The Self-Similarity of TCP Traffic: first results

Sam Manthorpe[‡] Ilkka Norros[§] Jean-Yves Le Boudec[¶]

November 9, 1995

Abstract

This paper presents on-going work to show how TCP contributes to the self-similar properties observed in data traffic. Preliminary results are presented from simulations of TCP that show how pseudo-self-similarity can be observed over up to 4 timescales in a system consisting of multiple traffic sources, the TCP flow-control algorithm, a shared-medium access network and an ATM network. The time-variance gradient method is used to estimate the Hurst parameter of the cell arrival process for different scenarios and it is shown that the TCP algorithm introduces pseudo-self-similar behaviour in the timescale between 100s of μs to seconds, this range being fairly indepedent of parameters such as the mean utilisation of the network, number of sources and the end-to-end delay.

1 Introduction

As has been shown by W.Leland's group at Bellcore, data traffic seen on today's LANs often demonstrates the so-called 'self-similarity' property, where the same arrival rate patterns are observed over a range of different timescales [1], [2].

The problems of dimensioning data networks using traditional mathematical approaches have been highlighted in [3]. Complex interaction between different layers of the protocol stack makes it practically impossible to find realistic models that can be used with traditional network dimensioning techniques. However, the apparently complex behaviour of data traffic can be simply described in terms of self-similarity. An important open question is: *does the fact that data traffic may be described by a simple mathematical concept mean that simple mathematical solutions may be found for dimensioning data networks?*

This paper does not attempt to answer this question, but merely to gain some insight into what causes the self-similarity of data traffic in the hope that a deeper understanding of the mechanics of the traffic may lead to some useful dimensioning techniques.

In this paper we examine the network in terms of timescales. Network protocols generally have their own 'timescale-layers' in which they influence the traffic; an example is given in figure 1, which shows not only the clearly defined protocol layers, but vaguer concepts such as the daily human cycle.

In this paper we examine the following layers: the ATM layer, packet layer (e.g. Ethernet, AAL5 frame, both of which are equivalent in terms of timescales) and the TCP layer. Self-similar traffic shows, by definition, the same behaviour over *all* timescales (up to infinity). When dealing with a real network, this is not an applicable concept. For the purposes of this discussion, we use the term pseudo-self-similar to describe self-similarity over a limited timescale [4]. Self-similar behaviour in a process can be observed using the so-called 'variance-time' plots. In our case, the arriving traffic stream is divided into equal sized windows, each w slots long. The number of cell arrivals in the window is divided to give the mean cell arrival rate within a window, ρ. The log of the variance is then plotted against the log of w; this is the variance-time plot. A full definition may be found in [5]. The variance-time curve for pseudo-self-similar processes is straight over the time scales where the process exhibits self-similar behaviour.

[‡]Laboratoire de Réseaux de Communication, EPFL, Switzerland
[§]VTT Research, Finland
[¶]Laboratoire de Réseaux de Communication, EPFL, Switzerland

Timescale	*Range*	*units*
yearly cycle	$10^5..10^8$	weeks
weekly cycle	10^5	days
daily human cycle	10^4	hours
human application use	$10^1..10^3$	minutes-hours
transport layer	$10^{-3}..10^0$	
network layer	$10^{-6}..10^{-3}$	
link layer	10^{-6}	

Figure 1: Timescales of a typical protocol stack

Such variance-time plots have a certain innaccuracy in that the estimation of the variance is biased. If the samples are independant, then the variance of a set of samples X_i can be estimated using

$$var(X_i) = \frac{\sum_{i=1}^{i=N} (X_i - E(X_i))^2}{N-1} \tag{1}$$

However, in our simulations, the samples are not independant, since they are obtained by measuring in successive windows on the same traffic stream. Work is ongoing to develop a method for correcting these errors. However, the errors will not make a big difference to the conclusions drawn from the results presented here, since they are multiplicative errors of order of magnitude 1, which will not change much on the logarithmic scales used.

Results will be presented that show that the TCP layer exhibits pseuedo-self-similar behaviour over several timescales. We do not consider higher layers, and ignore the 'human-influence' on the traffic, such as connection-retry effects and the daily usage cycle.

The following section briefly describes the simulation model used and section 3 presents the results and uses variance-time plots to show the self-similarity of the traffic and to estimate the Hurst parameters under a variety of scenarios. Section 3 concludes the paper with a discussion of the reasons for the self-similarity of data traffic.

2 Simulation Model

2.1 Network model

To study the behaviour of TCP/IP traffic in a broadband heterogeneous network, a simulation model consisting of a DQDB MAN and ATM was adopted. The model network is shown in figure 2. We consider the case where data traffic is generated on a DQDB MAN and routed to an Interworking Unit (IWU), which transmits the traffic on the ATM network [1]. The DQDB network is modelled as a processor sharing queueing system [6], which models the segment interleaving resulting from the MAC layer. Once at the IWU, the traffic is routed, via the CLON, to the destination network. The IWU works in cell-to-slot mode; that is, it does not reassemble frames, but converts segments directly into ATM cells. Before being sent on the ATM network, the traffic stream is shaped. The rest of the ATM network and the destination network were not modelled in detail, but were represented by a constant delay. Exact values for the model values that are not variable in these studies are presented below.

2.1.1 Traffic sources

In these studies we adopt an on-off source to isolate the effect of TCP on the traffic. The on-off model creates long bursts of data that are typical of applications such as file transfer over data networks. The

[1]The reason for this particular model is historical, since the simulator used was from the RACE 2032 COMBINE project; the authors believe that the choice of technologies (DQDB and ATM) in this model has little bearing on the conclusions of the paper

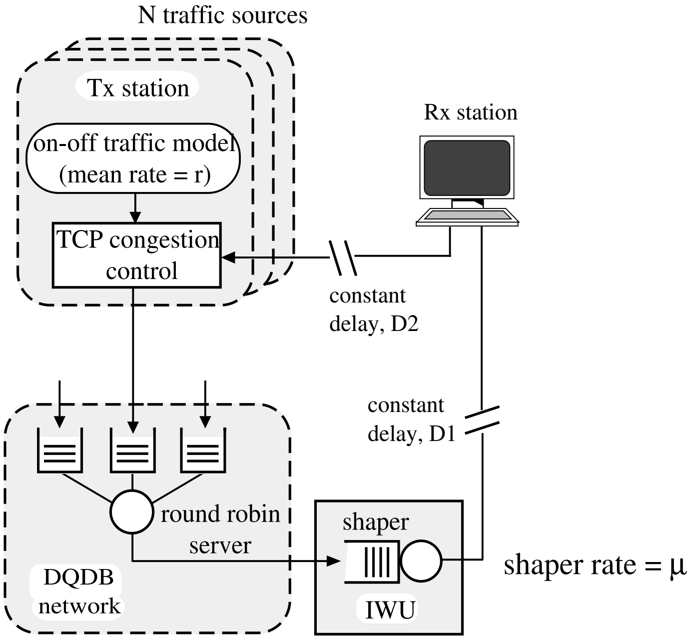

Figure 2: Network simulation model

file length distribution was approximated as a geometrical distribution, with the mean file length taken from measurements of file lengths on an EPFL file server. The geometrical distribution would not be a good choice if we were to study the behaviour at longer timescales than that of TCP, since the measured distribution of filelengths was actually heayy tailed; see A. The silence periods between file transmissions were also assumed to be geometrically distributed, with the mean silence length being chosen to give the mean traffic rate desired. Recent results have shown this assumption to be justified [7].

The TCP congestion and flow control algorithm was simulated at the transport layer of each source. The implementation adopted was that described by Van Jacobsen in [8] with a window based mechanism, the window size being controlled by the slow-start and congestion avoidance algorithms together. We assumed that all the receiving stations are sufficiently fast that they never have to slow the sources down using window advertising. For full details, the reader is referred to [9] and [8]. Unless otherwise stated, the number of traffic sources, N, is equal to 20.

2.1.2 Simulator

The simulations were run using software developed in the RACE 2032 COMBINE project, based upon the ATLAS-II library of C functions [10] for network simulation. Simulation model times varied between 5000 and 50000 seconds, depending on the time required to achieve sufficiently small confidence intervals for the packet loss probability in the shaper.

2.1.3 Simulation parameters

Simulation parameters that never changed for all the simulations run are as follows:

Parameter	*Value*
DQDB bus bitrate	100 Mbits/s
DQDB MAC access queues	1000 packets (effectively infinite)
Peak rate during TCP connection	10 Mbits/s
Maximum TCP window size	75 koctets
Maximum TCP timeout value	10 seconds
b used for TCP	2 (unitless)
Exponential timeout backoff	2 (unitless)

Note that the DQDB MAC access queues are effectively infinite, in that they never overflows, their service rates being bigger than the peak rate generated by a workstation. Other important parameters had a standard configuration, shown below, with one or more of the parameter being varied for the purposes of a particular study. Unless otherwise stated, the parameters used are those shown below.

Parameter	*Value*
Mean file length	34083 octets
TCP segment sizes	1500 octets
Number of sources, N	20
Mean shaper load	0.50

3 The self-similarity of data traffic

We start this section by determining over what timescales TCP influences the self-similarity of the data traffic in this model. In figure 3, we compare the time-variance plots for the standard simulation scenario, one with TCP and one without. Without TCP, there is a period of pseudo-self-similarity between 10 μs and 10 ms, with a Hurst parameter of 0.92. Timescales of greater than 10 ms see a gradient of 1, which is not self-similar.

Introducing TCP introduces a new part to the curve, between 100 μs and 1 second, which exhibits self-similarity with a Hurst parameter of 0.77. On timescales lower than 100 μ and greater than seconds, the curve follows very closely the behaviour of the simulation without TCP. This quite nicely shows over which timescales TCP has an influence on the traffic.

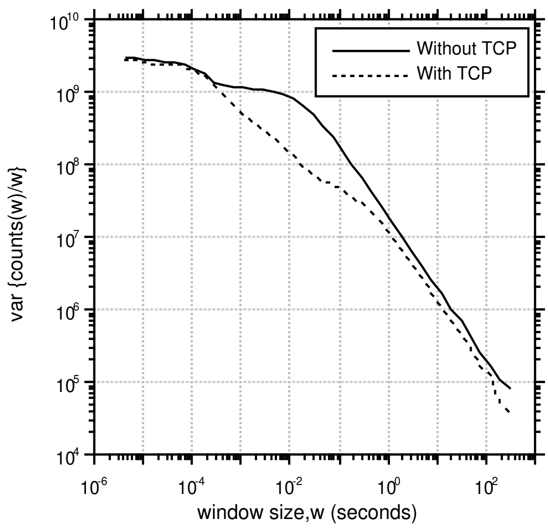

Figure 3: Influence of end-to-end delay on variance-time plots

In the following we see if varying any of the important simulation parameters (shaper load, mean burst length, number of sources, packet size and end-to-end delay) has an impact on this conclusion.

3.1 Influence of mean shaper load on self-similarity

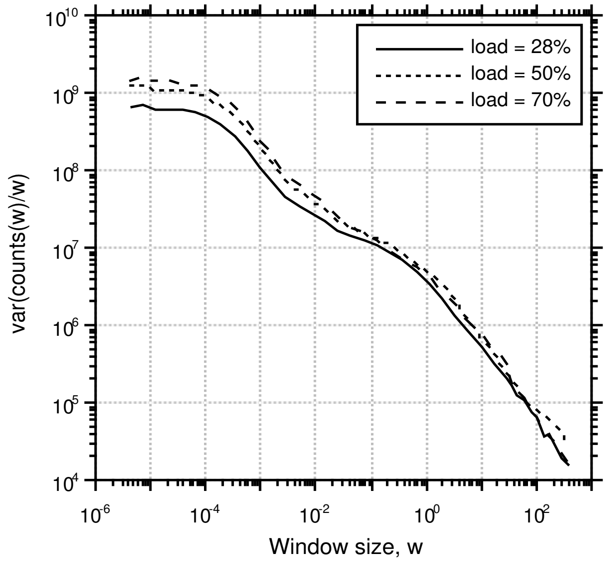

Figure 4: Time-variance curves for different mean shaper utilisations

Figure 4 shows the variance-time plots for the the simulated network, for 3 different loads. We define the load to be the mean offered traffic (see appendix A) divided by the bitrate of the traffic shaper. These results were obtained by varying the shaper rate. It can be seen that varying the load does not have much effect on the curves; each shows roughly pseudo- self-similar behaviour in the timescales between 100 μ s and seconds. We estimated the Hurst parameters to be:

load	*Hurst parameter*
.28	0.739
.50	0.701
.70	0.689

3.2 Influence of mean burst lengths (filelengths) on self-similarity

Figure 5 shows how changing the mean filelength (the mean burst length in the on-off model) influences the variance-time plots. The mean filelength is experessed in terms of a constant k:

$$filelength = 3.4083 * 10^k \tag{2}$$

In all cases, the plot is straight for window sizes 10^{-3} to 10^{-2} with a gradient of approximately 0.775 (Hurst parameter of 0.613). For longer window sizes, the gradient clearly depends on the mean burst length. As we saw in figure 3, the TCP algorithm had no influence on time scales longer than seconds. If the mean file length is shorter than seconds, (k=2 and k=3), TCP is rarely active long enough to have any influence on the traffic. For longer timescales, (k=5 and k=6), we see that the curve is dominated by the mean filelength parameter. There is self-similarity over at least two orders of magnitude of time scale, but the longer the burst length, the more the bursts seem to interfere with the straight variance plots caused by TCP. Whether the shape of these curves is due to the on-off model alone, or also to the TCP algorithm is for further study.

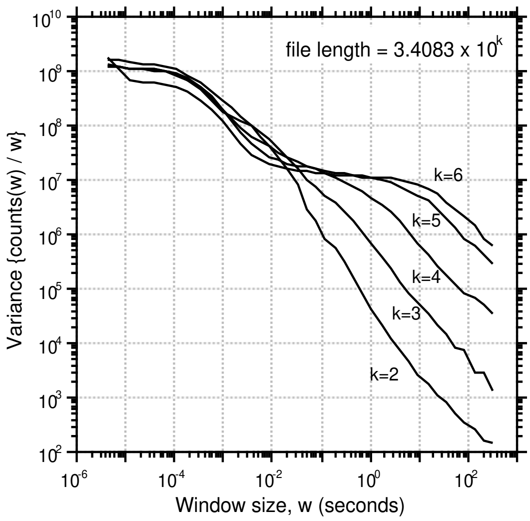

Figure 5: Influence of filelength on variance-time plots

3.3 Influence of number of sources on variance-time curves

The measurements made by Bellcore were made on aggregate traffic and it has been shown that *aggregate* traffic is self-similar. To see if TCP's contribution to data traffic self-similarity is only for an aggregation of sources, the variance curves were measured for different numbers of sources (n=1, 5, 10 and 50). The mean source load was adjusted so that the mean shaper load was always the same, independent to the number of sources. We see that although the curves are slightly displaced from each other, they all have basically the same shape and exhibit self-similar properties between 100 μs and seconds. These properties can be seen in slightly longer timescales as the number of sources increases.

From this we can speculate that TCP contributes to the self-similarity of the traffic whether it be applied to all the traffic as a whole, or individually to each traffic stream.

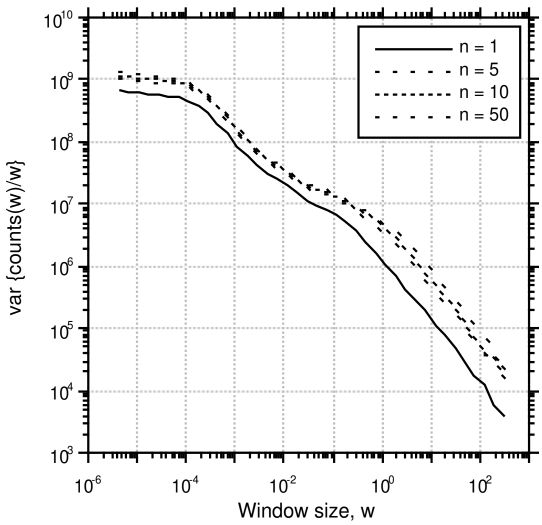

Figure 6: Variance-time plots for different number of sources

3.4 Influence of packet sizes on variance-time curves

Figure 7 shows how the size of the packet influences the variance time plots. We observe a slight change in the Hurst parameter as a function of the packet size. Interestingly, decreasing the packet size seems to introduce a slight knee in the otherwise straight line corresponding to TCPs timescale. For packet sizes of 1000 bytes, we see more or less the same result as for 1500 (in the rest of the simulations). Large packets of 9000 octets reduce slightly the timescale over which we see pseudo-slef-similarity. This seems to be a similar effect to that observed when increasing the file lengths; a low-gradient part of the curve interferes with the straight line of TCP. This phenomena would appear to be related to bursts. When the file lengths were increased, the bursts were file transmissions and here, when we increase the packet size, the bursts are bursts of cells resulting from packet transmissions. These ideas are nothing more than conjecture for the moment and have to mathematical justification.

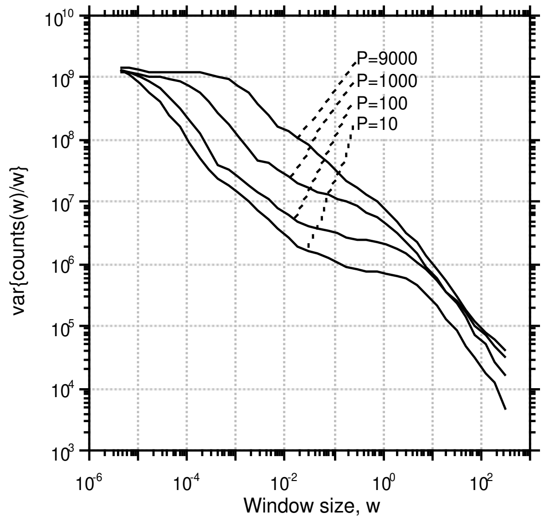

Figure 7: Influence of the packet size on the variance-time plots

3.5 Influence of end-to-end delay on variance-time curves

In figure 8 we see that varying the end-to-end delay does not significantly change the shape of the curves, but results in a displacement between them.

4 Conclusions

This work has shown how the TCP algorithm contributes to the pseudo-self-similarity of data traffic. We have seen that, in this scenario, TCP introduces self-similairty in the timescale between 100 μ s and seconds. This observation is reasonably independant of most of the parameters adjusted: shaper load, number of traffic sources and end-to-end delay. The mean file length (burst length of the on-off model) did influence the variance-time plots for the arriving aggregate traffic, but the interaction between the burst sent to the TCP algorithm and the behaviour of the TCP algorithm itself, in terms of the variance-time plots, is not yet understood.

From a general point of view, TCP can be seen as a specific way of organizing the *contention for transmission resources*. Since contention for resources is a generic aspect which appears in telecommunications in many contexts and at many time scales, it can be expected that contention in fact causes qualitatively similar effects as those considered in this paper also at other layers and time scales. This idea remains for further study.

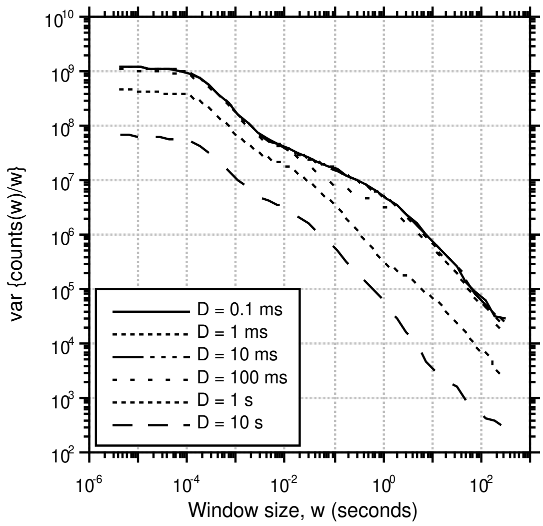

Figure 8: Influence of end-to-end delay on variance-time plots

Other parameters that may have a bearing on these results are those of the TCP algorithm itself and its associated operating system functions, in particular socket buffers [11]. The invesstigation of such parameters is for further study.

These results have shown that the self-simiilarity of data traffic is not due to one phenomena alone, but is a result of different processes at different timescales. Further work to identify these processes would allow one to decide if models for pseudo-self-similar traffic [12, 13, 14] may be used for network dimensioning.

5 Acknowledgements

Many thanks to Stephan Robert and Jean-Yves Le Boudec for many useful discussions. This work has been supported by the Swiss PTT Telecom.

A File length distribution

Figure 9 shows the distribution of file lengths measured on an EPFL file server. As can be seen, the distribution is heavy tailed and a geometric distribution would not be ideal if we wanted to study the behaviour at timescales corresponding to file transmissions. Here we use the geometric distribution for simplicity.

References

[1] W.Leland and D.Wilson, "High time-resolution measurement and analysis of lan traffic: Implications for lan interconnection," in *Proceedings of Infocom 91*, 1991.

[2] H.J.Fowler and W.E.Leland, "Local area network traffic characteristics, with implications for broadband network congestion manegement," *IEEE Journal on Selected Areas in Communications*, vol. 9, no. 7, pp. 1139–1149, 1991.

[3] S. Manthorpe, "The danger of data traffic models," Tech. Rep. 25, COST 242, April 1995.

[4] S. Robert, "Stochastic processes for self-similar traffic," Tech. Rep. 46, COST 242, September 1995. Contribution to COST242 meeting in Bratislava.

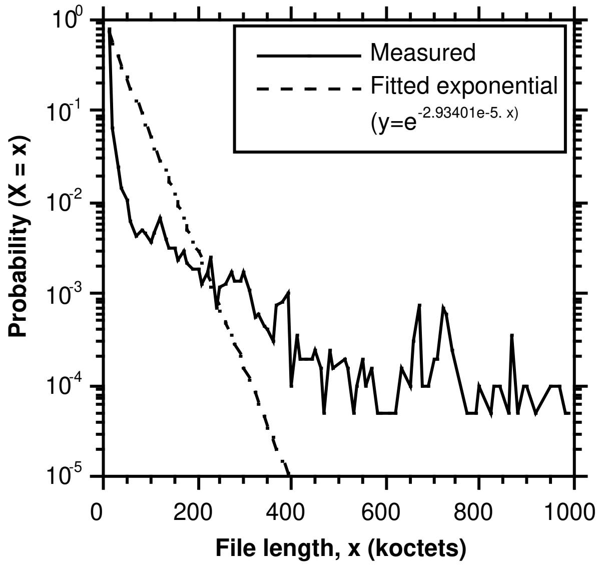

Figure 9: Distribution of file lengths on an EPFL file server

[5] W.E.Leland, M.S.Taqqu, W.willinger, and D.Wilson, "On the self-similar nature of ethernet traffic," tech. rep., Bellcore, March 1993. see Bellcore FTP server.

[6] S. Aalto, I. Norros, J. Virtamo, K. Kvols, and S. Manthorpe, "Performance aspects of streaming and message modes of interworking," in *Proceedings of the Interworking 92 Conference Bern*, nov 1992.

[7] V. Paxson and S. Floyd, "Wide area traffic: The failure of poisson modeling," *IEEE/ACM Transaction on Networking*, vol. 3, pp. 226–244, June 1995.

[8] V. Jacobsen and M.Karels, "Congestion avoidance and control," in *Proceedings of SIGCOM '88, Stanford*, August 1988.

[9] D. E. Comer, *Internetworking with TCP/IP Volume 1 - Principles, Protocols and Architecture.* 0-13-474321-0: Prentice-Hall Internation, 2 ed., 1991.

[10] M. Ruemekasten and A. Calcagno, *ATLAS Reference Manual.* University of Paderborn, Dept. Mathematics and Computer Science, 5 ed., december 1991.

[11] G. Wright and R. Stephens, *TCP/IP Illustrated: The implementation*, vol. 2. 0-201-63354-X: Addison-Wesley, 1995.

[12] I. Norros, "On the use of fractional brownian motion in the theory of connectionless networks," tech. rep., COST 242, September 1994.

[13] S. Robert, "Can self-similar traffic be modelled by markovian processes?," tech. rep., COST 242, May 1995. Contribution to COST242 meeting in Stockholm.

[14] A.T.Andersen, A.Jensen, and B.F.Nielsen, "Modelling of apparently self-similar packet arrival processes with markovian arrival processes (map)," Tech. Rep. TD(95)05, COST 242, January 1995. Cambridge meeting.