How to Characterize Hursty Traffic?*

S. Molnár, A. Vidács

High Speed Networks Laboratory,
Dept. of Telecommunications and Telematics,
Technical University of Budapest,
H-1111, Sztoczek u. 2., Budapest, Hungary,
Tel: (361) 463 3889, Fax: (361) 463 3107,
E-mail: {molnar,vidacs}@ttt-atm.ttt.bme.hu

Abstract

A discussion on the practical problems of fractal characterization of network traffic is presented. Our investigation is focused on the estimation and interpretation of the Hurst parameter in case of real traffic. We present an analysis based on measured ATM WAN traffic and we show that the estimation of the Hurst parameter depends on several factors, e.g. the estimation technique, sample size, time scale, level shifts, correlation structure. The dependencies are illustrated in examples including the effects of practical mechanisms like shaping or policing. It is shown that if use the Hurst parameter in practice we are faced with various misleading effects that can deceive our self- similarity tests and Hurst parameter estimation methods. Finally, we conclude that the estimated value of the Hurst parameter may be distorted in many practical cases and it may have no information for practical usage.

1 INTRODUCTION

During the last half decade a number of extensive studies of high resolution traffic measurements from a wide range of packet networks (Ethernet LANs, CCSN/SS7, ISDN, ATM) have been reported [1, 3, 6, 9, 10, 14, 15]. An interesting finding of these studies is the fractal nature of aggregate packet traffic which has opened a new venue in the teletraffic research. The most frequently identified properties are *long-range dependence* and *self-similarity*. Roughly speaking a weakly stationary process is called long-range dependent if the correlation between neighbouring disjoint blocks does not asymptotically vanish when the block size is increased. This property is in contrast to the exponentially decaying correlation of traditionally used stochastic processes (e.g. Markov processes, autoregressive processes, etc.). Moreover, a stochastic process is called self-similar if it behaves, up to a scaling factor, in exactly the same way at all time scales. This rather complex 'burst within burst' structure of the investigated packet traffic has also been reported by these studies which is in sharp contrast to the classical traffic modeling assumptions.

The self-similar traffic models seem to be quite successful [7, 8, 19, 18, 22, 24, 27]. They are very appealing from the point of view that a very complex bursty and correlated traffic stream can be described by only a few parameters. In most cases three parameters are appropriate for the traffic characterization (mean rate, peakedness and Hurst parameter). The basic idea behind using

*This report is based on our investigations published in [20].

self-similar traffic models is that self-similarity is assumed to be an *inherent nature* of the traffic. However, it should be noted that in spite of the facts that some fractal properties are identified and self-similar models fit these data better than conventional traffic models there is still no final conclusion whether we can consider self-similarity as a target from a modeling point of view or is it just a manifestation of an unknown traffic nature. Unfortunately, it is not possible in practice to check whether by definition a given traffic trace is self-similar or not and we can only look for some features of self-similarity. However, the detection of self-similarity only by identified properties could be misleading. Several nonstationary processes, e.g. level shifting processes [5] which can be observed in the superposed effects of different protocol levels [12] (ATM interface card based bursts, IP frames, window mechanisms, session procedures, etc.) can produce such properties. It means that if we found the traffic to be Hursty[1] it is due to long range dependence or nonstationarity.

Moreover, even if we choose a self-similar traffic model, the estimation and interpretation of these characteristics (especially the measure of self-similarity, the Hurst parameter) seems to be problematic in practice. The application of these models for practical purposes can also be questioned because real traffic can never be considered as pure self-similar traffic and in many dimensioning purposes (e.g. buffer dimensioning) the time scale of detected long range dependence is beyond the relevant time scales of the investigated queueing phenomena [11, 25]. The purpose of this paper is to discuss these problems, which are very important from a practical point of view, because what we would like to achieve is a method to correctly dimension our networks with practical methods.

In the paper we discuss and investigate the following problems based on measured data from an ATM WAN network:

- Hurst parameter estimation dependence on estimation technique, sample size, time scales, level shifts, correlation structure (shuffling, shaping, policing and cell disturbing effects).

- Interpretation of the estimated Hurst parameter.

The paper is organized as follows: in the next section we briefly describe our ATM measurements. In Section 3 we overview the basic definitions and features of fractal traffic focusing on the properties which are tested by the most popular self-similarity tests with their problems. The Hurst parameter estimation techniques with our results are described in Section 4. The impacts of several characteristics on the Hurst parameter and our discussion about the interpretation of Hurst parameter are reported in Section 5 and 6. We then conclude with a summary of the paper in Section 7.

2 ATM TRAFFIC MEASUREMENTS

The FUNET network The measurements were made on the FUNET ATM WAN network. 'FUNET' stands for 'Finnish University and Research Network', which provides primarily Internet services to its members based on TCP/IP-protocol. All these services are provided by CSC—Center for Scientific Computing which is a national service center that specializes in scientific computing and data communications providing modeling, computing and information services for universities, research institutes and industry. The FUNET long-distance 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 [16].

[1] Hursty traffic means that the classical Hurst parameter estimation methods, e.g. R/S tests, estimate a Hurst parameter which is bigger than 0.5.

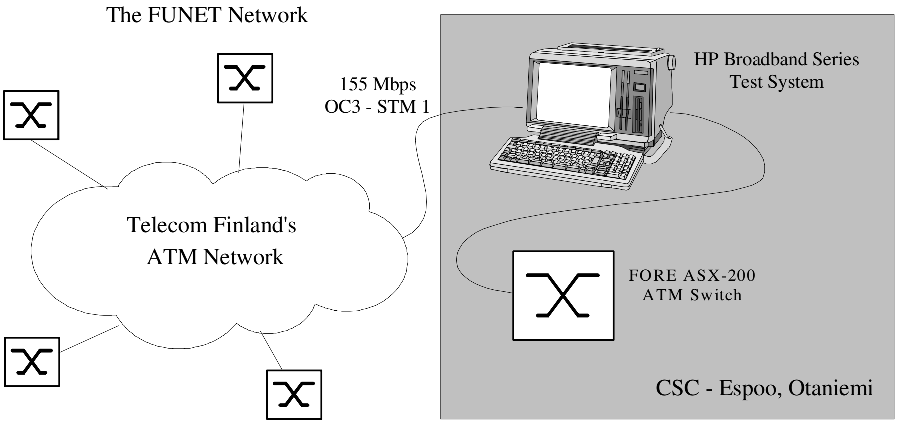

Figure 1: The FUNET measuring configuration.

Measuring tool and configuration The measurement was made at the CSC in Espoo, Otaniemi. This location is in the logical center of the whole network. 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 (see Figure 2). From that 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. During our work only the cell capture capability of the measurement unit was used: 131 072 ATM cells mapped into a 155 Mbps SONET/SDH signal can be recorded into the 8MB of capture memory of this equipment. All cells are timestamped with the calendar time with resolution 0.01 μs.

The measured data The aggregated traffic at the most heavily loaded point of the FUNET network was measured, including Internet traffic, data transfer and supercomputer usage. During the measurements, two types of data collections were made. In the first scenario the measured data was the time stamp of the arrival time instant for every single cell on the link. Because of the upper limit for the number of captured cells each measured data file contains 131 072 time stamps only, which corresponded to about 3–5 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 a one second interval. In this case the time interval of the observation could take several minutes long. A summary of these data sets is given in Table 1. The files FUNET1, FUNET2 and FUNET3 contain traffic data captured from the incoming traffic from the whole country to the CSC, and the FUNET4 measurement was made on the outgoing link. In the 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. The average traffic load was about 14 Mbps for the first three measurements, and about 8 Mbps in the case of the FUNET4 data. In the following, we refer to the data above listed as the 'FUNET measurements'.

The traffic structure As for the first four data sets, the measurement unit was able to register the VPI and VCI fields from the cell headers, too. Using this extra information we can reveal the structure of the aggregated traffic stream. Comparing the VPI/VCI fields the aggregated cell stream

Table 1: Qualitative description of the measured data sets and the values of Hurst-parameter H calculated from different statistical methods

Filename	#packets	Time (sec)	$\hat{H}_{idc}$	$\hat{H}_{var}$	$\hat{H}_{rs}$	$\hat{H}_{per}$
FUNET1	131,072	3.9	0.7	0.7	0.68	0.68
FUNET2	131,072	5.1	0.67	0.67	0.67	0.73
FUNET3	131,072	4.4	0.66	0.66	0.68	0.68
FUNET4	131,072	6.4	0.72	0.72	0.74	0.78
FUNETSTA.T3	14,807,546	425	0.70	0.70	0.82	0.94
FUNETSTA.T4	43,768,430	1964	0.67	0.67	0.79	0.90

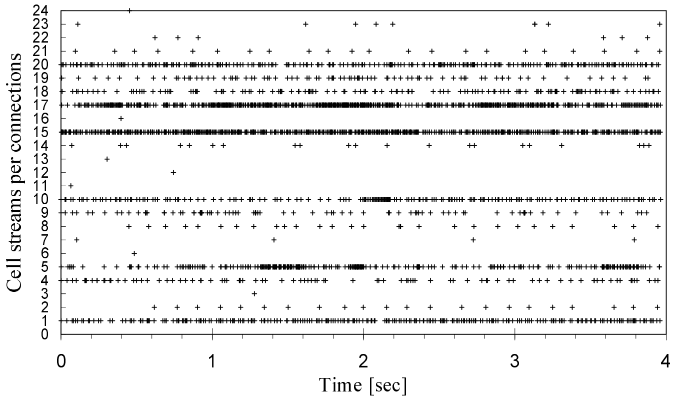

Figure 2: The structure of the FUNET1 traffic trace.

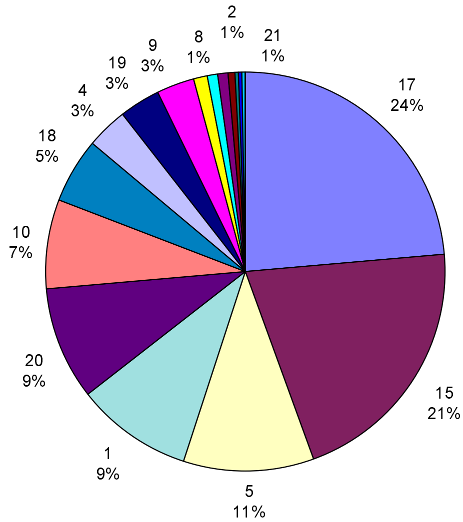

Figure 3: The bandwidths of connections in the FUNET1 data.

can be divided into independent connections. (Note that connection means a cell stream with common VPI and VCI fields in the headers. We do not have any information about the type of traffic carried by these cell streams.) The most important piece of information for us is the number of connections and their relative cell rate compared to each other. A detailed analysis was made for the FUNET1 data set. Figure 2 shows the separated cell streams schematically. (Because of the huge number of cells each hair-cross represents every 50th cell arrival in a connection.) As can be seen from the figure, during the 4 second measured time period 24 connections were in progress. The connection with highest rate contains about 30 000 cells which is about 24 percent of the whole aggregated traffic as well as the first dozen with highest intensity contain 99% of all the cells. Figure 3 shows the pie-chart of bandwidths of connections.

In our investigation the question of stationarity is fundamental. As far as it can be concluded from Figure 2 without a comprehensive stationary analysis, the cell streams are homogenous enough in time apart from the bursty nature of ATM traffic. There is no connection turned on or off in the middle of the measurement time and the rates apart from the burstiness are not changing considerably.

3 FRACTAL TRAFFIC

3.1 Definitions, Properties and Characterization

In this section we summarize the basic definitions, properties and characterization of fractal traffic. Consider a real valued continuos time stochastic process X_t ($t \in \mathbf{R}_+$) with stationary increments $Y_i = X_i - X_{i-1}$ ($i \in \mathbf{N}$).

Definition 1 The process X_t with self-similarity (or Hurst-) parameter $H > 0$ is called

- *strictly self-similar* if $X_{\alpha t}$ has the same finite dimensional distributions as $\alpha^H X_t$ for any $\alpha > 0$.
- *second-order self-similar* if the process $X_{\alpha t}$ and $\alpha^H X_t$ have the same second order characteristics for any $\alpha > 0$.
- *asymptotically second-order self-similar* if the second order characteristics of $X_{\alpha t}$ and $\alpha^H X_t$ are the same for $\alpha \to \infty$.

A nondegenerate self-similar process with stationary increments and finite variance has the following properties:

- *Hurst-parameter:* $0 < H \leq 1$.
- *Mean:* $E\{X_t\} = 0$.
- *Variance:* $Var\{X_t\} = t^{2H} Var\{X_1\}$ for $t > 0$.
- *Covariance:* $Cov\{X_{t_1}, X_{t_2}\} = \frac{1}{2}[t_1^{2H} + t_2^{2H} - (t_2 - t_1)^{2H}] Var\{X_1\}$ for $t_2 > t_1 > 0$.

In modeling, it is the increments of X_t that are of interest, because the sequence $\{Y_i\}$ displays both the "Joseph effect" (synonymous with *long-range dependence*, see below) and the "Noah effect" (synonymous with *slowly decaying variances* or *high variability*) of Mandelbrot [17].

The autocovariances $r(k) = Cov\{Y_i, Y_{i+k}\}$ of the increments Y_i are of the form

$$\begin{aligned} r(k) &= \frac{1}{2}[|k+1|^{2H} - 2|k|^{2H} + |k-1|^{2H}] Var\{Y_i\} \\ &= Var\{Y_i\} H(2H-1) k^{2H-2} + O(k^{2H-3}). \end{aligned} \tag{1}$$

For $H = 0.5$ the sequence of Y_i's is uncorrelated. The case $H \in (0, 0.5)$ is less interesting for statistical applications. For $H \in (0.5, 1)$, $\sum_k r(k) = \infty$. By definition:

Definition 2 A stationary process is called *long-range dependent* if its autocovariance function $r(k)$ is nonsummable, i.e., $\sum_k r(k) = \infty$.

Otherwise it is called *short-range dependent*. Equivalent definitions can also be given for specifying long-range dependence by the *nonvanishing correlations* between large blocks, or by the *singularity of the spectral density at zero* [3].

Now we describe the cumulating arrival process A_t representing the number of arrivals within the interval $(0, t]$ by $A_t = mt + X_t$ where X_t is a self-similar process with stationary increments with $Var\{X_1\} = 1$. The cumulating arrival process A_t has the following properties:

- *Mean:* $E\{A_t\} = mt$.

- *Variance:* $Var\{A_t\} = t^{2H}$.
- *Covariance:* $Cov\{A_{t_1}, A_{t_2}\} = \frac{1}{2}[t_1^{2H} + t_2^{2H} - (t_2 - t_1)^{2H}]$ for $t_2 > t_1 > 0$.

Furthermore, the following features can be derived from Definition 1 and the above three properties:

1. *Slowly decaying variances:* $Var\{A_{at}/a\} = a^{2H-2}t^{2H}$
2. *Long-range dependence:* For $H > 0.5$ the process A_t is long-range dependent.
3. *1/f-noise:* Long-range dependence manifests itself in the frequency domain as *1/f-noise*, which is the term used to refer to a sharp divergence in the power spectrum near the origin of the form $c_1|x|^{1-2H}$, where c_1 is a finite positive constant.
4. *Hurst effect:* The *rescaled-adjusted range statistics* (see Section 4) is characterized by $E\{R(d)/S(d)\} \sim c_2 d^H$, as $d \to \infty$ with $0.5 < H < 1$ where c_2 is a finite positive constant.

An important practical question is how we can characterise and model a given (e.g., measured) A_t traffic. A simple candidate based on second-order properties is the Gaussian process. We can use a Gaussian self-similar process called *Fractional Brownian Motion* (FBM) with Hurst-parameter H ($0.5 < H \leq 1$) for modeling X_t. In this case the cumulating arrival process A_t can be described as follows [23]:

$$A_t = mt + \sqrt{am}Z_t$$

, where Z_t is a normalized FBM characterized by the following properties:

- Z_t has stationary increments;
- $Z_0 = 0$, and $E\{Z_t\} = 0$ for all t;
- $Var\{Z_t\} = E\{Z_t^2\} = |t|^{2H}$ for all t, where $H \in [1/2, 1)$;
- Z_t has continuous paths;
- Z_t is a Gaussian process, i.e., all its finite-dimensional marginal distributions are Gaussian.

This model seems to be a good choice and has popularity in the literature [13, 15, 22, 23, 24] because it has the advantage of simple characterization of a rather complex 'burst within burst' traffic. There are only three parameters of this traffic model: m (input rate), a (variance coefficient) and H (Hurst parameter). The estimation of mean and variance coefficients is theoretically feasible but we note that in practice it is generally not very easy. However, the crucial point is to correctly determine the Hurst parameter. This problem is the main topic of this paper.

3.2 Problems

The analysis of testing for self-similarity and the estimation of the Hurst parameter are not easy in practice. The problem is that we are of course always dealing with finite data sets so it is not possible to check whether by definition a traffic trace is self-similar or not. We are therefore forced to look for different features of self-similarity (see the list of properties of A_t in Section 3.1) in our actual measured traffic.

More specifically, the *first problem* which should be addressed is that even if we detect some properties listed above we can not directly conclude that the analyzed data is taken from a self-similar process since there are other effects that can produce the same properties (e.g., nonstationary processes) so our analysis based only on these tests can be misleading as we will demonstrate it in Section 5. It is only reasonable to discuss about *self-similar behaviour over a given time-scale for a given data set.*

The *second problem* is that the estimated Hurst parameter depends on many factors (i.e., estimation technique, sample size, time scale, correlation structure, etc.) which makes it difficult to find the most relevant 'estimated H' for a given purpose. We also discuss this problem in Section 5.

Finally, the *third problem* with using the Hurst parameter for practical purposes (e.g., dimensioning of networks) is that the interpretation of the Hurst parameter (which is clear for pure self-similar processes) is not very clear for real traffic which can never be considered as a pure self-similar process. This problem is discussed in Section 5.

4 HURST PARAMETER ESTIMATION

In this section four different statistical tests are presented.

Index of Dispersion for Counts This commonly used measure for capturing the variability of traffic over different time scales is provided by the *index of dispersion for counts* [4]. For a given time interval of length t, the index of dispersion for counts (IDC) is given by the variance of the number of arrivals A_t during the interval of length t divided by the expected value of the same quantity:

$$IDC(t) = Var\{A_t\}/E\{A_t\}. \tag{2}$$

For a finite data set, the variance of A_t can be calculated by dividing the whole series into nonoverlapping blocks of length t and treat them as different instances of A_t.

Self-similar processes produce a monotonically increasing IDC of the form $m^{-1}t^{2H-1}$. Plotting $\log IDC(t)$ against $\log t$, this property results in an asymptotic straight line with slope $2H - 1$ [15]. (This behaviour is also based on Property 1 in Section 3.1, in contrast to traditional processes where the IDC is bounded.)

Figure 4(a) depicts 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 and is estimated to be about 0.4, resulting in an estimate $\hat{H}$ of the Hurst-parameter H of 0.7.

The same analysis was made for all the data sets. Table 1 shows the results: the values of the estimated Hurst-parameter $\hat{H}$. As can be seen from the table, the values of $\hat{H}$ are pretty much the same for all the data sets. It is remarkable that in the case of the last two data sets the analyzed process was the sequence of cell counts in each second 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 analysis This method is based on the property that a self-similar process has slowly decaying variances (see Property 1 in Section 3.1). The so-called variance-time plot is obtained by plotting $\log Var\{X_{at}/a\}$ against $\log(a)$ and by fitting a simple least squares line through the resulting points in the plane, ignoring the small values of a. Values of the estimated asymptotic slope $\hat{\beta}$ between -1 and 0 suggest self-similarity, and the estimate for degree of self-similarity is given by $\hat{H} = 1 + \hat{\beta}/2$ [15].

The corresponding plot for the FUNET1 data set can be seen in Figure 4(b). The estimated values of $\hat{H}$ are listed in Table 1. Since the variance-time plots and the IDC diagrams are closely related

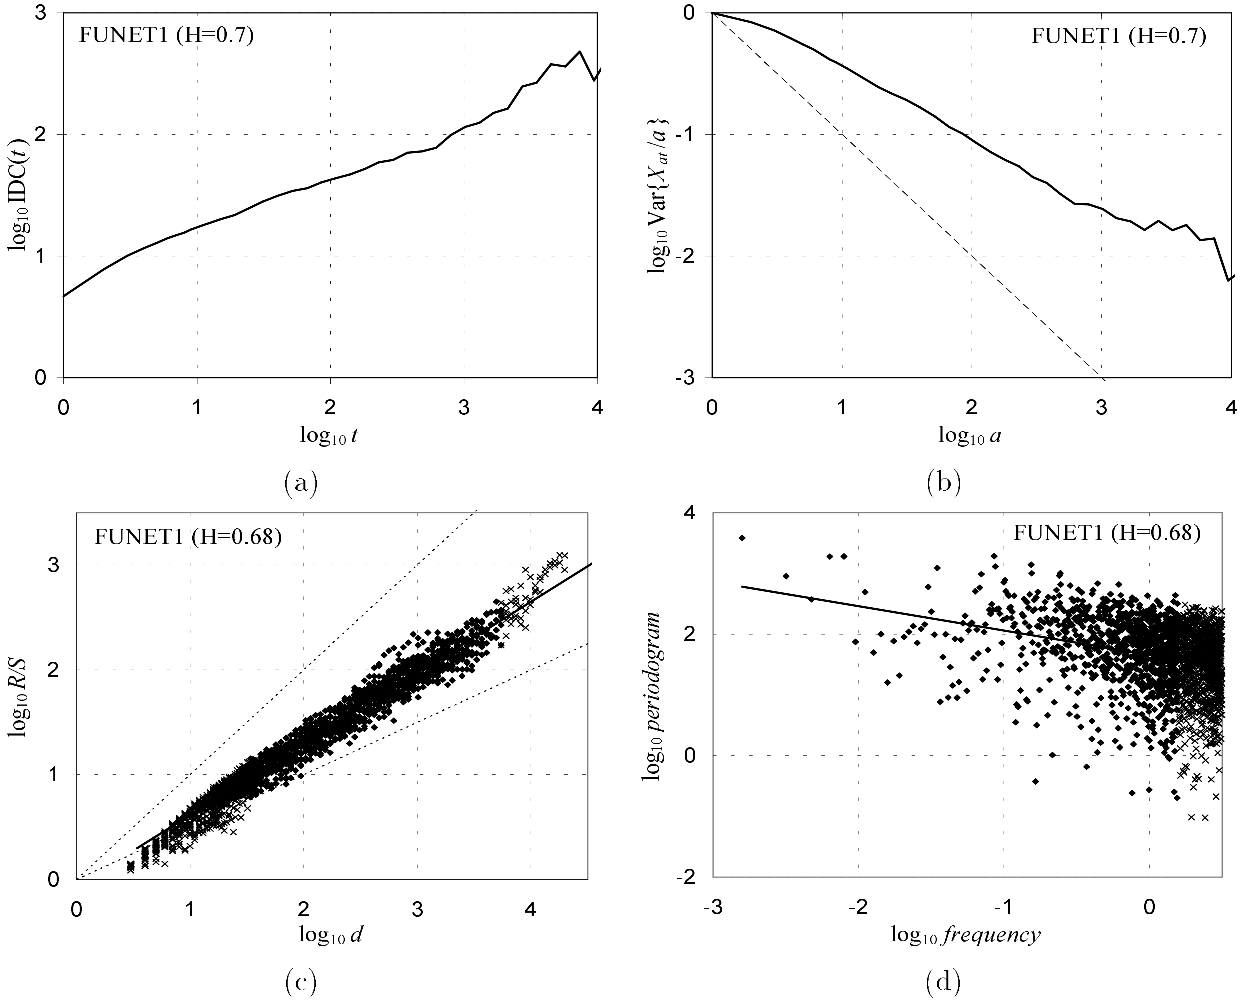

Figure 4: (a) IDC plot, (b) variance-time plot, (c) R/S diagram, and (d) periodogram plot for FUNET1 data.

statistical methods, the results obtained from this method are the same as in the previous subsection.

R/S analysis This method tries to capture the Hurst parameter (see Property 4 in Section 3.1) based on the rescaled adjusted range statistics.

Given an empirical time series of length N ($X_k : k = 1, \ldots, N$), the whole series is subdivided into K nonoverlapping blocks. Now, we compute the rescaled adjusted range $R(t_i, d)/S(t_i, d)$ for a number of values d, where $t_i = \lfloor N/K \rfloor (i-1) + 1$ are the starting points of the blocks which satisfy $(t_i - 1) + d \leq N$.

$$R(t_i, d) = \max\{0, W(t_i, 1), \ldots, W(t_i, d)\} - \min\{0, W(t_i, 1), \ldots, W(t_i, d)\}, \tag{3}$$

where

$$W(t_i, k) = \sum_{j=1}^{k} X_{t_i+j-1} - k \cdot \Big(\frac{1}{d} \sum_{j=1}^{d} X_{t_i+j-1}\Big), \quad k = 1, \ldots, d. \tag{4}$$

$S^2(t_i, d)$ denotes the sample variance of $X_{t_i}, \ldots, X_{t_i+d-1}$. For each value of d one obtains a number of R/S samples, which decreases from K for larger values of d. One computes these samples for logarithmically spaced values of d, i.e., $d_{l+1} = m \cdot d_l$ with $m > 1$, starting with d_0 of about 10. Plotting $\log R(t_i, d)/S(t_i, d)$ vs. $\log d$ results in the R/S plot, 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 and largest values of d are omitted. The slope of the regression line is an estimate for H [2].

Figure 4(c) 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 (see Table 1).

Periodogram-based analysis This method is used to identify the manifestation of self-similarity by frequency domain analysis of the measured data (see Property 3 in Section 3.1). Let $I(\cdot)$ denote the sample periodogram (i.e., power spectrum as estimated using a Fourier transform) defined by

$$I(\lambda) = \frac{1}{2\pi N} | \sum_{j=1}^{N} X_j e^{ij\lambda} |^2, \qquad \lambda \in [0, \pi). \tag{5}$$

As mentioned in Section 3.1, 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 $1 - 2H$. In most of the cases this will lead to a wrong estimate of H since the periodogram estimation method is unbiased and inconsistent. However, this method can reveal the power spectrum near the origin. The periodogram plot is obtained by plotting $\log(I(\lambda))$ against $\log \lambda$.

Figure 4(d) 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—in the present context, the regression line was fitted to the lowest 50% of all frequencies—is clearly different from zero, the slope estimate is about -0.36 which yields $H = 0.68$. This result corresponds to the previously calculated values of H.

The analysis was made for all the data sets, and the results are listed in Table 1.

Summary To summarize the results listed in Table 1, we conclude that:

- The estimated values of the parameter H are definitely greater than 0.5 for all cases.
- The values of H are nearly the same for all of the four analysis methods and for all the data sets. The common value for it is about 0.7. (Apart from the last two values for the FUNETSTA data sets.)

In spite of this, it would be too early to say that it follows from the results above that the measured traffic is self-similar with self-similarity parameter 0.7. To establish such a statement, we should carefully examine the applied analysis methods with their preliminary conditions and confidence intervals as well as the structure of the analysed data sets in more details.

In the next section, we investigate the problems arising during the calculation of the parameter H and determine those effects which can influence the results considerably.

5 IMPACTS ON THE HURST PARAMETER

In practice, using measured data sets the estimated values of H obtained from different analysis methods are influenced by the following circumstances:

- dependence on estimating technique
- dependence on sample size
- dependence on time scales
- dependence on data structure

Next, we discuss these points and examine how the calculated self-similarity parameter is affected. In our investigation we focus on the robustness of H against these effects.

5.1 Estimating technique

In Section 4 we presented four different statistical methods testing for and estimating the degree of self-similarity, but this list is still far from complete.

Calculating variance-time plots and IDC plots is easy, the resulting diagram is descriptive. That is why these so called "eyeball tests" are popular and widely used in practice. However, these methods are not reliable for empirical records with small sample sizes.

In practice, R/S analysis is based on a heuristic graphical approach. The feature that makes R/S statistics particularly attractive is its relative robustness against changes in the marginal distribution [26]. With respect to the effectiveness of R/S analysis dependent on the sample size, similar comments as in the case of the previous two methods are also relevant.

The absence of any results for the limit laws of the previously mentioned statistics make them inappropriate when a more refined data analysis is required. In contrast, a more refined data analysis is possible using periodogram-based methods in the frequency domain. Several periodogram-based estimators can be found in the literature, such as maximum likelihood type estimates (MLE) and related methods [26]. In particular, for Gaussian processes Whittle's estimate MLE has been studied extensively. Using these approaches, more information can be collected on the H-estimate, such as confidence intervals.

In practice, when the required preliminary conditions for the statistical tests are not fully satisfied, the different methods can give slightly different estimates of H. (As can be seen in Table 1.) To give a fully detailed analysis for all four statistical methods is beyond the scope of this paper. Next, we choose one of them, the IDC plot for further study. This method, being simple and effective, requires small computation power and the resulting plot is descriptive enough to visualise the changes when the analysed data set is modified. (Note, when using this method the value of H is calculated by fitting a straight line to the plot assuming that the curve is linear. In practice, deviations from linearity introduce additional uncertainty to the estimation.)

5.2 Sample size

For a given time interval t the IDC value is defined by Equation 2 in Section 4. For a self-similar process, this value is increasing without limit as t tends to infinity. In practice, we only have finite data sets and the value of t can not exceed the sample size. Furthermore, to get a reliable estimate of IDC(t) the maximum window size is limited to the the 10% of the sample size, using nonoverlapping windows of length t we need at least about 10 values to calculate the variance with acceptable confidence. Thus,

Table 2: Calculated values of Hurst-parameter H estimated from different subsets of the FUNET1 data.

FUNET1 (100%)	50%	25%
$H = 0.70$	$H_1 = 0.69$	$H_1 = 0.61$, $H_2 = 0.57$
	$H_2 = 0.70$	$H_3 = 0.71$, $H_4 = 0.62$

the calculated IDC(t) value is getting more and more inaccurate as t increases. As a result, the IDC plot becomes more and more noisy as t increases.

Example In our analysis the critical issue is to examine the robustness of the estimated values of H against the sample size. To do this, we estimated the Hurst-parameter for different subsets of the whole FUNET1 data set. First we divided the data into two and four equal parts and calculated the Hurst parameter in each case. Table 5.2 shows the calculated values.

Discussion When the full data set was split into two parts, the estimates of H remained the same in both cases. However, when the data was split further, the calculated values for three subsets out of four differ from the original H.

5.3 Time scales

In our case the analysed process is the number of cells arriving in a fixed time interval of length t. The t represents the *time scale* at which the process is examined. Our task now is to find the appropriate time scales at which the IDC curve of the process increases monotonically. (As for the upper limit of t, the same holds as in the previous section, namely, the IDC is getting more and more inaccurate as t increases.)

Example 1 This terminology makes it possible for us to take advantage of the longer but not so detailed FUNETSTA measured data sets. Calculating the IDC values for FUNETSTA.T3 and plotting the result together with the IDC curve of the FUNET1 data set—taking into consideration the different time scales(!)—we get Figure 5. (Note, that there was a time gap between the FUNET1 and FUNETSTA.T3 measurements (about 30 min). Here we assume again, that the measured traffic was stationary.) Another drawback of bringing together the two plots is that the middle part of the curve is noisy and uncertain.

Discussion 1 Using a simple linear regression is a rough guess for this plot and estimating $\hat{H}$ gives the rather approximate value of 0.81. But the main result is that the curve increases monotonically throughout a time span that covers 6 orders of magnitude.

Example 2 In the previous example we tried to cover as many time scales as possible. Here we examine the Hurst parameter when estimated at shorter but different time scales. To do this, the regression line was fitted to different parts of the IDC curve using a sliding window which covers only one decade of time. The resulted values of H estimated from $(t, 10t)$ interval and plotted at the middle of each interval can be seen on Figure 6.

Discussion 2 The results show that estimating the Hurst parameter at different time scales yields to values of $\hat{H}$ from 0.68 to 0.8. This is a clear indication of the importance of the problem of 'what time scale to choose for Hurst parameter estimation' in practice.

5.4 Level shifting

During our analysis we always assumed that the statistical properties of the measured cell stream are independent from time. This assumption is questionable, but—since the greatest part of statistical

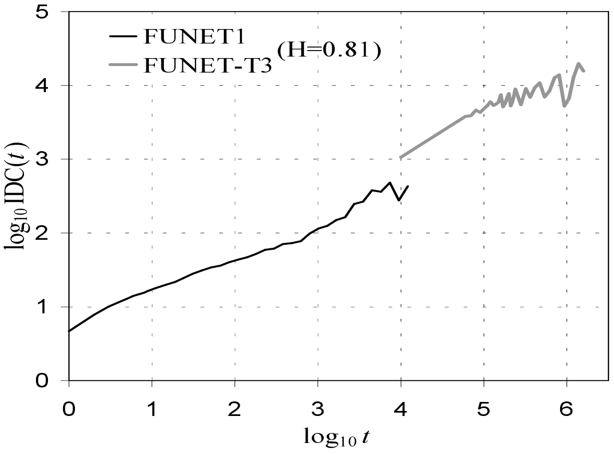

Figure 5: IDC plot for FUNET1 and FUNET-STA.T3 data sets.

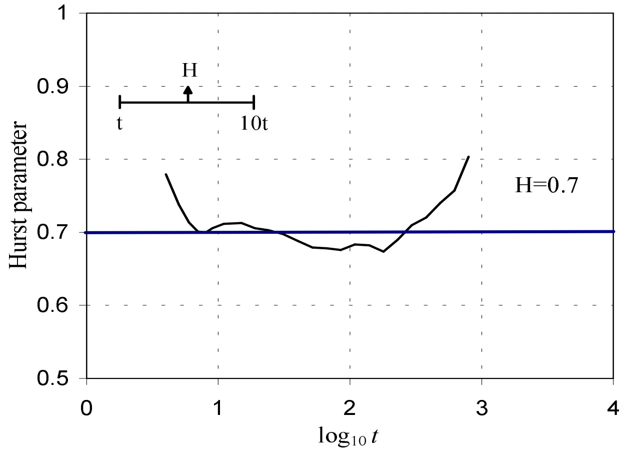

Figure 6: Hurst parameter estimates for different time scales.

analysis methods require stationarity as a basic preliminary condition—it cannot be avoided. Strictly speaking, we assume that the measured process is stationary in the wide sense which means that its mean is finite and independent of time, and its autocorrelation function is finite and is invariant of time shift. (If we decided to treat our measured data sets as nonstationary sequences it would be almost impossible to make a comprehensive analytical study with meaningful results general enough to use elsewhere. Furthermore, in the case of finite data sets it is not possible to discriminate a stationary long-range dependent sequence from a nonstationary one.)

In this section we investigate the case when the assumption of stationarity does not hold (i.e., there is a level shift present in the measured traffic traces.) We examine how robust our statistical test is in case of a nonstationary cell sequence with a change in the mean as a function of time.

Example 1 In this first simple model nonstationarity is introduced by adding a CBR traffic to the second half of the measured data set. (Note, that this example represents not just a theoretical problem but a possible event in practice: while measuring the network traffic suddenly a new source may start to emit cells with constant cell rate.) Figure 7 shows the calculated IDC plots for these new multiplexed data sets. 2.8Mbps (CBR20—20% of the load of FUNET1) and 7Mbps (CBR50—50% of the load of FUNET1) CBR rates were applied.

Discussion 1 The effect on the IDC plot is clearly visible. For the FUNET+CBR20 plot the upper part of the curve is moved up a bit as well as the lower segment shifted down slightly. As a result, the calculated Hurst parameter is greater, about 0.72. For the FUNET+CBR50 case the effect is sharper, the calculated value of $\hat{H}$ being 0.8.

Example 2 To understand the effects of level shift on the IDC plot more deeply, we investigated a simple CBR model in this example. As a starting point we chose a CBR traffic trace with the same rate as the mean rate of the FUNET1 traffic. The nonstationarity was introduced by increasing the CBR rate by 10, 20 and 50 percent abruptly at half time of the investigated time period. The IDC(t)value for an ideal CBR source without jitter is zero for all t which cannot be plotted on a logarithmical scale. The calculated IDC plots for the CBR traces with level shift can be seen on Figure 8.

Discussion 2 The calculated IDC plots for the CBR traces with level shift can be seen on Figure 8. All the IDC curves are straight lines with slope 1. The only difference is that the IDC values are higher when the level shift is stronger.

The simplicity of the examined CBR model makes it possible to calculate the IDC(t) values analytically. (In practice, calculating the IDC plot for a finite data set means evaluating a double sum to

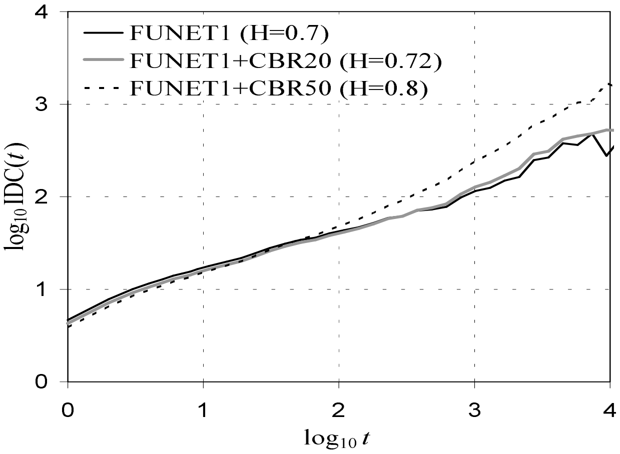

Figure 7: IDC plot for FUNET1 multiplexed with nonstationary CBR traffic.

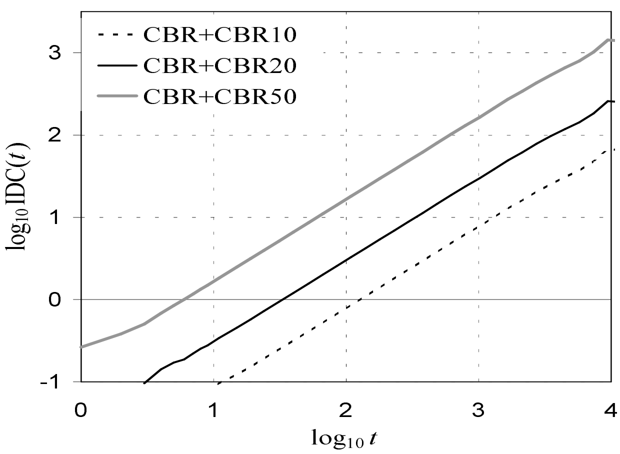

Figure 8: IDC plot for a CBR cell stream with level shift.

estimate the mean and the variance. The following results are derived from this IDC estimator.) The IDC(t) for the above data sets is of the form:

$$IDC(t) = \frac{(a_1 - a_2)^2}{2(a_1 + a_2)} t,$$

where a_1 and a_2 are the cell rates for the first and second half of the data respectively. For $a_1 \neq a_2$ the IDC plot is given by:

$$\log IDC(t) \simeq const + \log t, \quad \text{where} \quad const = \log \frac{(a_1 - a_2)^2}{2(a_1 + a_2)}$$

which gives us a straight line with slope 1.

The main result here is the fact that although the CBR data with level shift has nothing to do with self-similarity, the estimated IDC is a monotonically increasing straight line with slope 1.

Example 3 The first example is generalized here by replacing the CBR traffic with a Poisson process. Again, the FUNET1 data was modified by adding a Poisson traffic to the second half of the measured data to increase the mean rate by 20 and 50 percent. The calculated IDC plots can be seen in Figure 9.

Discussion 3 The effect of nonstationarity in the plots is the same as in Example 1. The upper part of the curves moved up and the lower-left segments are shifted down simultaneously, resulting in higher Hurst parameter estimates.

Example 4 To make the effect of level shifts on the IDC plot clearer, in this example a simple but inhomogeneous Poisson process is examined which changes its intensity in time. Here we consider the case when the Poisson source emits cells with rate λ_1 and suddenly changes its intensity to λ_2. Figure 10 presents the analysis result for these data sets. (For every process λ_1 was set to 1 and λ_2 changes as noted in the figure.)

Discussion 4 For such simple inhomogeneous Poisson processes the IDC estimate can be derived analytically. Let λ_1 and λ_2 denote the intensity parameters of the process for the two halves. Then, the IDC(t) value can be calculated as follows:

$$IDC(t) = 1 + \frac{(\lambda_1 - \lambda_2)^2}{2(\lambda_1 + \lambda_2)} t.$$

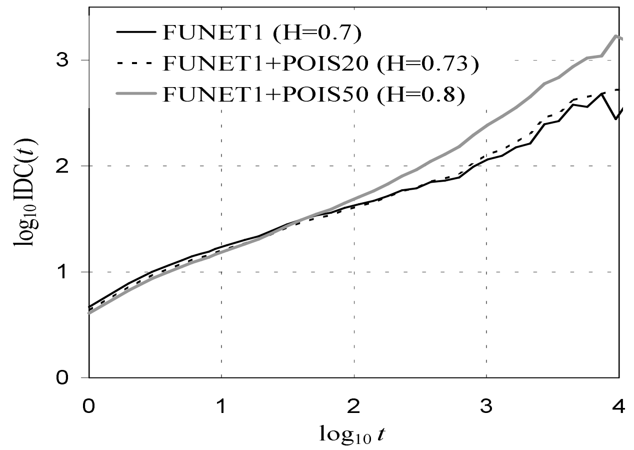

Figure 9: IDC plot for FUNET1 multiplexed with a nonstationary Poisson traffic.

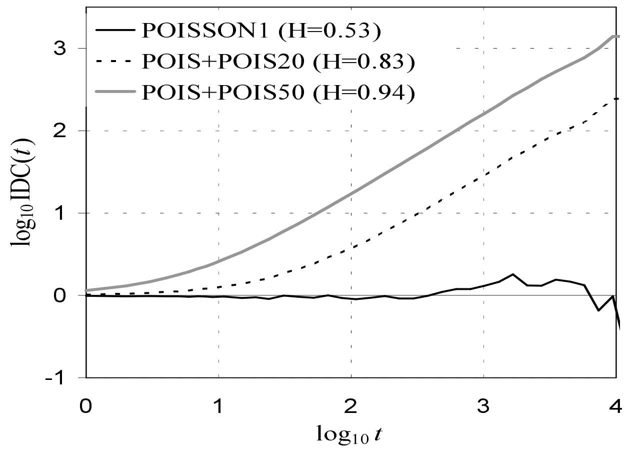

Figure 10: IDC plot for inhomogeneous Poisson processes.

For the appropriate IDC plot for $\lambda_1 \neq \lambda_2$ and $t \to \infty$ we get:

$$\log IDC(t) \simeq const + \log t, \quad \text{where} \quad const = \log \frac{(\lambda_1 - \lambda_2)^2}{2(\lambda_1 + \lambda_2)}.$$

This equation gives a straight line with slope 1 as an asymptote.

The main message from this example is again that a monotonically increasing IDC does not necessarily come from the self-similar nature of the analysed data. Instead, it comes from the nonstationarity present in the sample trace. We mention that a linearly growing IDC curve over many time scales can also be created even with a simple stationary Markovian model (e.g. with an Interrupted Poisson Process). In this case the increasing IDC curve again nothing has to do with self-similarity.

5.5 Correlation structure

In this section we next examine how the calculated Hurst-parameter is affected when the structure of the measured data set is modified by shuffling, shaping, policing, and disturbing the cell stream.

Example 1 This first example is just a theoretical one without any practical meaning but gives us useful information about the correlation structure of the analysed traffic. In order to show the difference from a short term correlated data we generated a new data set from the original FUNET1 data by mixing the sequence of cell interarrival times randomly thus building a new cell stream. The long-term correlations were obviously removed by this random shuffling which can be investigated in Figure 11.

Discussion 1 As for the shuffled FUNET1 data the IDC curve starts as in the case of the original plot but soon it stops increasing and remains constant for values of $\log t$ greater than 1. The self-similar feature disappeared but the resulting process is still bursty and possesses short-range dependence.

Example 2 This example investigates the effects of traffic shaping. Our shaping algorithm was the *leaky bucket shaping* which forces nonconforming cells to be delayed. 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 33 072 cell/s which is about 14Mbps) was shaped with parameters $M = 0$ (measured in cells), $r = 50Mbps$ (F1-3-0); $M = 0$, r=20Mbps (F1-8-0) and $M = 8$, r=20Mbps (F1-8-128). The analysis results are illustrated in Figure 12.

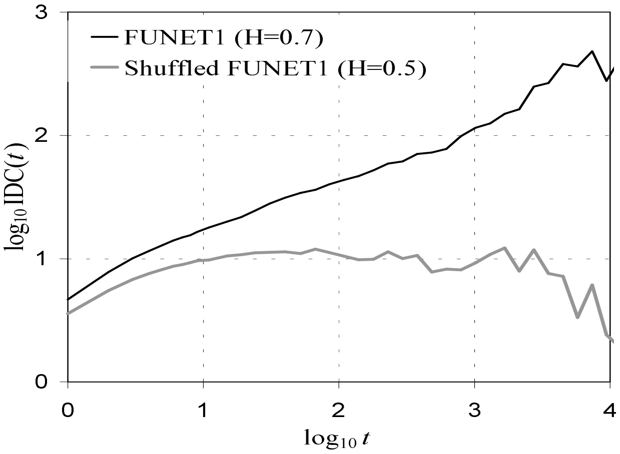

Figure 11: IDC plot for shuffled FUNET1 data.

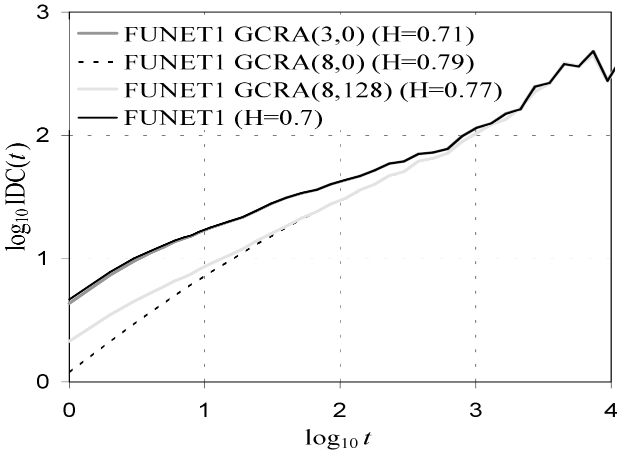

Figure 12: IDC plot for shaped FUNET1 data.

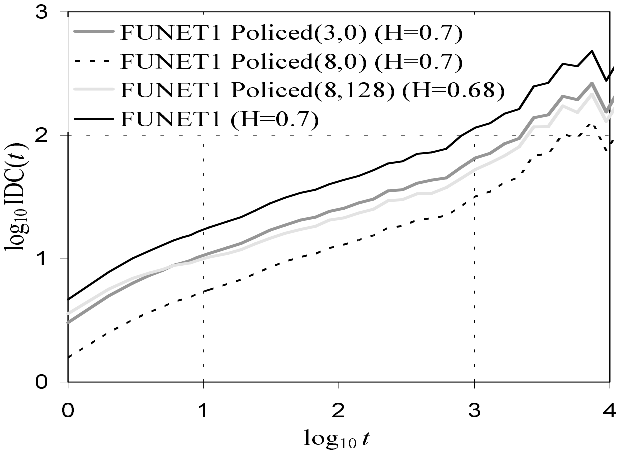

Figure 13: IDC plot for policed FUNET1 data.

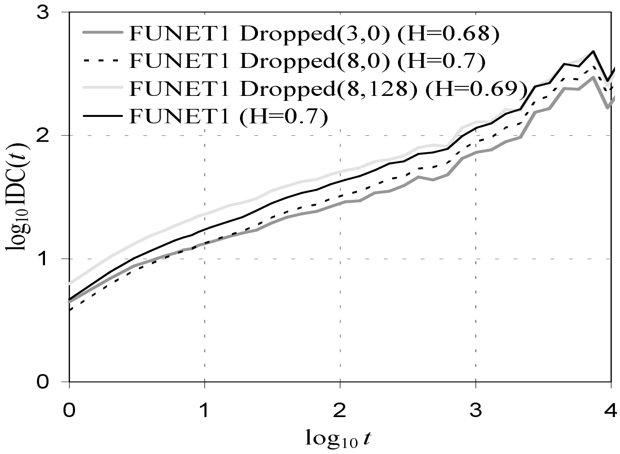

Figure 14: IDC plot for the overflow process of policed FUNET1 data.

Discussion 2 The correlation of the shaped cell streams are slightly affected due to the shaping procedure. The IDC plot demonstrate the remained long-term correlation in the data. These results are in consistent with the results in [21]. This means that for the purpose of removing long-range dependence a drastic shaping is needed which means that shapers would have to use very large buffers which cannot be used in many applications due to the extreme delay introduced. The shaping effect, however, resulted in even higher values for the estimated Hurst parameter. The explanation and further discussion about this results are reported in Section 6.

Example 3 Instead of shaping, here we use a policing algorithm which is nearly the same as our shaper but discards every nonconforming cell instead of delaying it. As a result, the policed cell stream contains less cells but will be conforming with the specified rate r and bucket size M. The parameters r and M were chosen as previously. The IDC plots for the policed cell streams can be seen in Figure 13.

Another interesting question is the structure of the discarded cell stream. Figure 14 represents the IDC curves for this kind of overflow process.

Discussion 3 As the Figures show the IDC curves are almost the same in the case of both the policed cell streams and the overflow (dropped) cell streams with approximately the same Hurst

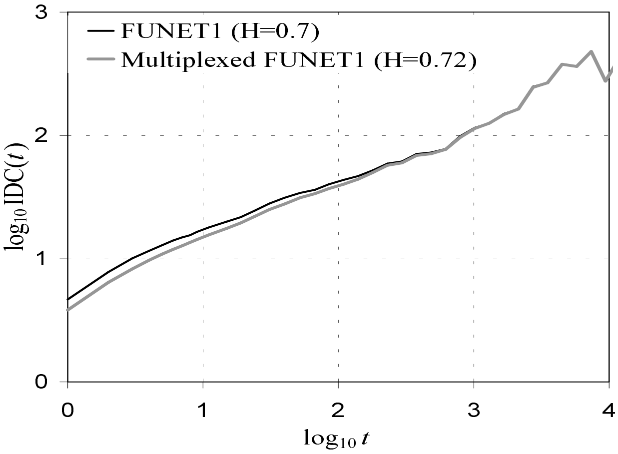

Figure 15: IDC plot for the multiplexed FUNET1 data.

parameter. We can conclude that the self-similarity feature is more robust for policing than for shaping. Our results are consistent with the engineering intuition that FIFO queues behave like low-pass filters and the long-range correlations of the self-similar traffic (having power spectra divergence at low frequencies) are not affected.

Example 4 Finally, we examine how the traffic trace is modified when flowing through a multiplexing stage. The FUNET1 data (whose rate is about 9% of the link speed) was multiplexed with Poisson traffic with rate 90% of the link speed. Next, the IDC plot was calculated for those cells of the aggregated cell stream which cells belonged to the original FUNET1 data. The result can be seen on Figure 15.

Discussion 4 The results show that the IDC curve is only slightly modified by multiplexing a Poisson traffic with even a high load compared to our measured traffic. This also indicates the robust nature of the detected feature of self-similarity.

6 DISCUSSION ON THE HURST PARAMETER

In this section we discuss the problem of interpretation of estimated Hurst parameter based on our investigations.

In our shaping example we found that the long-range dependence feature persisted but the estimated Hurst parameter value was different and the direction of change in this value was very interesting. *The estimated Hurst parameter is increased due to shaping.* This can be explained as follows. On short time scales there is effect of the shaper and it smoothes the cell stream. That is why the variance of the number of cells in a given window is decreasing on short time scales 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 Figure 12. Therefore, it is obvious that the estimated Hurst parameter will be higher. This result is contrary to the usual interpretation of estimated Hurst parameter because the Hurst parameter is believed to be a measure of burstiness. For Poisson traffic, which is a smooth process, it is 0.5 and when increasing burstiness the Hurst parameter is increasing. However, our example shows that if we are smoothing the traffic the Hurst parameter is increasing! So how can we interpret the Hurst parameter? If the process is a pure self-similar process there is a good interpretation of the Hurst parameter, see e.g. [24]. However, it is not obvious what the interpretation of the Hurst parameter is in practice where the traffic structure is modified by several mechanisms (shaping, queueing, multiplexing, etc.) and the process is not a pure self-similar process. Can it be used as a burstiness measure? Our investigations suggest the answer to be negative. Can we gain any information from the estimated Hurst parameter? Can we use it for dimensioning purposes?

We believe that the Hurst parameter could represent an important and compressed information (the degree of self-similarity) about a pure self-similar process but in practice when several effects modify the structure of the traffic they may distort the Hurst parameter so much that its original meaning is hidden. *It could also happen that there is no useful information that we could gain from the estimated and distorted Hurst parameter.* In that case more characteristics are needed to describe the complex traffic and the appealing fractal characterization with only a few parameters will not be appropriate.

7 CONCLUSION

In this paper we have considered several questions of fractal characterization of measured traffic. We presented our ATM traffic measurements with their self-similarity analysis and the problems of testing self-similarity.

We investigated the estimation of the self-similarity parameter and found that the estimated Hurst parameter depends on many characteristics including the estimation technique, sample size, time scale, level shifts, correlation structure. These results highlight the problem of 'how can we get the correct value for the Hurst parameter in practice?'. We also investigated the problem of 'deceiving self-similar tests' and found that several effects e.g. level shifting can result in high values of Hurst parameter in the case of non self-similar processes.

We investigated the problem of interpretation of the estimated Hurst parameter and found that in practice one should be very careful useing the Hurst parameter because it may be distorted and may fail to give a useful characterization for real traffic. We also showed that the Hurst parameter cannot even be used as a burstiness measure in practice. Based on our investigations the main message of our paper is that it has been demonstrated that in practice the Hurst parameter alone gives very little useful information about the traffic. Our research motivates further and deeper research of the problem and indicates that the practical applicability of fractal traffic characterization is not as simple as it was believed so far.

In fact, there are various open issues and living discussions in teletraffic research about the findings of fractal properties and still there is no definite conclusion about whether the most promising modeling approach is to consider the observed traffic being self-similar *in nature* or not. Future research will answer this exciting question.

References

[1] J. Beran, R. Sherman, M. S. Taqqu, and W. Willinger. Long-range dependence in variable-bit-rate video traffic. *IEEE Transactions on Communications*, 43(2/3/4):1566–1579, 1995.

[2] COST242 Management Committee. Methods for the performance evaluation and design of broadband multiservice networks (the COST242 final report). Technical report, COST242, June 1996.

[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, and N. O'Connell. Statistical issues raised by the Bellcore data. In *Proc. 11th Teletraffic Symposium*, Cambridge, March 1994.

[6] A. Erramilli, J. Gordon, and W. Willinger. Applications of fractals in engineering for realistic traffic processes. In *Proc. 14th ITC*, pages 35–44, Antibes Juan-les-Pins, France, 1994.

[7] A. Erramilli and R. P. Singh. Application of deterministic chaotic maps to model packet traffic in broadband networks. In *Proc. 7th ITC Specialists Seminar*, pages 8.1.1–8.1.3., Morristown, NJ, 1990.

[8] A. Erramilli et al. Chaotic maps as models of packet traffic. In *Proc. 14th ITC*, pages 329–338, Antibes Juan-les-Pins, France, 1994.

[9] H. J. Fowler and W. E. Leland. Local area network traffic characteristics, with implications for broadband network congestion management. *IEEE Journal on Selected Areas of Communications*, 9(7):1139–1149, 1991.

[10] M. W. Garret and W. Willinger. Analysis, modeling and generation of self-similar VBR video traffic. In *Proc. ACM SIGCOMM'94*, pages 269–280, London, UK, 1994.

[11] D. P. Heyman and T. V. Lakshman. What are the implications of long-range dependence for VBR-video traffic engineering? *IEEE/ACM Transactions on Networking*, 4(3):301–317, June 1996.

[12] J. Jormakka. On self-similar models for ATM-traffic. In *Proc. ITC Specialists Seminar on Control in Communications*, pages 277–288, Lund, Sweden, September 1996.

[13] W-C. Lau, A. Erramilli, J. L. Wang, and W. Willinger. Self-similar traffic parameter estimation: A semi-parametric periodogram-based algorithm. 1995.

[14] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traffic. In *ACM SIGCOMM'93*, pages 183–193, San Francisco, CA, USA, September 1993.

[15] W. E. Leland, M. S. Taqqu, W. Willinger, and D. W. Wilson. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2(1):1–15, February 1994.

[16] K. Lindberg. The FUNET network connection to foreign countries. *CSC News*, 7(4):19, December 1995.

[17] B. B. Mandelbrot. Long-run linearity, locally Gaussian processes, H-spectra and infinite variances. *International Economic Review*, 10:54–60, 1994.

[18] S. Molnár and A. Vidács. On modeling and shaping self-similar ATM traffic. In *Proc. 15th International Teletraffic Congress*, Washington, D.C., USA, June 1997.

[19] S. Molnár, A. Vidács, and A. Hegedűs. Meeting a challenge: Modeling self-similar LAN/MAN traffic. In *Proc. 8th IEEE Workshop on Local and Metropolitan Area Networks*, Berlin/Potsdam, Germany, August 1996.

[20] S. Molnár and A. Vidács and A. A. Nilsson. Bottlenecks on the way towards fractal characterization of network traffic: Estimation and interpretation of the Hurst parameter. In *Proc. 15th Int. Conference on the Performance and Management of Complex Communication Networks—PMCCN'97*, pages 125–144, Tsukuba, Japan, November 1997.

[21] A. L. Neidhardt and A. Erramilli. Shaping and policing of fractal traffic. In *Proc. ITC Specialists Seminar on Control in Communications*, pages 253–264, Lund, Sweden, September 1996.

[22] I. Norros. Studies on a model for connectionless traffic, based on fractional Brownian motion. Technical Report 242TD(92)041, COST, 1992. Presented also in June 1993 at the Conference on Applied Probability in Engineering, Computer and Communication Sciences, Paris.

[23] I. Norros. A storage model with self-similar input. *Queueing Systems*, 16:387–396, 1994.

[24] I. Norros. On the use of fractional Brownian motion in the theory of connectionless networks. *IEEE Journal on Selected Areas in Communications*, 13(6):953–962, August 1995.

[25] B. K. Ryu and A. Elwalid. The importance of long-range dependence of VBR video traffic in ATM traffic engineering: Myths and realities. In *Proc. SIGCOMM'96*, pages 3–14, CA, USA, August 1996.

[26] W. Willinger, M. S. Taqqu, W. E. Leland, and D. V. Wilson. Self-similarity in high-speed packet traffic: Analysis and modeling of Ethernet traffic measurements. *Statistical Science*, 10(1):67–85, 1995.

[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 *Proc. of the ACM/SIGCOMM'95*, 1995.