On Modeling and Shaping Self-Similar ATM Traffic

Sándor Molnár and Attila Vidács

[a] High Speed Networks Laboratory, Dept. of Telecommunications and Telematics, Technical University of Budapest, H–1111, Sztoczek u. 2., Budapest, Hungary
E-mail: molnar@bme-tel.ttt.bme.hu, vidacs@ttt-atm.ttt.bme.hu

A measurement study of ATM WAN traffic has been carried out and it is shown that the recorded data exhibit self-similar features. The conclusions are supported by a comprehensive analysis examining four statistical methods. Our results are validating one of the most striking findings of the present teletraffic research: a broad range of packet network traffic has fractal-like behaviour. We also investigate three self-similar traffic models for the measured traffic together with their performance analysis. Finally, we discuss issues on shaping and simulated queueing performance of the ATM traffic and show the strong robustness of self-similar properties identified.

1. INTRODUCTION

In the last decade a number of extensive studies of high resolution traffic measurements from a wide range of packet traffic networks have been reported [1,6,7,10,11,14,22]. The most important finding of these studies is the identified fractal-like behaviour implying the so called *long-range dependence* and *self-similarity* properties. As a result of intensive research at Bellcore a series of papers reported these findings in Ethernet LAN [7,9,11,12]. The comprehensive study of Leland's group with the conclusion that this traffic is self-similar was published in detail in [11]. The study of Duffy et al. [6] revealed the self-similarity traffic property in common-channel signalling network. Meier-Hellstern et al. [14] found that the Pareto distribution with infinite variance is applicable for characterizing the D-channel traffic in N-ISDN. Paxson et al. [20,22] reported the self-similar features of TCP traffic. The fractal properties also appeared in the analysis of video traffic (see the work of Garett et al. [10], and Beran et al. [1]). These traffic measurements demonstrated a tremendous traffic burstiness at several time scales with the properties of self-similarity. In fact, there are various open issues and living discussions in teletraffic research about these findings and still there is no definite conclusion about that the observed traffic are self-similar *in nature* or these properties can be explained by non-stationarity of the processes in larger time scales [5].

The modeling of self-similar traffic appeared as an emerging and challenging field of the present teletraffic research. It seems that there are different promising approaches to capture this complex fractal-like behaviour. Norros [17,19] used a Gaussian self-similar process known as the Fractional Brownian Motion. Willinger et al. [27] applied the superposition of on/off sources with heavy-tailed on and off periods. Erramilli et al.

[8,25] studied different chaotic maps. Another hot topic of the present research is to investigate the impacts of different network control mechanisms (e.g., shaping [16,23]) on self-similarity properties and the dimensioning aspects of fractal traffic.

The purpose of this paper is threefold. First, we present an analysis study based on measurements taken from a WAN ATM network with four methods testing for self-similarity. We show that the self-similar properties are present in this traffic. Our conclusions do not provide new results but validating these findings by the analysis of actual ATM traffic in a real WAN network. Second, we investigate the mentioned three promising self-similar modeling approaches to capture the observed properties, namely, the *fractional Brownian traffic* [17,19], *superposed on/off sources* [13,26] and *chaotic maps* [8,25]. Third, we give analysis results of *shaped self-similar traffic*. We investigate the practically important question: can self-similarity be removed from the traffic by shaping or not? Queueing analysis of the shaped and original traffic are also presented about the robustness of the self-similarity property.

2. ATM TRAFFIC MEASUREMENTS

2.1. The FUNET network

The measurements were made on the FUNET WAN ATM network. "FUNET" stands for "Finnish University and Research Network" which network is built on Telecom Finland's ATM-network. All the Nordic national networks (FUNET, DENnet, ISnet, SUNET and UNINETT) are connected to the Nordic Backbone Network (NORDUnet) which has a connection point in Stockholm, Sweden. NORDUnet has connections to the US backbones, the European backbones and to networks in central and eastern Europe.

2.2. Measuring tool and configuration

The measurement was made in the logical center of the whole network at the Center for Scientific Computing (CSC) in Espoo, Otaniemi. All the international links start from here, including the main crosslink to Stockholm. Our measurement equipment was inserted between the network and the high-capacity ATM switch situated in Espoo. At this point all the ATM traffic from the FUNET network transported through the switch and the traffic generated at the CSC and transmitted to the rest of the world could be monitored. The measurements were made by an HP Broadband Series Test System equipment.

2.3. The measured data

The measured traffic was the aggregated traffic at the most heavily loaded point of the network. During the measurements, two types of data collections were made. In the first scenario the measured data was the time stamp (with resolution 0.01μs) of the arrival time instant for every single cell on the 155Mbps link. Because of the upper limit for the number of captured cells each measured data file contains 131,072 time stamps only, which was about 3–6 seconds according to the network load. For the long-term analysis longer measurement periods were needed, so in the second measurement scenario the recorded data was the number of cells received in one second intervals. In this case the time interval of the observation could take several minutes long. A summary description of these data sets is given in Table 1.

Table 1
Data sets characteristics and estimated Hurst parameters

Filename	#cells	Length (sec)	Rate (Mbps)	$\hat{H}_{idc}$	$\hat{H}_{var}$	$\hat{H}_{rs}$	$\hat{H}_{per}$
FUNET1	131,072	3.9	14	0.69	0.69	0.68	0.68
FUNET2	131,072	5.1	11	0.67	0.67	0.67	0.73
FUNET3	131,072	4.4	13	0.66	0.66	0.68	0.68
FUNET4	131,072	6.4	9	0.72	0.72	0.74	0.78
FUNETSTA.T3	14,807,546	425	15	0.70	0.70	0.82	0.94
FUNETSTA.T4	43,768,430	1964	9	0.67	0.67	0.79	0.90

The files FUNET1–3 contain traffic data captured from the incoming traffic, and the FUNET4 measurement was made on the outgoing link. In case of the last two measurements in Table 1, the registered data was the number of cells received in every second on the incoming link.

3. THE SELF-SIMILARITY PHENOMENA AND ITS TESTING

In the following, we summarize the definitions and properties related to self-similarity. Consider a weakly stationary discrete time stochastic process $X = (X_1, X_2, X_3, \ldots)$ with constant mean, finite variance and autocorrelation function $r(k)$. Further, let $X^{(m)}$ denote a new time series obtained by averaging X over non-overlapping blocks of size m. That is, for $k = 1, 2, 3 \ldots$ the time series $X_k^{(m)} = (1/m)(X_{km-m+1} + \cdots + X_{km})$ and let $r^{(m)}$ denote the autocorrelation function of $X^{(m)}$. The process X is called *strictly self-similar* with self-similarity (or Hurst-) parameter H if $mX^{(m)}$ has the same finite dimensional distributions as $m^H X$ for all $m \geq 1$. (The process X is *exactly second-order self-similar* if $r^{(m)}(k) = r(k)$, $(k \geq 0)$ and *asymptotically second-order self-similar* if $r^{(m)}(k) \to r(k)$ as $k, m \to \infty$.)

3.1. Properties

A self-similar process has several properties:

- *Long-range dependence:* A self-similar process with parameter $H < 0.5$ is long-range dependent. By definition, the process X is called *long-range dependent* if $\sum_{k=1}^{\infty} r(k) = \infty$. Otherwise it is called *short-range dependent* [3]. This non-summability results in a hyperbolically decaying autocorrelation function.

- *1/f-noise:* Long-range dependence manifests itself in the frequency domain as the so called *1/f-noise*, which is the term used to refer to a sharp divergence in the power spectrum near the origin.

- *Slowly decaying variances:* The variance of the arithmetic mean decreases more slowly than the reciprocal of the sample size: $var\{X^{(m)}\} \sim am^{-(2-2H)}$, as $m \to \infty$, where a is a finite positive constant independent of m, and $0.5 < H < 1$.

- *Hurst effect:* The *rescaled-adjusted range statistics* (see next section) is characterized by $E\{R(d)/S(d)\} \sim const \cdot d^H$, as $d \to \infty$ with $0.5 < H < 1$.

Since we are always dealing with finite data sets, it is in principle not possible to check whether by definition a traffic trace is self-similar or not. Instead, we check for different features of self-similarity present in actual packet traffic based on the properties listed above [2,7,11]. The following most popular self-similarity tests are used to capture some of the listed properties:

3.2. Tests

Indices of dispersion A commonly used measure for capturing the variability of traffic over different time scales is provided by the *index of dispersion for counts* (IDC) [4], i.e., $IDC(L) = var\{\sum_{j=1}^{L} X_j\}/E\{\sum_{j=1}^{L} X_j\}$. Self-similar processes produce a monotonically increasing IDC. Plotting $\log IDC(L)$ against $\log L$, this property results in an asymptotic straight line with slope $2H-1$ [11]. A similar dispersion index is also defined based on the interarrival times of the cell stream referred to as *index of dispersion for intervals* (IDI) [4] and given by $IDI(k) = k \cdot var\{\sum_{j=1}^{k} Y_j\}/E^2\{\sum_{j=1}^{k} Y_j\}$ where $\{Y_j\}$ is the series of interarrival times. It can be used to observe different phenomenons more easily, e.g., to check how close a process is to a renewal process.

Variance-time plot This method is based on detecting the slowly decaying variance property of self-similar processes (see Section 3.1). The so-called variance-time plots are obtained by plotting $\log var(X^{(m)})$ against $\log m$ and by fitting a simple least squares line (with slope $2-2H$) through the resulting points, ignoring the small values of m [11]. Checking self-similarity by variance-time plot is mathematically equivalent to testing the IDC curve.

R/S analysis Given an empirical time series $(X_k : k = 1, \ldots, N)$ of length N, the rescaled adjusted range statistics $R(d)/S(d)$ for a number of values d is given by $R(d) = \max\{0, W_1, W_2, \ldots, W_d\} - \min\{0, W_1, W_2, \ldots, W_d\}$, with $W_k = (X_1 + X_2 + \cdots + X_k) - k\overline{X}(d)$, $(k = 1, 2, \ldots, d)$, $S^2(d)$ is the sample variance, and $\overline{X}(d)$ is the sample mean. One computes these R/S samples for logarithmically spaced values of d, and plotting $\log R(d)/S(d)$ vs. $\log d$ results, also known as *pox diagram*. Next, a least squares line is fitted to the points of the R/S plot, where both the R/S samples of the smallest values of d are not considered (because they are dominated by short-range correlations) and those R/S samples of large values of d where the number of samples are less than say 5 (because they are statistically insignificant). The slope of the regression line for these R/S samples is an estimate for the Hurst parameter H [11].

Periodogram-based analysis Let $I(\cdot)$ denote the sample periodogram (i.e., estimated power spectrum using a Fourier transform) defined by $I(\lambda) = (1/2\pi N)|\sum_{j=1}^{N} X_j e^{ij\lambda}|^2$, $\lambda \in [0, \pi)$. As mentioned in the properties list, the spectral density of self-similar processes obeys a power law near the origin. Thus, the first idea to determine the Hurst parameter H is simply to plot the periodogram in a log-log grid, and to compute the slope of a regression line which is fitted to a number of low frequencies. This should be an estimate of $1-2H$. Note, that in most of the cases this will lead to a wrong estimate of H since the periodogram is *not appropriate* to estimate accurately the spectral density. Thus, more sophisticated methods have to be applied to obtain useful estimates of H [11]. However, this method can reveal the power spectrum near the origin.

4. ANALYSIS

IDC and IDI plots Figure 1 shows the IDC curve corresponding to the trace FUNET1. The sequence of cell counts in every 100μs interval was analyzed. The IDC curve for the FUNET1 file increases monotonically throughout a time span that covers 3–4 orders of magnitude and shows an asymptotic slope that is strictly different from the horizontal line resulting in an estimate $\hat{H}$ of 0.69. This behaviour is in stark contrast to conventional traffic models such as Poisson processes, where the IDC is constant ($H = 0.5$). The positive correlations present in the FUNET1 data is also indicated by the increasing IDI curve (Figure 2). To verify the result, plot the IDC diagram for the "shuffled" traffic

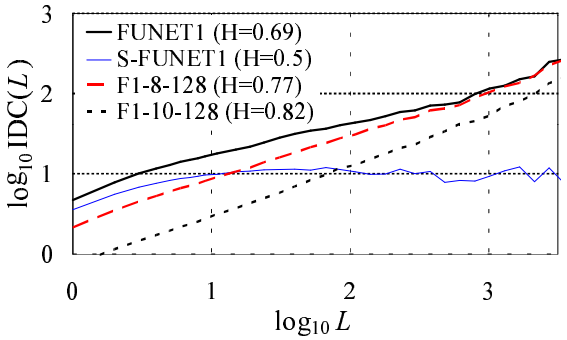

Figure 1. IDC diagrams for the original-, shuffled-, and shaped FUNET1 data.

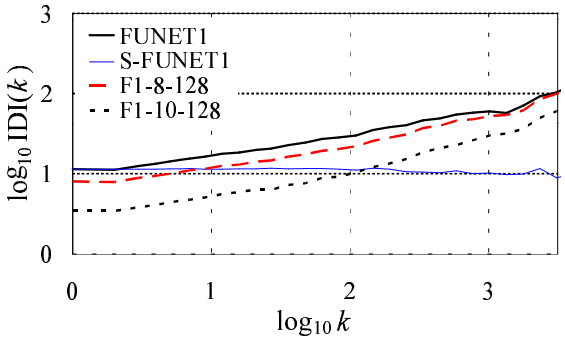

Figure 2. IDI diagrams for the original-, shuffled-, and shaped FUNET1 data.

trace which was generated by mixing the sequence of the cell interarrival times randomly and rebuilding the cell stream (see Figure 1 S-FUNET). The curve starts as in the case of the FUNET1 plot but soon it stops increasing and remains constant for values of $\log L$ greater than 1. This means that removing the long-range dependencies from the data set, the self-similar feature disappears. But the resulting process is still different from the Poisson process. For the values of L lower that 10 the slope of the curve is positive which indicates the presence of short-range dependence. For the values of L greater than 10 the slope of the diagram is zero but the IDC value is about 10. This means that the traffic is bursty but not self-similar. This property is also illustrated by the almost horizontal IDI curve at a value of higher than 10, which equals to the squared coefficient of variation of the cell stream (a good measure of burstiness), plotted in Figure 2. (For a Poisson process the $\log IDI$ is a horizontal line at the value of 0.)

The same analysis was made for all the data sets from the FUNET measurements (see Table 1). As can be seen from the table, the values of $\hat{H}_{idc}$ are pretty much the same for all the data sets. It is remarkable, that in case of the last two data sets the analyzed process was the sequence of cell counts in each seconds instead of 100μs as in the case of the first four sets. In spite of the fact that the time scale was four orders of magnitude higher the Hurst-parameter remained the same.

Variance-time plots The variance-time plot for the FUNET1 and the shuffled data sets can be seen in Figure 3. The estimated values of $\hat{H}_{var}$ are listed in Table 1. Since

the variance-time plots and the IDC diagrams are closely related statistical methods, the results obtained from this method are the same as in the previous subsection.

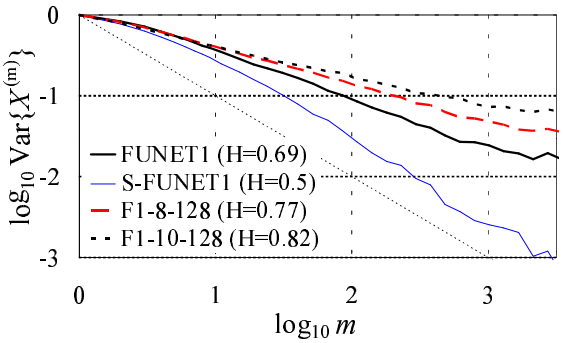

Figure 3. Variance-time plot for the original-, shuffled-, and shaped FUNET1 data.

R/S diagrams Figure 4 shows the R/S plot for the FUNET1 data. The analyzed process was the sequence of cell counts in every 100μs. The estimated value of H for this data set is 0.68, which is nearly the same as the values calculated by the two previous methods.

The same analysis was made for all the FUNET measurement data sets. The resulting values for parameter H are listed in Table 1. For the first four data sets we get what we expected. But in case of the last two data sets, the estimated value of H is higher. The difference between the results and expected values could follow from the fact that these data sets do not contain statistically appropriate amount of data.

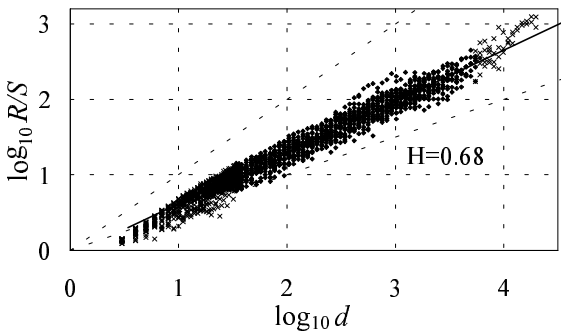

Figure 4. R/S plot for FUNET1.

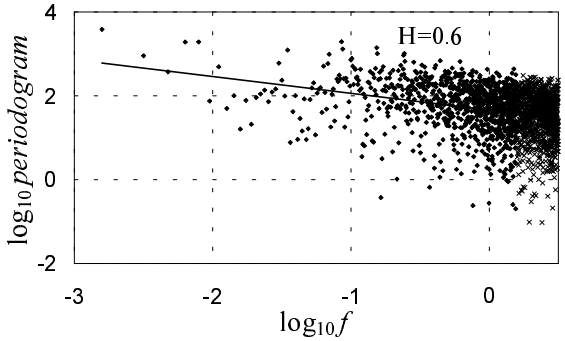

Figure 5. Periodogram plot for FUNET1.

Periodogram plots Figure 5 presents the periodogram plot for the FUNET1 data set, where the analyzed time series was the number of cells in every 1 msec. The slope of the low frequency part is clearly different from zero, and yields $\hat{H} = 0.68$. This result corresponds to the previously calculated values of H.

The analysis was made for all the data sets again, the results are listed in Table 1. For the FUNET1 and FUNET3 data sets the estimated value of H is the same as was

before. But for sets FUNET2 and FUNET4 the values of H are a little bit higher than the previous estimates. As for the last two data sets, the same holds as in the case of the R/S statistics.

Discussion We can conclude, that for the FUNET measurements the estimated value of H is about 0.7 for all the data sets and the measured WAN traffic definitely exhibit the features of self-similarity. To estimate the self-similarity parameter H more precisely, more refined statistical methods with confidence intervals for H and longer measurements with more data are needed.

5. MODELING

In this section three different self-similar traffic generators based on different modeling approaches are presented briefly. These models are all evaluated in the next section.

Fractional Brownian traffic This model was first introduced in [17] and published in [18]. The *fractional Brownian traffic* is a process of the form $A_t = mt + \sqrt{am}Z_t$, where A_t is the amount of traffic in $[0, t)$, and Z_t is a normalized fractional Brownian motion, i.e., a strictly self-similar process. The process has three parameters: the mean input rate m, the variance coefficient a and the Hurst-parameter H of Z_t [19]. We generated the self-similar sample path known as fractional Gaussian noise as presented in [21], and it is referred to as 'FGN' in this paper.

Superposed on/off sources The model was first introduced by Mandelbrot [13] and later extended by Taqqu and Levy [26]. The basic idea of the model is the construction of self-similar processes based on aggregating many simple on/off processes with heavy-tailed on and off periods [26]. This sequence is referred to as the 'ON/OFF' data.

Chaotic maps Erramilli and Singh proposed *chaotic maps* for fractal traffic modeling. The underlying idea is based on a nonlinear map that describes the evolution of a state variable over discrete time [8,25]. We used a simple, two parameter nonlinear chaotic map, the *intermittency map* to generate the 'CHAOS' data.

6. PERFORMANCE EVALUATION OF THE MODELS

The analysis results for the synthetic traffic traces generated by the above mentioned traffic models are presented in this section. The whole methodology and the statistical algorithms used are the same as were in case of the analysis of the FUNET measurement data sets. Figure 6 presents the IDC plots calculated from the ON/OFF, CHAOS and FGN data sets. The resulting diagrams seem to be nearly ideal. For all the plots, the linear parts are clearly visible and it is easy to fit the regression line to the points and to calculate its slope. Table 2 lists the estimate of parameter H for all three cases. (The Hurst-parameter for all the models was set to 0.75.)

The IDI curves (Figure 7) reveal more differences among the models. The IDI curve of the on/off model has similar increasing tendency as in the FUNET1 data showing the presence of positive correlations. This tendency cannot be seen in the CHAOS curve which is horizontal. The reason for this is that the IDI reflects correlations among cell interarrival times and not among the number of cells in a given time interval as in the case

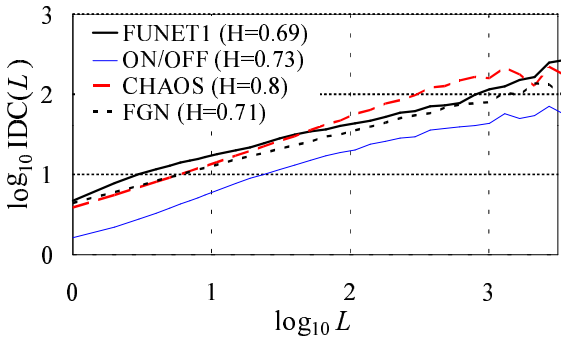

Figure 6. IDC plots for the models.

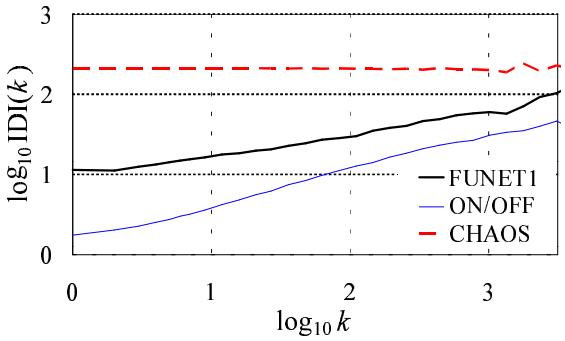

Figure 7. IDI plots for the models.

Table 2
Summary of test results estimating the Hurst-parameter of synthetic traffic traces

Filename	$\hat{H}_{idc}$	$\hat{H}_{var}$	$\hat{H}_{rs}$	$\hat{H}_{per}$
ON/OFF	0.73	0.73	0.75	0.80
CHAOS	0.80	0.80	0.75	0.83
FGN	0.71	0.71	0.71	0.72

of IDC. The CHAOS model generates cell streams with uncorrelated interarrival times and it results in a constant IDI curve.

Figure 8 shows the variance-time plots calculated from the ON/OFF, CHAOS and FGN data sets. As can be seen from Table 2, the results obtained are the same as previously.

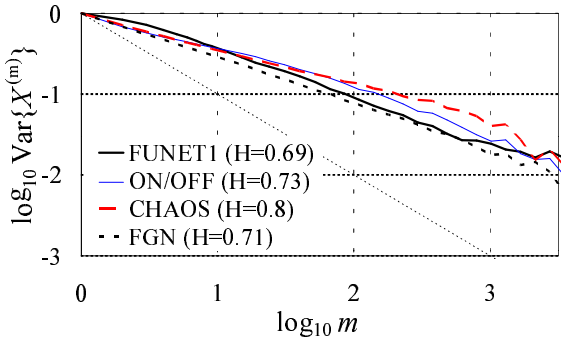

Figure 8. Variance-time plots for the models.

Table 2 lists the calculated values of $\hat{H}$ from the R/S diagrams for all models. For the ON/OFF and CHAOS data sets the calculated value is 0.75, which was the input self-similarity parameter for the model. The value of $\hat{H}$ of the FGN trace is lower than the expected 0.75, but the results is in good agreement with the values estimated previously.

The results obtained from the periodogram plots clearly indicate the presence of long-range dependence in each case (see Table 2). The values of $\hat{H}$ for the ON/OFF and CHAOS sets are higher than before.

6.1. Discussion

As for the traffic model based on the superposition of independent on/off sources, we can say that the generated synthetic traffic trace—as far as the above analyzed statistical properties are concerned—is a good approximation of the modeled traffic. The synthetic traffic trace possesses the self-similar feature, and the calculated self-similarity parameter H (apart from the periodogram) is nearly the same as was the model input parameter.

The chaotic map model possesses the self-similar property also, but the other statistical features—based on our preliminary studies [15]—are different from the expected ones. The model could be a good source model for a single traffic source emitting self-similar cell streams, but not appropriate to model aggregated network traffic.

The most promising modeling approach in our case is the FGN traffic model. The synthetic traffic trace generated by the model has nearly the same statistical properties and self-similar features than the measured traffic trace.

7. SHAPING AND QUEUEING ANALYSIS

From engineering point of view there are various impacts of self-similarity in ATM networks. An important issue is the buffer sizing. A heavier than exponential decays of queue length distributions are expected which imply larger buffer requirements. It has great importance concerning ATM switch designers because the first generation ATM switches have under-dimensioned buffers. In this section queueing analysis results are reported to investigate this issue.

Another important question is related to the robustness of self-similar features as the traffic flows through several queueing stages in the network. Moreover, an important question is whether the self-similarity can be removed from the traffic by shaping it or not. This question is important because if self-similarity can be removed or reduced by shaping at the entrance of the network then engineers could be free from the complex problem of self-similarity in the network concerning buffer sizing, traffic and congestion control, protocol design, etc. We analysed shaped self-similar traffic and report these results in the following.

7.1. Analysis of shaped self-similar traffic

Our shaping method was the *leaky bucket shaping* which forces to delay non-conforming cells. Consider a leaky bucket with leak rate R and bucket size M. Cells which find the bucket content smaller than M are directly admitted to the network; otherwise, they are queued with FIFO discipline and admitted to the network with rate R.

The FUNET1 data (with average rate 14Mbps) was shaped with parameters $M = 128$ cells, $R = 20$Mbps (F1-8-128) and $R = 15.5$Mbps (F1-10-128). The analysis results of the shaped data sets are illustrated in Figure 1, 2 and 3. The correlation of the shaped cell streams are slightly affected due to the shaping procedure. The IDI curve is shifted down but also shows the same correlation structure (Figure 2). The IDC and variance-time plots (Figure 1 and 3) demonstrate the remained long-term correlation. Our results concerning the robustness of self-similar features are in agreement with the the shaping studies in [16,23]. However, the shaping effect surprisingly resulted in even higher values for the estimated Hurst parameter (See Figure 1 and 3). *The estimated Hurst parameter is increased due to shaping.* It can be explained as follows. On short time scales the shaping

procedure smooths out the cell stream. That is why the variance of the number of cells in a given window is decreasing which results in a shifting IDC curve to smaller values. However, on large time scales there is no significant effect of the shaping so there is no change in the IDC curve as can be seen in the figures. Therefore, it is obvious that the estimated Hurst parameter will be higher. This result is questioning the interpretation of the estimated Hurst parameter because it is believed as a measure of burstiness. On the contrary, our example shows that if we are smoothing the traffic the Hurst parameter is increasing. If the process is a pure self-similar process there is a good interpretation of the Hurst parameter (see e.g. [24]), but in practice where the traffic structure is modified by several mechanisms (shaping, queueing, multiplexing, etc.) the process is not pure self-similar. The question is, what is the interpretation of the estimated Hurst parameter? These results motivate our future research in that direction.

7.2. Queueing analysis

A queueing analysis based on simulation with the original (FUNET1) and the shaped cell stream (F1-8-128) was also performed. The investigated cell stream was fed to a FIFO queue with service rate $r = 17$Mbps and $r = 15.5$Mbps and the complementary queue length distributions are plotted in Figure 9. For comparison the distributions related to the Poisson process with the same mean rate and the shuffled FUNET1 cell streams are also plotted. As we can see, both the original and the shaped cell streams resulted in significantly higher queue lengths compared to the Poisson or shuffled cell arrival processes. These queueing results also support the conclusion that the self-similarity properties (e.g., the long-term correlation) are quite robust and cannot be removed by shaping.

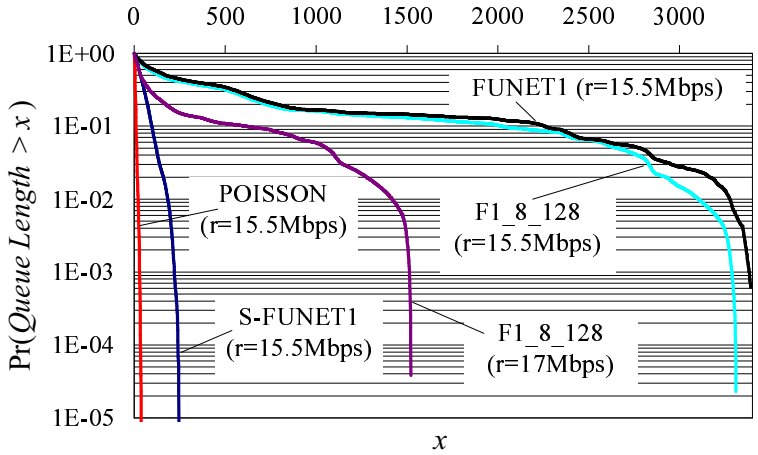

Figure 9. Queue length distributions for the original-, shaped-, and shuffled FUNET1 data and for a simple Poisson process.

Our results are consistent with the engineering intuition that FIFO queues behave as low-pass filters and the long-range correlations of the self-similar traffic (having power spectra divergence at low frequencies) are not affected. It also means that shapers would have to incorporate very large buffers which cannot be used in many applications because of the extreme introduced delay.

8. CONCLUSIONS

In this paper we presented ATM traffic measurements, analysis and modeling results focusing on capturing the self-similarity phenomenon. Moreover, we investigated the issue of shaping self-similar ATM traffic with queueing analysis. On the basis of the considerations presented in our paper, the following conclusions can be drawn:

(1) the self-similarity properties are present in the investigated WAN ATM traffic;

(2) the degree of self-similarity of the investigated ATM traffic is about 0.7 measured in Hurst parameter;

(3) the investigated three different models (*fractional Brownian traffic*, *superposed on/off sources* and *chaotic maps*) are all capable to capture the self-similar properties of ATM traffic but the *fractional Brownian traffic* was found to be the superior;

(4) strong robustness of self-similarity properties was identified which cannot be removed from the traffic by shaping;

(5) the interpretation of the estimated Hurst parameter is problematic in practice.

Acknowledgements The ATM measurements and part of the work reported here were carried out at the Laboratory of Telecommunications Technology at Helsinki University of Technology. We particularly want to thank Prof. Kauko Rahko for the opportunity to work at his laboratory. The traffic traces were gathered with the help of Marko Luoma and Jorma Jormakka. We gratefully acknowledge helpful discussions with Ilkka Norros about his FGN approach to packet traffic modeling.

REFERENCES

1. J. Beran, R. Sherman, M. S. Taqqu and W. Willinger. Long-range dependence in Variable-Bit-Rate video traffic. *IEEE Trans. on Comm.*, 43(2/3/4):1566–1579, 1995.
2. J. Beran. Statistical methods for data with long-range dependence. *Statistical Science*, 7(4):404–427, 1992.
3. D. R. Cox. Long-range dependence: A review. In H. A. David and F. T. David, editors, *Statistics: An Appraisal*, pages 55–74, Ames, Iowa, 1984. The Iowa State University Press.
4. D. R. Cox and P. A. W. Lewis. *The Statistical Analysis of Series of Events*. Menthuen, 1966.
5. N. G. Duffield, J. T. Lewis, N. O. Connel, R. Russel and F. Toomey. Statistical issues raised by the Bellcore data. In *11th Teletraffic Symposium*, UK, 1994.
6. D. E. Duffy et al. Statistical analysis of CCSN/SS7 traffic data from working CCS subnetworks. *IEEE JSAC*, 12(3):544–551, 1994.
7. A. Erramilli and W. Willinger. Fractal properties in packet traffic measurements. In *ITC Regional Seminar*, pages 144–158, St. Petersburg, Russia, 1993.
8. A. Erramilli et al. Chaotic maps as models of packet traffic. In *ITC-14*, 1994.
9. H. J. Fowler and W. E. Leland. Local area network traffic characteristics, with implications for broadband network congestion management. *IEEE JSAC*, 9(7):1139–1149, 1991.
10. M. W. Garret and W. Willinger. Analysis, modeling and generation of self-similar VBR video traffic. In *ACM SIGCOMM'94*, pages 269–280, London, UK, 1994.

11. W. E. Leland, M. S. Taqqu, W. Willinger and D. V. Wilson. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Trans. on Networking*, 2(1):1–15, February 1994.
12. W. E. Leland and D. V. Wilson. High time-resolution measurement and analysis of LAN traffic: Implications for LAN interconnection. In *INFOCOM'91*, pages 1360–1366, 1991.
13. B. B. Mandelbrot. Long-run linearity, locally Gaussian processes, H-spectra and infinite variances. *International Economic Review*, 10:54–60, 1994.
14. K. Meier-Hellstern et al. Traffic models for ISDN data users: Office automation application. In *ITC-13*, pages 167–172, Copenhagen, Denmark, 1991.
15. S. Molnár, A. Vidács and A. Hegedűs. Meeting a challenge: Modeling self-similar LAN/MAN traffic. In *8th IEEE Workshop on Local and Metropolitan Area Networks*, Berlin/Potsdam, Germany, August 1996.
16. A. L. Neidhardt and A. Erramilli. Shaping and policing of fractal traffic. In *ITC Specialists Seminar on Control in Communications*, pages 253–264, Lund, Sweden, September 1996.
17. I. Norros. Studies on a model for connectionless traffic, based on fractional Brownian motion. In *Conference on Applied Probability in Engineering, Computer and Communication Sciences*, Paris, France, June 1993
18. I. Norros. A storage model with self-similar input. *Queueing Systems*, 16:387–396, 1994.
19. I. Norros. On the use of fractional Brownian motion in the theory of connectionless networks. *IEEE JSAC*, 13(6):953–962, August 1995.
20. V. Paxson. Empirically derived analytic models of wide-area TCP connections. *IEEE/ACM Trans. on Networking*, 2(4):316–336, 1994.
21. V. Paxson. Fast approximation of self-similar network traffic. Technical Report LBL-36750, Lawrence Berkeley Lab. and EECS Division, Univ. of California, Berkeley, Berkeley CA 94720, April 1995.
22. V. Paxson and S. Floyd. Wide-area traffic: The failure of Poisson modeling. *IEEE/ACM Trans. on Networking*, 3(3):226–244, June 1995.
23. H. Saito and T. Tsuchiya. Upper bound of loss probability for self-similar traffic In *ICC'96*, Dallas, Texas, USA, June 1996.
24. G. Samorodnitsky and M. S. Taqqu. *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance*. CHAPMAN & HALL, 1994.
25. R. P. Singh and A. Erramilli. Application of deterministic chaotic maps to model packet traffic in broadband networks. In *7th ITC Specialists Seminar*, pages 8.1.1–8.1.3., Morristown, NJ, 1990.
26. M. S. Taqqu and J. Levy. Using renewal processes to generate long-range dependence and high variability. Birkhäuser (ed.), *Dependence in Probability and Statistics*, pages 73–89, Boston, 1986.
27. 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. In *ACM/Sigcomm'95*, 1995.