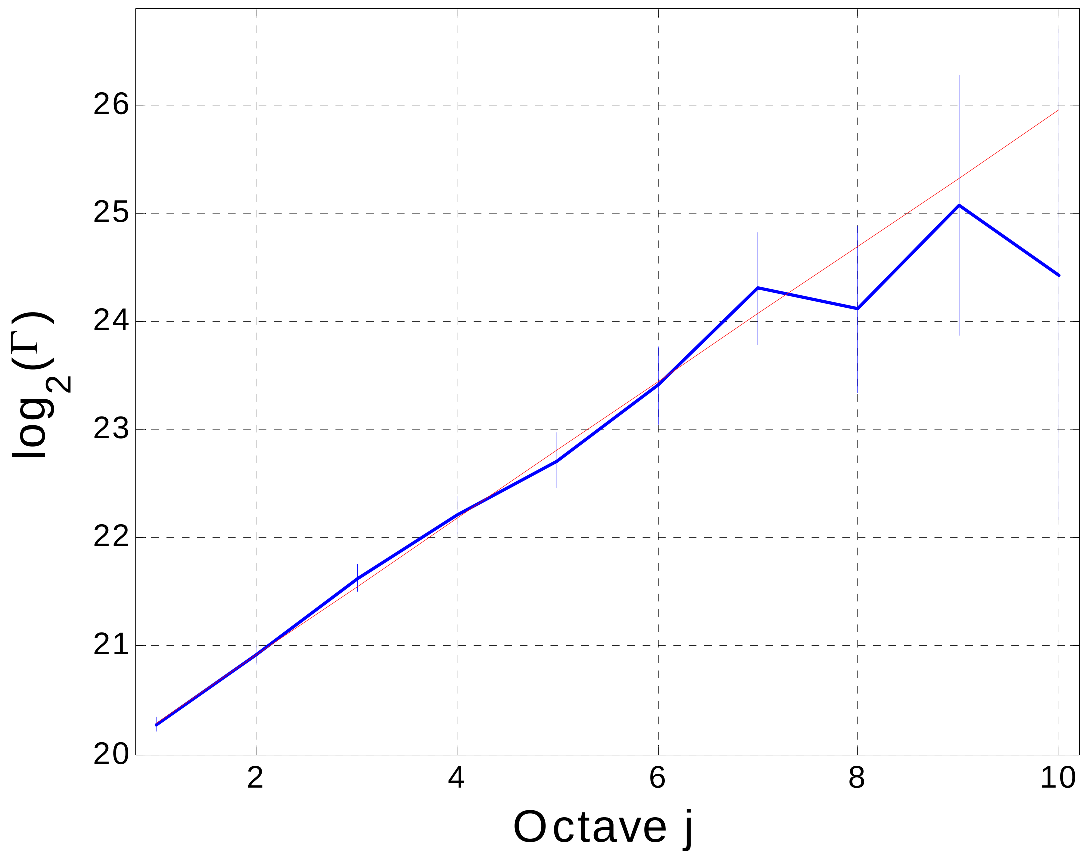

Analysis of loss traces

- 100 sources, various queue sizes
- Buffer size: 46, 100, 200 packets (552 bytes)
- The same video traces were used as traffic sources
- 10-minute loss processes
- UDP transfers
- Whether and how queue size affects the LRD
- Binning: 0.1 msec

Loss trace, 46 packets, monofractal

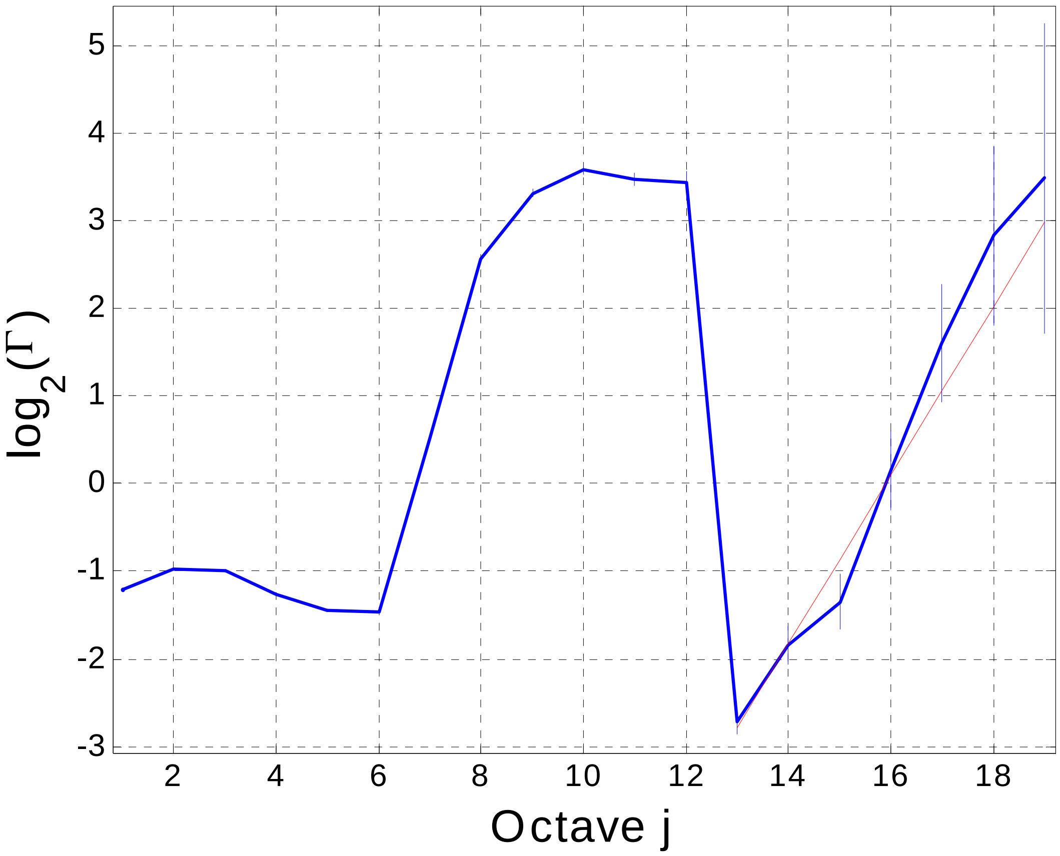

Observations

- Regardless of the buffer size, LRD is present with the same break-point and similar estimated values of the H parameter.
- Various wavelet estimators yield different values of H.
- While the obtained graphs are fairly consistent, the H parameter estimation may yield non-physical values.
- The difference in performance may be attributed to:
  - trends introduced by video coders
  - choice of the algorithm.

Concluding remarks

- Time-scales are important for estimating behavior of packet loss.
- Wavelet-based analysis proved useful for finding time-scale break-points beyond which long-range dependency can be detected.

Current work: queue management

- Impact of various queue management policies on packet loss and delay patterns:
  - FIFO/DropTail
  - Random Early Drop (RED)
  - Fair Queuing (FQ)
  - Stochastic Fair Queuing (SFQ)
  - Deficit Round Robin (DRR)
  - Class Based Queuing

Publications

- V. Markovski, F.Xue, and Lj. Trajkovic, "Simulation and analysis of packet loss in User Datagram Protocol transfers," *The Journal of Supercomputing*, vol. 20, pp. 175 – 196, 2001.
- F. Xue, V. Markovski, and Lj. Trajkovic, "Packet loss in video transfers over IP networks," *Proc. IEEE Int. Symp. Circuits and Systems*, Sydney, Australia, May 2001, vol. II, pp. 345 – 348.
- V. Markovski and Lj.Trajković, "Analysis of loss episodes for video transfers over UDP," *Proc. Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS 2000)*, Vancouver, BC, Canada, July 2000, pp. 278 – 285.
- F. Xue and Lj.Trajković, "Performance analysis of a wavelet-based Hurst parameter estimator for self-similar network traffic," *Proc. Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS 2000)*, Vancouver, BC, Canada, July 2000, pp. 294 – 298.
- F. Xue, V. Markovski, and Lj. Trajković, "Wavelet analysis of packet loss for video transfers over UDP," *Proc. First International Conference on Internet Computing (IC 2000)*, Las Vegas, NV, USA, June 2000, pp. 427 – 433.

References: on traffic

- W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Trans. Networking,* vol. 2, pp. 1 – 15, 1994.
- M. W. Garrett and W. Willinger, "Analysis, modeling and generation of self-similar VBR video traffic," *Proc. ACM SIGCOMM'94*, London, UK, Aug. 1994, pp. 269 – 280.
- J. Beran, R. Sherman, M. S. Taqqu, and W. Willinger, "Long-range dependence in variable bit rate video traffic," *IEEE Trans. Commun.,* vol. 43, pp. 1566 – 1579, 1995.
- C. W. J. Granger and R. Joyeux, "An introduction to long-memory time series models and fractional differencing," *J. of Time Series Analysis*, vol. 1, pp. 15 – 29, 1980.

References: on traffic

- O. Rose, “Statistical properties of MPEG video traffic and their impact on traffic modeling in ATM systems,” Technical Report 101, University of Wuerzburg, Institute of Computer Science, Am Hubland, 97074 Wuerzburg, Germany, Feb. 1995.
- L. Breslau, “Example traffic trace for ns”: http://www.research.att.com/~breslau/vint/trace.html.
- D. S. Swayne, D. Cook, and A. Buja, “XGobi: Interactive dynamic data visualization in the X window system,” *Journal of Computational and Graphical Statistics*, vol. 7, no. 1, pp. 113 – 130, Mar. 1998.

References: on packet loss

- J.-C. Bolot, "End-to-end packet delay and loss behavior in the Internet," *Proc. ACM SIGCOMM '93 Conference on Communications Architectures, Protocols and Applications*, pp. 289 – 298, San Francisco, CA, USA, Sept. 1993.
- M. S. Borella, D. Swider, S. Uludag, and G. B. Brewster, "Internet packet loss: measurement and implications for end-to-end QoS," *Proc. Int. Conference on Parallel Processing*, Minneapolis, MN, Aug. 1998, pp. 3 – 15.
- J. M. Boyce and R. D. Gaglianello, "Packet loss effects on MPEG video sent over the public Internet," *Proc. ACM Multimedia '98*, Bristol, UK, Sept. 1998, pp. 181-190.
- M. Yajnik, S. Moon, J. Kurose, and D. Towsley, "Measurement and modeling of the temporal dependence in packet loss," *Proceedings IEEE INFOCOM*, pp. 345 – 352, New York, NY, USA, Mar. 1999.

References: on queuing schemes

- A. Demers, S. Keshav, and S. Shanker, "Analysis and simulation of a fair queuing algorithm," *Proc. ACM SIGCOMM '89*, Sept. 1989, pp. 1 – 12.
- S. Floyd and V. Jacobson, "Random early drop gateways for congestion avoidance," *IEEE/ACM Transactions on Networking,* vol. 1, no. 4, August 1993.
- P. McKenney, "Stochastic fairness queuing," *Internetworking: Research and Experience,* vol. 2, pp. 113 – 131, Jan. 1991.
- M. Shreedhar and G. Varghese, "Efficient fair queuing using deficit round robin," *Proc. of ACM SIGCOMM '95*, Aug. 1995, pp. 231 – 242.

References: on fractal estimators

- P. Abry and D. Veitch, "Wavelet analysis of long-range dependent traffic," *IEEE Trans. Information Theory*, vol. 44, pp. 2 – 15, Jan. 1998.
- P. Abry, P. Flandrin, M. Taqqu, and D. Veitch, "Wavelets for the analysis, estimation and synthesis of scaling data," *Self Similar Network Traffic Analysis and Performance Evaluation*, K. Park and W. Willinger, Eds., John Wiley, 2000, pp. 39 – 84.
- D. Veitch: "Matlab code for the estimation of scaling exponents": http://www.emulab.ee.mu.oz.au/~darryl/MS_code.html
- A. Feldmann, A. C. Gilbert, and W. Willinger, "Data networks as cascades: Investigating the multifractal nature of Internet WAN traffic, *Proc. ACM SIGCOMM '98*, Aug. 1998, pp. 42 – 55.
- A. Feldmann, A. C. Gilbert, P. Huang, and W. Willinger, "Dynamics of IP traffic: A study of the role of variability and the impact of control, *Proc. ACM SIGCOMM '99*, Aug. 1999, pp. 301 – 313.

Vancouver

Vancouver

Vancouver