Characterization of Internet Traffic

S-38.149 Postgraduate Course in Teletraffic Theory

Esa Hyytiä
email: esa.hyytia@hut.fi

11th October 1999

Abstract

The aim of this paper is to give an overview of traffic seen in Internet and what kind of models can be used to describe it. The traffic in Internet is an aggregate of traffic from many local area networks (LAN). Measurement results strongly suggest that the traffic is self-similar with huge variation of timescales, from order of milliseconds to even days.

This report is based mostly on ten articles written during the years 1991-1999. The most of the articles consider the measurement and analysis of local area network traffic, what goes wrong with traditional models (Poisson) and how well self-similar processes model the traffic. Also the traffic caused by world wide web is analyzed individually.

Glossary:

B-ISDN	broadband ISDN
CD	collision detection
CSMA	carrier sense multiple access
FBM	fractional brownian motion
IP	internet protocol
LLCD	log-log complementary plots
MAC	medium access control
RFC	request for comments, internet standards
TCP	transmission control protocol
URL	Universal Resource Locator

1 Introduction

The internet consists of huge number of local area networks (LAN) connected together (fig. 1). The traffic in WAN is an aggregate of of LAN traffic and in order to understand the nature of the traffic existing in the internet one must first study the traffic in LAN and even applications used there. Understanding the nature of LAN traffic helps network designers

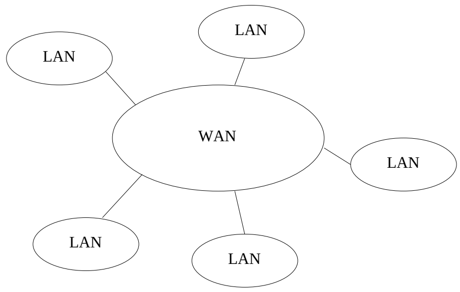

Figure 1: Local area networks connected to wide are network (like internet)

in their work when planning the future networks. Also prediction of possible bottlenecks in network is important. The subject also includes understanding the consequences of used protocol (TCP/IP mostly) and applications, as well as user behavior.

1.1 Material

This report is based on ten articles listed below. The articles are in the order of the publication date and their main results are pointed out in the table. For relations between articles see also figure 2.

[LW91]	Leland, Wilson	*High-Time-Resolution Measurement and Analysis of LAN Traffic: Implications for LAN interconnection* : The first high-resolution measurements of LAN traffic and discovery of the burstiness of traffic across time domains of many orders
[FL91]	Fowler, Leland	*Local Area Network Traffic Characteristics, with Implications for Broadband Network Congestion Management*: modest increases in buffer size are not significant due the burstiness of traffic
[WTLW95]	Leland, Taqqu, Willinger, Wilson	*On the Self-Similar Nature of Ethernet Traffic (Extended Version)*: This is apparently an extended version of a paper presented in ACM Sigcomm '93 and largely referenced. It collects earlier results of authors and extends them, and contains a traffic models based on self-similar processes.
[LWTW94]	Leland, Willinger, Taqqu, Wilson	*Statistical Analysis and Stochastic Modeling of Self-Similar Datatraffic*: -

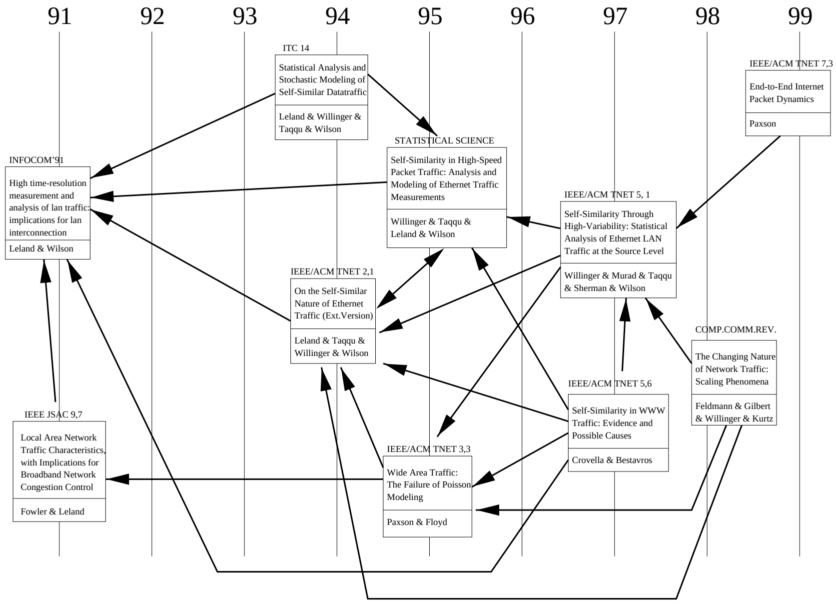

Figure 2: Used articles and references between them in graph-form.

[PF95]	Paxson, Floyd	*Wide Area Traffic: The failure of Poisson modeling*: The first empirical evidence for large-time scaling phenomena in measured WAN traffic.
[WTLW95]	Willinger, Taqqu, Leland, Wilson	*Self-Similarity in High-Speed Packet Traffic: Analysis and Modeling of Ethernet Traffic Measurements*: This paper also uses same measurements first presented in [LW91]. The paper includes sections about statistical tests for self-similarity discusses about possible use of distributed computing with traffic analysis.
[WTSW97]	Willinger, Taqqu, Sherman, Wilson	*Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level*: The last common paper from the authors (from ones listed here). Its main results are that a superposition of many ON/OFF sources (with certain requirements) exhibit the self-similar nature observed in Ethernet LAN traffic and while network traffic is complex in nature still quite simple models can be used to generate it.
[FGWK98]	Feldmann, Gilbert, Willinger, Kurtz	*The Changing Nature of Network Traffic*: Authors use wavelet-based techniques in analysing LAN traffic data.
[LTW94]	Willinger, Taqqu, Leland, Wilson	*Self-Similarity in Hihg-Speed Packet Traffic: Analysis and Modelling of Ethernet Traffic Measurements*: A mathematical presentation of the problem including definitions of self-similar processes, their properties, statistical tests for self-similarity and analysis of ethernet trffic measurements from [LW91].
[CB97]	Crovella, Bestavros	*Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes*: In this paper authors present measurements about WWW traffic and show it has self-similar characteristics. They also try to explain possible reasons for self-similarity.
[Pax99]	Paxson	*End-to-End Internet Packet Dynamics*: This paper contains description of measurements between 35 sites and discusses the about packet level problems/characteristics.

The process starts from report vt Leland and Wilson presented in INFOCOMM'91. Leland and Wilson are (co)authors of many other papers listed in table too. Later Taqqu and Willinger have joined the group and their papers contain many major results presented here.

1.2 Ethernet

The ethernet is clearly the most popular local area network technology in use today. It was originally designed for coaxial cable but currently new networks are mostly built using twisted pair cable. The todays ethernet LANs have usually the transmission rate of 10 Mbit/s, while 100 Mbit/s networks are been deployed with increasing rate. Ethernet networks use carrier sense multiple access with collision detection protocol (CSMA/CD) for medium access control (MAC). Another popular MAC layer is token passing, both described in IEEE 802 standars [HO88].

Basicly the transmission of data packet in ethernet network is very simple. Assume a node A wants to transmit a packet to another node B in the same network.

1. The node A listens the bus to see whatever there is a transmission under way. If there is, the node A waits until the transmission ends and after a small delay (min. $9.6\ \mu$s) tries sending again.
2. The bus is not in use, and node A starts sending the packet.
3. At same time the node also listens the bus to see if someone else is sending too. If that happens there is a *collision*, and transmission is repeated after a random amount of time (ie. back to step 1).

Under normal operation with moderate load the packet collisions are quite rare events. The IEEE 802 standard recommends of data field lengths between 46 and 1500 bytes. The first measurements of ethernet traffic are from 80s by Shoch and Hupp [LW91].

1.3 TCP/IP protocol

The TCP/IP protocol is used in internet. So a brief description of TCP/IP is in place here to understand possible consequences of the used protocol and the applications to packet level traffic behaviour. The whole TCP/IP protocol (or internet protocol) is defined in internet standards called requests for comments (RFC) freely available for example from http://www.ietf.org/as well as from many mirrors.

1.3.1 IP - Internet protocol

The internet protocol is connectionless protocol used to deliver a single packet trough the network. IP header contains fields for source and target address (both 32 bits long), the type of packet and number of hops. While travelling through network the hop counter is reduced at each node and if it reaches zero, the packet is dropped and source node is notified about it.

The IP protocol does not quarantee that the order of packets is same at the recieving end, or even that the packet ever reaches the destination. Also duplicate packets may be recieved. The upper level protocol (TCP or application) must take care of such situations and handle them appropriately.

The upper level protocols UDP and ICMP are essentially a way to send certain kind of IP packet. The ICMP packets contains information about the state of network while UDP packets are used by some applications.

1.3.2 TCP - Transmission control protocol

The transmission control protocol (TCP) is a connection oriented protocol where a connection is first established between source and target nodes and after that packets may be sent. Once connection is set up, data can be transmitted “safely” through the network in both directions. If a connection is lost (ie. some packets have not reached the destination end) the application is notified about it. TCP handles automatically time-outs and retransmissions. TCP protocol also includes a sliding window flow-control mechanism to adjust the transmission rate. There may be maximum N transmitted packets out at a time and no new ones will be sent until one of them is acknowledged by reciever. The application is identified in both ends by a *port number*.

The range of application using TCP in internet is large: World Wide Web (WWW) or HTTP (port 80), File Transfer Protocol (FTP, port 21), Telnet (port 23) etc. Especially WWW traffic has been lately the major factor in the growth of the internet traffic.

2 Measurements

Measuring the network traffic in a local area network (LAN) is not a straightforward task. Ideally we want to record the whole data packet (including application, source and destination) and the time when the packet was sent (timestamp). The amount of data seen in LANs is very huge. Thus if long time intervals are suppose to be captured, packing of data is necessary. What is important data depends on what is been analysed.

2.1 The Measurements of Ethernet traffic

The traffic traces can be recorded with a software solution where a workstation is used to record the data packets. In this method received data packets are first handled by operating system, and then handled to application (after possible buffering). Thus the time when packet was sent is recorded in a little inaccurate way. The operating system can in some cases be modified to improve the accuracy. The generally the software approach has some limitation:

- during the heavy traffic load some packets might be lost, typically 5-10 %
- the accuracy of timestamps is not high, while writing to disk etc. the arrived packet might not be served right away and the timestamps are inaccurate.

Another way is to use special hardware for the job.

In [LW91] authors describe an ethernet traffic monitoring system consisting of a special hardware and a workstation. The ethernet interface was controlled by 20 MHz MC68030-based single board computer (SBC) which was used to timestamp and preprocess incoming packets. As SBC has no other tasks the accuracy of time stamps was good. The preprocessed packets were then turn over to the workstation which saves them first on disk and from there to a tape. With this configuration authors were able to store several days traces of LAN traffic without packet losses. The later work of authors use these same measurements.

The measurements used by Paxson in [Pax99] are available from http://ita.ee.lbl.gov/. The recording of traffic traces is done by *tcpdump*, which a widely available in unix environments.

2.2 The measurements of WWW traffic

In [CB97] authors describe a system where they used a modified version of publicly available world wide web browser NCSA Mosaic to collect better statistics. The earlier reports had used the log files from www-proxies or servers to analyze the www-traffic. The additional information authors got from their measurements were the behavior of the user itself. The modified version of NCSA Mosaic browser recorded the uniform resource locator (URL) of the files user retrieved and the time required to transfer the file. In the analysis authors only used the files which were not in the cache of the browser at the time of request. The accuracy of time stamps was 10 ms.

The measurements took place in the Department of Computer Science at the University of Boston. The site consisted of 37 workstations connected to local network. The data contains the WWW sessions of about 600 different users and during the measurements 4700 sessions were recorded giving total number of about 600000 URL requests. The total number of bytes transferred were about 2 gigabytes.

3 Traffic Models

Good models for data traffic are essential for simulations of different network architectures. If poor models are used one might end up with wrong conclusions.

3.1 Traditional models for arrivals

In the following several traditional models for packet arrivals are listed. The models are analytically tractable, but have some flaws when fitting them with experimental data.

1. Poisson arrivals, the number of arrivals in time interval T follows the exponential distribution with parameter λT
2. Batch Poisson, where inter-arrival times are still exponentially distributed, but the number of customers arriving at same time follows some random distribution instead of being constant one.
3. Hyperexponential, where system consists of K parallel Poisson-servers. Each time one of them is chosen randomly with some probability. Ie. $f(x) = \sum_{i=0}^{K} p_i \mu_i e^{-\mu_i x}$.

4. Markov-modulated Poisson process (MMPP), where the arrival rate λ of Poisson process is determined by the state of background Markov-process

A more detailed definitions of models listed above can be found for example from [HP93]. These models work well for example in case of traditional telephone network.

3.2 Burstiness

The major flaw with traditional traffic models is that they do not model the burstiness of internet traffic correctly. There are three commonly used definitions for burstiness:

1. the ratio of peak bandwidth to mean bandwidth
2. the coefficient of variation c: $c^2 = \frac{\text{Var}[X]}{\text{E}[X]}$. For given time interval coefficient of variation is the ratio of standard deviation in the number of arrivals to mean number of arrivals. The coefficient of variation can be defined to inter-arrival times as well. For Poisson arrivals the coefficient of variation $c \propto 1/\sqrt{L}$.
3. the index of dispersion (for arrivals) for given time interval is the ratio of the variance of arrivals to mean number of arrivals. To Poisson-arrivals this is equal to one. The index of dispersion can be defined to inter-arrival times as well.

Results from measurements suggests that internet traffic is indeed non-Poisson. The burstiness exists in every time-scale while with traditional models like Poisson, the burstiness disappears in the long time intervals (order of the time constant of model). Especially the index of dispersion converges to some fixed value, while with the experimental data it seems to grow infinitely.

3.3 Self-Similarity

Fractals and self-similarity have been a hot topic during the last years. Generally self-similarity means that when a discrete time or continuous time process is scaled in time similar patterns can be seen. The process in larger scale is a copy of itself in smaller timescales. Another term often used with this kind of processes is long-range dependence which generally generally is not about same phenomenan but means processes whose autocorrelation function has long tail behavior [WTSW97].

3.3.1 Definition of Self-Similarity

Let $X = (X_t : t = 1, 2, 3, \dots)$ be wide-sense stationary stochastic process with mean μ, variance of σ^2 and $r(k)$ as its autocorrelation function. In [CB97] authors consider zero-mean process and have somewhat simpler formulas. It is also assumed that autocorrelation function $r(k)$ has form

$$r(k) \sim k^{-\beta} L_1(k) \qquad \text{as } k \to \infty, \tag{1}$$

where $0 < \beta < 1$ and $L_1(x)$ is any function for which $\lim_{t\to\infty} L_1(tx)/L_1(t) = 1$ for all $x > 0$ (ie. at infinity L_1 is grows slowly).

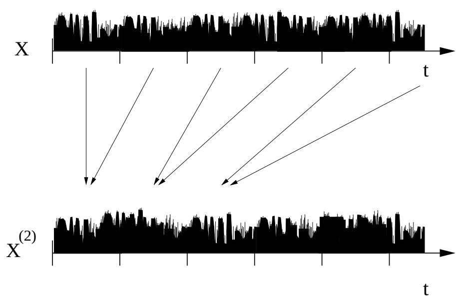

Figure 3: Aggregated process of size 2.

Define a m-aggregated series $X^{(m)} = (X_k^{(m)} : k = 1, 2, 3, \dots)$ by summing the original series X over non-overlapping blocks of size m:

$$X_k^{(m)} = \frac{1}{m}(X_{km-m+1} + \dots + X_{km}) \qquad k = 1, 2, 3, \dots \tag{2}$$

And denote by $r^{(m)}(k)$ the autocorrelation function of aggregated process.

The time series X is said to be (exactly) *H-self-similar* if for all $m > 0$ it holds

$$X_k \overset{d}{=} m^{-H} \sum_{i=(k-1)m+1}^{km} X_i, \qquad \forall\, m \in N. \tag{3}$$

The parameter $H = 1 - \beta/2$ is called Hurst-parameter. Equation (3) means that a scaled versions of same process are distributionally self-similar.

Moreover the process X is said to be *(exactly second-order) self-similar* with parameter H if for all $m = 1, 2, 3, \dots$ if, $m^{-H} \sum_{i=(k-1)m+1}^{km} X_k$ has the same variance and autocorrelation function as X. In terms of aggregated function this means that for all $m = 1, 2, 3, \dots$ [WTLW95]

$$\begin{aligned} \mathrm{V}\left[X^{(m)}\right] &= \sigma^2 m^{-\beta} && (4) \\ r^{(m)} &= r(k) = 1/2\sigma^2(|k|^{2-\beta}) \qquad k = 0, 1, 2, \dots, && (5) \end{aligned}$$

where $\sigma^2(f)$ denotes the second central difference operator applied to function f:

$$\sigma^2(f(k)) = f(k+1) - 2f(k) + f(k-1).$$

Furthermore, the process X is called *(asymptotically second-order) self-similar* if

$$r^{(m)}(k) \to 1/2\sigma^2(|k|^{2-\beta}) \tag{6}$$

3.3.2 Hurst's law and R/S-statistic

Let X_k be a set of observations ($k = 1, 2, \dots, n$ of some process, $\overline{X}(n)$ the sample mean and $S^2(n)$ th sample variance. Also let $W_k = (X_1 + X_2 + \dots + X_k) - k * \overline{X}(n)$ ie. X_k is difference of first k samples from observed mean. Now *R/S-statistic* is given by

$$R(n)/S(n) = 1/S(n) \cdot [max(0, W_1, W_2, \dots, W_n) - min(0, W_1, W_2, \dots, W_n)]\,. \tag{7}$$

Hurst has found that many time series in nature follow the Hurst's law:

$$\mathrm{E}\,[R(n)/S(n)] \sim cn^H \qquad \text{as } n \to \infty, \tag{8}$$

where $H > 0.5$ is Hurst parameter and $c > 0$ is a finite constant.

For a short range-model it holds [WTLW95]

$$\mathrm{E}\,[R(n)/S(n)] \sim dn^{0.5} \qquad \text{as } n \to \infty, \tag{9}$$

where d is a finite constant. The difference between (8) and (9) is called the Hurst effect.

3.3.3 Tests for Self-Similarity

Before using self-simular models one must justify the assumption of self-similar nature. In [CB97] Crovella and Bestavros use the following tests:

1. *variance-time plot*: the estimated variance of $X^{(m)}$ is plotted against m on log-log scale. A straight line with $(-\beta)$-slope indicates self-similarity.
2. *R/S plot (rescaled adjusted range)*: This method relies on formulas (8) and (9) where the estimator of Hurst parameter H is obtained by observing the slope of plotted sample values ($\log(R(t_i, n)/S(t_i, n))$). For details see [WTLW95].

3. *periodogram method*: the slope of the observed power spectrum is plotted on log-log scale. The slope of the curve should be close to a straight line near origin. In fact the slope should be $\beta - 1$ and this is estimated from the plot [CB97].
4. *Whittle estimator*: Unlike the other three methods which were more or less heuristics, this estimator is clearly defined. The drawback with Whittle estimator is the assumption of underlying process. For details see [WTLW95].

The same tests are also used in [WTLW95]

3.4 Heavy-Tailed Distributions

A distribution of a random variable is said to be heavy-tailed if

$$\mathrm{P}\{X > x\} \sim x^{-\alpha}, \qquad x \to \infty, 0 < \alpha < 2.$$

For example for the exponential distribution holds

$$\mathrm{P}\{X > x\} = e^{-\lambda x},$$

and it does not have this feature. The simplest heavy-tailed distribution is the Pareto-distribution for which is holds

$$\mathrm{P}\{X > x\} = (k/x)^{\alpha}, \qquad \alpha, k > 0,\ x \geq k.$$

The parameter k is the smallest possible value of the random variable.

If the parameters α is smaller than or equal to 2, the distribution has infinite variance. Furthermore if $\alpha \leq 1$ it has also infinite mean. So as the value of α decreases, the more of the probability mass is located in the tail of the distribution [PF95] [CB97].

3.4.1 log-log complementary plots

When there is a doubt that some distribution possesses heavy-tailed characterictics, it is common practice to employ the log-log complementart plots (LLCD). The cumulative distribution function $F(x)$ and complementary distribution function $\bar{F}(x)$ are

$$F(x) = \mathrm{P}\{X \leq x\} \qquad \bar{F}(x) = 1 - F(x) = \mathrm{P}\{X \geq x\}.$$

The log-log complementary plots are simply plots of $\bar{F}(x)$ on log-log axis. If the distribution is heavy-tailed, then

$$\frac{d\ \log \bar{F}(x)}{d\ \log x} = -\alpha, \qquad x > \theta$$

for some θ. In practice heavy-tailed behavior of distribution can be simple checked by plotting the estimated distribution on log-log axis and observing possible linear behavior from the figure. The plot should be linear over a three orders of magnitude or more [CB97]. In figure 4 is an example plot of Pareto distribution.

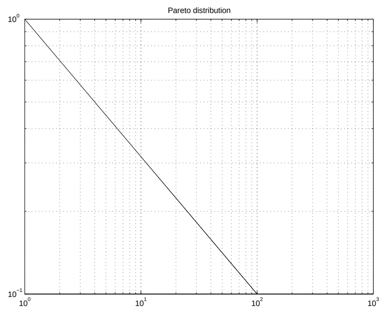

Figure 4: Pareto distribution with $k = 1$ and $\alpha = 0.5$.

3.5 Self-similarity and long-range dependence

The theory of self-similarity and long-range dependencies is useful when studying the traffic at internet. In fact the self-similarity seems to be a fundamental phenomena in internet traffic (see eg. [LTW94]): the number of packets seen in network during a short time period seems to behave similarily than the number of packets during much longer periods of time.

From self-similarity process it follows [LTW94]:

1. *slowly decaying variances*: the variance of the sample mean decreases more slowly than the reciprocal of the sample size: $\mathrm{V}[X^{(m)} \sim m^{-\beta}$ as $m \to \infty$ and $0 < \beta < 1$
2. *long-range dependence*: Also called the Joseph effect by Mandelbrot. the autocorrelation function decays hyperbolically rather than exponentially and $\sum_k r(k) = \infty$.
3. $1/f$*-noise*: the spectral density $f()$ follows a power-law near origin i.e. $f(\lambda) \sim \lambda^{-\gamma}$ as $\lambda \to 0$ and $0 < \gamma < 1$ where $\gamma = 1 - \beta$

4 Results briefly

The traffic seen in WANs is consistent with self-similarity, originally reported in [PF95] and later studied and confirmed in other reports like [FGWK98] [WTSW97].

In [WTSW97] authors also present the results that aggregated traffic from simple ON/OFF-sources with certain characteristics exhibits a self-similar nature.

4.1 WWW-traffic

In [CB97] Crovella and Bestavros have shown that World Wide Web transfer have self-similar characteristics. Firstly transmission times of files follow heavy-tailed distribution (the sizes of files are heavy-tailed too), and secondly the silence times follow a heavy-tailed distribution as well (user thinking time). The comparison of distributions show that ON-time distribution is heavier tailed than OFF-time.

4.2 Generating self-similar traffic

Several models have been proposed to generate traffic which is self-similar. In [PF95] contains a very simple $M/G/\infty$-model while [WTSW97] contains rather long discussion about how self-similar traffic can be generated by using simple ON/OFF sources. Also exact methods exist for generating self-similar traffic from fractional Gaussian noise (FBM) and fractional ARIMA models but they become infeasible when long simulations are considered [LTW94].

References

[CB97] Mark E. Crovella and Azer Bestavros. Self-similarity in world wide web traffic: Evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, December 1997.

[FGWK98] A. Feldmann, A.C. Gilbert, W. Willinger, and T.G. Kurtz. The changing nature of network traffic: Scaling phenomena. *Computer Communication Review*, April 1998. Also available from http://www.research.att.com/~agilbert/papers.html.

[FL91] H.J. Fowler and W.E. Leland. Local area network traffic characterics with implications for broadband network congestion management. *IEEE Journal on selected areas in communications*, 9(7):1139–1149, September 1991.

[HO88] Joseph L. Hammond and Peter J.P. O'Reilly. *Performance analysis of Local Computer Networks*. Addison-Wesley, 1988.

[HP93] Peter G. Harrison and Naresh M. Patel. *Performance Modelling of Communication Networks and Computer Architectures*. Addison-Wesley, 1993.

[IET] The internet engineering task force. http://www.ietf.org/.

[ITA] The internet traffic archive. http://ita.ee.lbl.gov/.

[LTW94] Will E. Leland, Murad S. Taqqu, and Daniel V. Wilson. On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2(1):1–15, Feb 1994.

[LW91] W.E. Leland and D.V. Wilson. High time-resolution measurement and analysis of lan traffic: Implications for lan interconnections. In *IEEE Infocom'91*, pages 1360–1366, Bal Harbour, FL, April 1991.

[LWTW94] W. E. Leland, W. Willinger, M. S. Taqqu, and D. V. Wilson. Statistical analysis and stochastic modeling of self-similar datatraffic. In J. Labetoulle and J.W. Roberts, editors, *ITC 14*, pages 319–328, Antibes Juan-les-Pins, France, June 1994. Elsevier Science B.V.

[Pax99] Vern Paxson. End-to-end internet packet dynamics. *IEEE/ACM Transactions on Networking*, 7(3):277–292, June 1999.

[PF95] Vern Paxson and Sally Floyd. Wide area traffic: The failure of poisson modeling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, June 1995.

[WTLW95] Walter Willinger, Murad S. Taqqu, Will E. Leland, and Daniel V. Wilson. Self-similarity in high-speed packet traffic: Analysis and modeling of ethernet traffic measurements. *Statistical Science*, 10(1):67–85, 1995.

[WTSW97] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson. Self-similarity through high-variability: Statistical analysis of ethernet lan traffic at the source level. *IEEE/ACM Transactions on Networking*, 5(1):71–86, February 1997.