Scaling Analysis of IP Traffic Components

Sándor Molnár, Trang Dinh Dang

High Speed Networks Laboratory
Department of Telecommunications and Telematics
Budapest University of Technology and Economics
H–1117, Pázmány Péter sétány 1/D, Budapest, Hungary
Tel: (361) 463 3889, Fax: (361) 463 3107
E-mail: {molnar,trang}@ttt-atm.ttt.bme.hu

Abstract

In this paper a comprehensive scaling analysis of IP traffic and its components is presented. We focus on the correlation and scaling behavior of IP traffic components on both transport and application layers. It is shown that the correlation structure of the aggregation is mainly determined by the component with the highest variance and correlations at the investigated time scale. It is also demonstrated that WAN traffic can exhibit complex multifractal structure even at large time scales. The aim of the paper is to understand the characteristics of the aggregated IP traffic by analyzing its individual components.

1. Introduction

Significant research has been carried out recently to understand the *scaling phenomena* in network traffic [19, 18, 15, 7, 6, 9, 14, 16, 11]. This is in close relationship to traffic burstiness because bursts should be defined in terms of time scales over which clustering activities occur. The surprising scaling phenomenon observed in data traffic is that these clustering activities are present over several time scales [19]. This phenomenon triggered a new modeling approach, called *fractal modeling*, which can offer parsimonious models (e.g. self-similar or long-range dependent (LRD) traffic models) to capture this behavior [19, 10, 12, 11].

A number of studies have reported that aggregated LAN traffic is consistent with exact self-similarity and aggregated WAN traffic is asymptotically self-similar (long-range dependent) [13, 7]. Moreover, it has been found that the scaling structure in measured WAN traffic can be categorized into two main regions: a large time scaling phenomenon with self-similarity and a small time scale phenomenon with multifractal scaling [7, 6, 9]. It has been also pointed out that the transition from the multifractal to self-similar scaling occurs around time scales of a typical packet round-trip time in the network [7, 6]. The physical explanations and engineering implications are also addressed in several papers [6, 16].

However, most of the studies have investigated the aggregated WAN or LAN traffic and only a few papers have examined the nature of IP traffic components [4]. We think that the understanding of the characteristics of the individual components in the aggregation, i.e. a comprehensive analysis of different protocol layers with distinct traffic components is vital for establishing a correct physical understanding and modeling methodology.

In this paper we present a scaling analysis of IP traffic focusing on the characteristics of the components in the aggregation on both transport and application layers and also investigate their impacts on the characteristics of the aggregation. The analysis is based on a series of ten-hour traffic measurements. The analyzed traffic originates from flows from the Internet to our University campus representing various IP applications. The scaling structure of the components and the aggregation on large time scales, i.e. over a few seconds have been studied.

The objective of this paper to contribute to the characterization of IP traffic, especially to improve the understanding of the nature of IP traffic components and their influence on the aggregated IP traffic.

The paper is organized as follows. The details of our measurements are described in Section 2. A global view about the structure of the measured IP traffic concerning the bandwidth share of transport and application layer protocols is presented in Section 3. The characteristics of the chosen interval to be analyzed based on a stationary and trend analysis study are also described in this section. In Section 4 the results of our scaling analysis are shown based on LRD tests (variance-time plots, R/S plots, periodograms) and also scaling analysis tools (wavelet-based methods and partition function method). Finally, the paper is concluded by Section 5 with a summary of our main findings.

2. IP traffic measurements

Figure 1 presents our traffic measurements setup. A number of LANs located at the Informatics Building of the Budapest University of Technology and Economics are connected to the outside world by a 100MB FDDI and a 155MB ATM link. These networks are composed of several Ethernet based LANs which are referred to as Department Groups (DG) and each DG consists of about 100 workstations. Connections between DGs and between a DG and the outer world are guaranteed by an ATM backbone. Ethernet frames are transmitted over the ATM backbone using LAN emulation. Workstations belong to staff members, PhD students, and student laboratories using a variety of operating systems and network interfaces ranging from 10Base2 (BNC) through 100BaseT (UTP) to 100VGAnyLAN. In the measurements both the incoming and outgoing traffic of the DG1 have been collected but in this paper the aggregated incoming IP traffic mainly generated by a WAN environment is analyzed. The traffic monitoring tool, called Captie [17], captured the IP traffic on the ATM link between DG1 and the switch. The analyzer can monitor the data traffic transmitted over the link and record different statistics of data flows according to user requests.

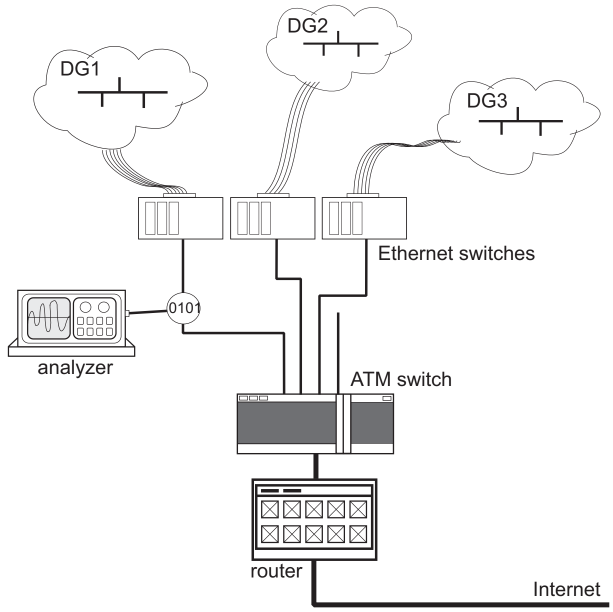

Figure 1: The configuration of the IP traffic measurements

During weekdays of April 1999 statistics of IP and non-IP traffic were continuously collected in log files every day from 8am to 6pm. Our measurements concerned the traffic volume transferred under IP packets, transport layer protocols such as TCP, UDP, ICMP, and OSPF, as well as application layer protocols such as HTTP, FTP, SSH, and SMTP. Data series representing the traffic measured in bytes per second used in most of our analyzes are gained from these measurements.

3. Structure of the IP traffic

3.1. Overview of the traffic

Figure 2 shows the traffic intensity in bytes per second for a typical IP traffic flow and its main components from our measured traces. Traffic bursts can be observed over the whole period of the measurement. It can be concluded from Figure 2 and our more detailed investigations that huge peaks over a short period can occur at any time of the day so the well-known busy period concept from telephone traffic engineering cannot be applied. Visual investigation suggests that at the transport layer the TCP data is dominant and its behavior determines the characteristics of the IP traffic.

Besides the TCP traffic there is a considerable amount of UDP traffic as well but this sort of traffic is much smoother than the TCP and it does not seem to have any influence on the IP traffic nature. At application layer the HTTP and the FTPdata traffic take the main roles (we differentiated the FTPdata traffic from the FTPcontrol traffic because FTP control messages are transported over the IP network by separate packets) . However, the FTPdata traffic behavior is rather complex since the most of the FTPdata connections are transferred at low speed but traffic intensity also contains some extremely high jumps during short periods. The size of these jumps is often $10^3 - 10^5$ times greater than the typical FTPdata connection speed. Therefore the empirical autocorrelation function computed from a finite series with such a peak may give misleading results. However, the analysis of a short time period of the FTPdata traffic (in order to validate the assumption about stationarity) may not give us the complete characteristics of this traffic type.

3.2. Bandwidth share

Figure 3 and Figure 4 present the ratio of different protocols contributing to the overall load at the transport layer and the application layer, respectively. These results are the average values of several measurement days.

The majority of data is carried by the TCP protocol at the transport layer which takes about 90% bandwidth of the whole volume of transferred data. The rest of the load mainly corresponds to the UDP protocol. The ICMP and OSPF control messages share only 1-2% of the transport layer traffic volume. On the other hand, based on the results in Figure 4 we can conclude that among the applications the HTTP and the FTP traffic are dominant in volume.

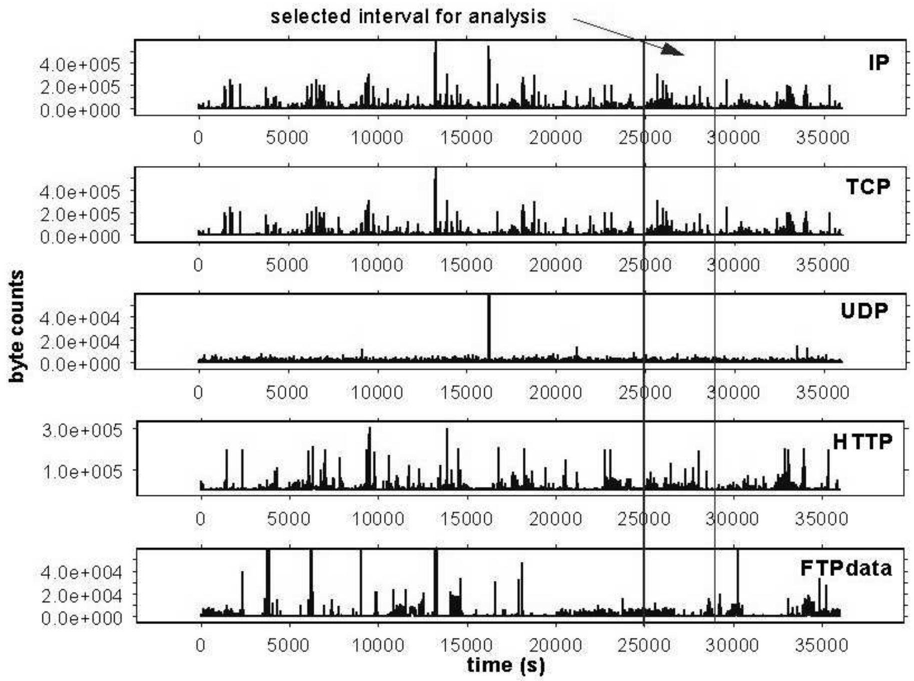

Figure 2: Traffic intensity of IP and some higher layer protocols

3.3. Stationarity analysis

An important assumption in traffic modeling is stationarity. However, it is rather difficult to justify completely this assumption on the investigated data series [3, 11]. The real traffic data over a longer time period often appear to have local trends, load jumps, cycles, etc., which are the characteristics of non-stationary processes. Thus a straightforward approach to overcome this problem is to select time periods where the stationarity seems to be acceptable.

A simple test to detect stationary periods in the data is to slide a window along the measured data and investigate the variations of the data averages in the windows. The datagram of this series may give some information about level-shift, trends, etc. However, for bursty data like our measured traffic this method does not provide appreciable results. To achieve this goal we also applied another tool which is based on a *change point detection method* [5]. The main idea of this method is to slide a window along the data

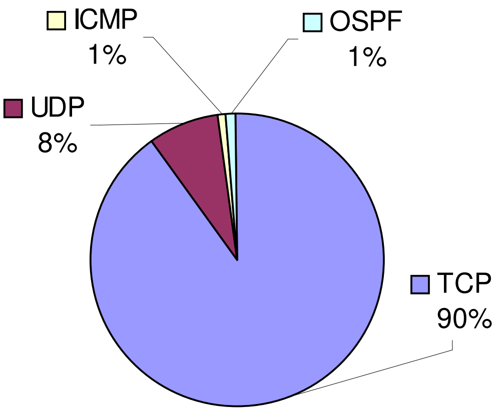

Figure 3: Bandwidth share of transport layer protocols

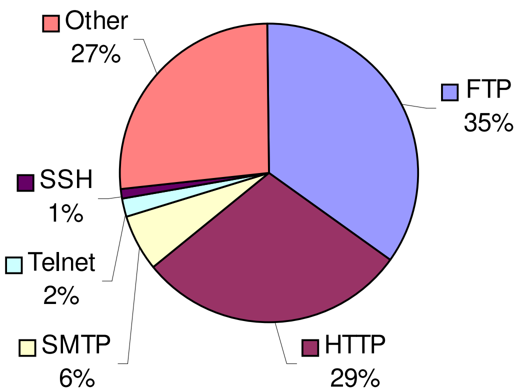

Figure 4: Bandwidth share of application layer protocols

series and then compare the distribution of data samples in two equal halves of the window. If the two distributions are significantly different then the assumption of the stationarity covered by the window is rejected. The comparison task of distributions of the two series with equal size is performed by applying the Kolmogorov-Smirnov test [5].

Based on our stationarity tests we have selected several subsets from the whole measured data for analysis. The subsets are obtained from time intervals where the IP traffic and also each component of both transport layer protocol traffic and application layer protocol traffic can be justified to be stationary.

In this paper we present the analysis of a one hour trace chosen from a collection of selected series by our test (see the selected part in Figure 2).

4. Correlation and scaling analysis

In this Section we present the characteristics of the investigated IP traffic and its components. First, we study the autocorrelation function and also discuss how the different component autocorrelation functions produce their results in the aggregation. Second, the long-term scaling is investigated and our LRD test results are presented and discussed. Third, the detailed scaling of each component of the IP traffic is analyzed.

4.1. Correlation structure

There are a number different components with different contribution effects on the correlation structure of the aggregated IP traffic. The identification of the characteristics of the components which mainly determine the characteristics of the aggregated IP correlation structure is vital for the understanding of IP traffic structure.

Consider the superposition of a number of independent traffic streams, i.e. $A = \sum_{i=1}^{N} A_i$. Denote the autocorrelation function of A_i by $r_{A_i}(k)$. It can be shown that the autocorrelation of the aggregated traffic stream is given by

$$r_A(k) = \frac{1}{\sum_{i=1}^{N} \sigma_{A_i}^2} \sum_{i=1}^{N} \sigma_{A_i}^2 r_{A_i}(k) \qquad (1)$$

where $\sigma_{A_i}^2$ is the variance of the traffic volume in the chosen time unit for stream i. As $k \to \infty$ the autocorrelations of short-range dependent (SRD) streams vanish rapidly and the autocorrelations of LRD streams decay asymptotically as $k^{-\beta_i}$. The autocorrelation of A is determined by the LRD stream which decays at the lowest rate, i.e. $r_A(k) \sim k^{-\min \beta_i}$. Therefore the LRD stream with the highest H parameter will dominate ($\beta = 2 - 2H$) and the aggregation will be LRD with this parameter. However, in practice we investigate k for large values instead of infinity. For this case we can also consider the variance of the streams because the variance is the weight in the sum in Eq. 1. Consequently it may happen that there is a stream with a faster decaying autocorrelation function but with a high variance and this stream will dominate in the autocorrelation function of the aggregated traffic stream on the investigated time scale. Moreover, it also follows from Eq. 1 that the volume of traffic has no influence, so a small fraction of traffic with high variance and slowly decaying autocorrelation can determine the autocorrelation of the whole aggregation.

Observe now the above discussed properties in our measured IP traffic. We present in Figure 5 the sample autocorrelation functions of different measured traffic flows. In the evaluation of these functions (and especially investigating correlation coefficients at large lags) we can observe that the correlation coefficients often have small values. In these cases we have to take care of the confidence interval which can be roughly estimated by the $\pm 2/\sqrt{n}$ rule corresponding to significance 0.05. In our cases we used 3600 samples which gives about 0.03 for this confidence interval. We can observe a slow decay in correlation of the IP traffic which indicates a possible presence of LRD. The transport layer protocols, TCP, UDP, ICMP, and OSPF, work above the IP layer so the IP traffic is the aggregation of these flows. Among these components the TCP traffic takes the dominant role since the form of its correlation absolutely determines the correlation structure of the IP. This is so because the TCP series has sample variance which is much greater than that of the other transport traffic series (see in Table 1). This observation is in accordance with Eq. 1 and our discussions above. *Note that TCP has a significant impact on the correlation structure of the IP aggregation because it has the highest variability on the investigated time scale and not because it has the highest bandwidth share in the IP aggregation.* Besides the TCP we also observed a possible long-term decay of the UDP. The ICMP and the OSPF seem to be SRD. However, the final conclusions about correlation structure are stated just after detailed investigations of LRD presented in the next subsection.

At the application layer traffic carried by the HTTP, FTP, SMTP, and the Telnet protocols is considered. All of these traffic flows are components of the TCP traffic aggregation. Figure 5 shows that the autocorrelation functions of HTTP, FTPdata, FTPcontrol, and Telnet all seem to have a long-term decay. Moreover, we also see that the correlation structure of the TCP inherits the form of the correlation structure of HTTP traffic in spite of the fact that the FTP-data appears to have stronger correlation and slower decay. By investigating the sample variance of these components in Table 1 we see that the HTTP traffic has the greatest variance value, which is at least 60 times greater than the variances of the other traffic flows and this is the reason for its dominance to form the autocorrelation of the TCP aggre-

Traffic type	volume	variance $(.10^5)$	VT plot	R/S	Per.	correlation structure
IP	100%	1951.4	0.72	0.73	0.75	LRD H=0.73
TCP	88.5%	1946.6	0.72	0.73	0.75	LRD H=0.73
UDP	8.87%	4.3	0.67	0.68	0.7	LRD H=0.68
OSPF	0.88%	2.1	-	-	-	SRD
ICMP	1.68%	2.9	0.63	-	-	SRD
TCP	100%	1946.6	0.72	0.73	0.75	LRD H=0.73
HTTP	47.64%	800.4	0.72	0.74	0.72	LRD H=0.73
FTPdata	12.29%	14.6	0.85	0.78	0.86	LRD H=0.85
FTPcontrol	0.26%	0.004	0.72	-	0.74	LRD H=0.72 (artifact)
SMTP	1.69%	5.8	-	-	-	SRD
Telnet	1.99%	1.5	0.7	-	-	SRD
Others	36.13%					

Table 1: Summary of the LRD analysis of the IP based protocols

gation. *Note again that HTTP is the dominant protocol to influence the correlation structure of the TCP aggregation because its high variability on the investigated time scale.* Of course, the correlation structures of these protocols are the results of the interactions between these protocols. For example, the correlation structure of TCP is the joint result of the contributing protocols (mainly HTTP) and the TCP mechanism rather than simply the "HTTP forms the correlation structure of TCP".

4.2. Long-range dependence analysis

The *long-range dependent* (LRD) property of a traffic flow is revealed in the power law decay of the autocorrelation function at large lags, i.e. $r(k) \sim c|k|^{2H-2},\ k \to \infty,\ H \in (0.5, 1)$ and c is a constant [3]. The degree of this slow decay is determined by the Hurst parameter (H). There are several statistical methods for LRD testing and parameter estimation [3, 2]. In our case we choose the variance-time plot, the R/S analysis, and the periodogram plot [3] for this goal and use the Logscale Diagram based on the wavelet transform [2] to verify the results.

Results of our LRD analysis are summarized in Table 1. It should be noted that traffic volumes of the transport layer protocols and of the application layer protocols are compared to the IP and the TCP volume, respectively. The LRD behavior of the IP traffic is determined by the TCP traffic because it has both the largest Hurst parameter and variance among the transport protocols. This finding is in accordance with our discussions in the previous subsection. In the case of general protocols used for network control (ICMP and OSPF) the LRD tests failed. Considering also their correlation structure shown in Figure 5 we can conclude that they are in the class of SRD traffic. UDP was found to be LRD but with smaller variance and Hurst parameter compared to TCP.

In the application layer traffic the HTTP is dominant to form the LRD characteristic of the TCP traffic in spite of the fact that the FTPdata has larger Hurst parameter (see Table 1). As we discussed above this is due to the fact that HTTP has a significantly larger variance compared to FTP. We note that some other analyzed FTPdata subsets do not exhibit LRD property so we still do not make a general conclusion about the correlation structure of this traffic type. We have found that the SMTP is SRD. Concerning the Telnet data series our comprehensive analysis has shown that this traffic is SRD in spite of the fact that the variance-time test suggests LRD. In the case of the FTPcontrol our LRD analysis indicates the presence of LRD with $H = 0.72$. By analyzing the datagram of this series we have found that the occurred long-term dependence is caused by the periodicity in control messages sending. Beyond that, the FTP control messages take only a negligible amount of the whole traffic with small variance. Therefore from a traffic engineering point of view we can disregard the presence of this flow.

4.3. Scaling analysis

In our scaling analysis we looked for multifractal and monofractal (e.g. self-similar) scaling structures. In our terminology a time series $\{X_i,\ i = 1, 2, \ldots, n\}$ is called *multifractal* if the logarithms of the *partition function* $S^m(q)$ (or equivalently the absolute moments) scale linearly with the logarithm of the aggregation level m [18, 8], i.e. $\log S^m(q) = \tau(q) \log m + c_1(q)$ where $S^m(q) = \sum_{k=1}^{n/m} |Z_k^{(m)}|^q$ with $Z_k^{(m)} = \sum_{i=1}^{m} X_{(k-1)m+i}$ and $c_1(q)$ is a constant. Furthermore, if the *scaling function* $\tau(q)$ is a linear function of q we call it *monofractal*. A special and well-known case of monofractals is the *self-similarity* in the case of $\tau(q) = qH - 1$ where H is the self-similarity parameter (Hurst parameter) [8].

Scaling behavior can also be tested by wavelet-based methods [1]. The discrete wavelet transform (DWT) represents a data series X of size n at a scaling level j by a set of wavelet coefficients $d_X(j,k),\ k = 1,2,\dots,n_j$, where $n_j = 2^{-j}n$. Define the q^{th} *order Logscale Diagram* (q-LD) by the log-linear graph of the estimated q^{th} moment $\mu_j(q) = 1/n_j \sum_{k=1}^{n_j} |d_X(j,k)|^q$ against the octave j. Linearity of the LDs at a different moment order q suggests the scaling property of the series, i.e. $\log_2 \mu_j(q) = j\alpha(q) + c_2(q)$ where $\alpha(q)$ is the *scaling exponent* and $c_2(q)$ is a constant. The plot of $\alpha(q)$ against q can reveal the type of scaling [1].

In our analysis both the partial function method and the wavelet-based method were used. In [15] it was found that in some cases multifractal scaling is more convincing without subtracting the mean from the time series because the centered data has several disadvantages, see [15]. We have performed scaling analysis both with and without subtracting the mean but in our case no significant differences were found in the results. In this paper the results are related to the centered data because it is the case when we have any hopes to find possible self-similar scaling [18]. A summary of our scaling analysis results are shown in Table 2

Traffic type	**volume**	**scaling structure**
IP	100%	multifractal
TCP	88.5%	multifractal
UDP	8.87%	none
OSPF	0.88%	none
ICMP	1.68%	none
TCP	100%	
HTTP	47.64%	multifractal
FTPdata	12.29%	monofractal $h = 0.74$
SMTP	1.69%	none
Telnet	1.99%	none
Others	36.13%	

Table 2: Summary of the scaling analysis of the IP based protocols

The results are presented in Figure 6. Concerning the LD of moment order $q = 2$ in Figure 6(a) a nearly linear interval of the LD plot at octaves $1 \le j \le 5$ can be observed[1]. The larger values of j were not considered because of the limited size of the data set and also because the set of the wavelet coefficients at large scaling levels contains only a few values, which cannot give a reliable approximation. (These considerations are also taken into account in scaling analysis of other flows.) A linear regression to the interval gives an estimation of LRD parameter $H = 0.76$ with confidence interval $(0.73, 0.8)$. The result deviates slightly from estimates of H provided by other LRD tests in Table 1 but the confidence interval still includes those values.

LDs of the IP traffic computed at different q provide the estimation of the scaling exponent $\alpha(q)$ presented in Figure 6(b). The non-linear curve of the scaling function suggests that IP traffic has a multifractal structure on these time scales. The renormalized partition functions of the IP traffic are depicted in Figure 6(c). (Renormalization was performed in order to have the same intercepts of the curves for all q.) Our estimation of the scaling function $\tau(q)$ in Figure 6(d) based on the partition functions shown in Figure 6(c) also confirms our findings concerning multifractal scaling.

Investigating data series from the transport layer protocols our scaling analysis showed that although the UDP traffic may have LRD property as discussed but the scaling tests of UDP failed. We have found that the OSPF and the ICMP traffic flows also do not have scaling structure. In the case of the TCP we have found that its scaling structure is similar to the scaling structure of IP traffic. The estimated $\alpha(q)$ and $\tau(q)$ shown in Figure 7 rather resemble these functions of the IP traffic. We can conclude that the TCP traffic also exhibits multifractal scaling.

We have also analyzed the scaling structure of application layer protocols. Scaling analysis results of the HTTP can be seen in Figure 8. Both the scaling exponent $\alpha(q)$ and the scaling function $\tau(q)$ have convex curves which suggest the presence of multifractality. Moreover, the estimated $\tau(q)$ is nearly the same as in the case of the IP and the TCP traffic. Analysis results of the FTPdata are shown in Figure 9. It can be seen that both the scaling exponent $\alpha(q)$ and the scaling function $\tau(q)$ are linear functions of q. Therefore, our selected FTPdata set reveals clear evidence of monofractality. However, according to our earlier discussions the identified monofractal structure is not our general conclusion about the characteristics of the traffic carried by the FTP protocol. Finally, we have found that the SMTP and the Telnet traffic do not exhibit a scaling structure.

Our results confirm the results presented in [7, 6, 9] showing that WAN traffic is LRD. However, we also complement these findings by demonstrating that *WAN traffic can exhibit a complex multifractal structure not only at small but also at large time scales*. Moreover, the analysis revealed that *the aggregation is composed of components with very different scaling behavior* (no scaling, multifractal scaling, monofractal scaling).

5. Conclusions

In this paper we have presented traffic analysis results from a wide range of IP traffic measurements. It has been found that the IP traffic is bursty during the whole day and sometimes it also contains extremely high traffic peaks over a short time period. These observations are questioning the concept of busy period in case of this kind of IP traffic.

[1] In this LD plot an improved estimation was used ($y_j = \log \mu_j + g_j$) with a correcting factor g_j described in [1]

We have carried out a stationarity analysis prior to our correlation and scaling studies and selected a time period of the IP traffic in which the aggregated IP traffic and all of its analyzed components at both transport and application layers can be considered to be stationary.

We have investigated the impacts of different characteristics of the correlation structures of the components on the correlation structure of the aggregation. We have found that the ruling impact is due to the component which has the highest variance and also significant correlations on the investigated time scale. Among the transport layer and the application layer protocols the TCP and the HTTP were found to have these properties, respectively, which is the main reason (and not the high bandwidth share!) of the dominance of these protocols.

At the transport layer both TCP and UDP exhibit LRD but only the TCP has detectable multifractal structure. Other protocols at this layer were found to be SRD with no scaling properties.

At the application layer both HTTP and FTP have LRD properties. Multifractal and monofractal scaling have been identified for HTTP and FTP, respectively. However, we observed that the scaling of FTP traffic is not a general characteristic.

We conclude that the investigated IP traffic is a LRD aggregate of components with different scaling properties resulting in a complex multifractal structure for the aggregated WAN traffic even at large time scales.

Acknowledgement

Authors would like to thank István Maricza and Attila Vidács for their comments on the paper, András Gefferth and Zoltán Pusztai for their contributions during the measurements and Darryl Veitch for the code of the wavelet-based method. The research was supported by the Inter-University Centre of Telecommunications and Informatics (ETIK).

References

[1] P. Abry, P. Flandrin, M.S. Taqqu, and D. Veitch, "Wavelets for the analysis, estimation and synthesis of scaling data", to appear in *Self-Similar Network Traffic Analysis and Performance Evaluation*, K. Park and W. Willinger, Eds., 1999.

[2] P. Abry and D. Veitch, "Wavelet analysis of long range dependent traffic", *IEEE Trans. Inform. Theory*, 44(1):2–15, January 1998.

[3] J. Beran, *Statistics for LongMemory Processes*, Chapman & Hall, One Penn Plaza, New York, NY 10119, 1995.

[4] V. Bolotin, J. Coombs-Reyes, D. Heyman, Y. Levy, and D. Liu, "IP traffic characterization for planning and control", In P. Key and D. Smith, editors, *Proc. ITC 16*, pages 425–436. Elsevier Science B. V., 1999.

[5] T. Éltető and S. Molnár, "On the distribution of round–trip delays in TCP/IP networks", In The 24th Annual Conference on Local Computer Networks (*LCN'99*), Lowell, Boston, MA, USA, October 1999.

[6] A. Feldmann, A.C. Gilbert, and W. Willinger, "Data Networks as Cascades: Investigating the Multifractal Nature of Internet WAN Traffic", *ACM Computer Communication Review*, 28:42–55, September 1998.

[7] A. Feldmann, A.C. Gilbert, W. Willinger, and T.G. Kurtz, "The Changing Nature of Network Traffic: Scaling Phenomena", *ACM Computer Communication Review*, 28:5–29, April 1998.

[8] A. Fisher, L. Calvet, and B.B. Mandelbrot, "Multifractality of Deutschmark/US Dollar Exchanges Rates", Yale University, 1997. Working Paper.

[9] A.C. Gilbert, W. Willinger, and A. Feldmann, "Scaling analysis of conservative cascades, with applications to network traffic", *IEEE Trans. Inform. Theory*, 45(3):971–991, April 1999.

[10] D.L. Jagerman, B. Melamed, and W. Willinger, "Stochastic Modeling of Traffic Processes", In *Frontiers in Queueing*, pages 271–370. CRC Press, 1997.

[11] S. Molnár and I. Maricza eds, "Source Characterization in Broadband Networks", Interim report, COST 257, Vilamoura, Portugal, January 1999.

[12] S. Molnár and A. Vidács, "On Modeling and Shaping Self-Similar ATM Traffic", In 15th International Teletraffic Congress, Washington, DC, USA, June 1997.

[13] V. Paxson and S. Floyd, "Wide-area traffic: The failure of Poisson modeling", *IEEE/ACM Transactions on Networking*, 3(3), June 1995.

[14] R.H. Riedi, M.S. Crouse, V.J. Ribeiro, and R.G. Baraniuk, "A multifractal wavelet model with application to network traffic", *IEEE Trans. Inform. Theory*, 45(3):992–1018, April 1999.

[15] R.H. Riedi and J. Lévy–Véhel, "Multifractal properties of TCP traffic: a numerical study", INRIA research report 3129, Rice University, February 1997.

[16] R.H. Riedi and W. Willinger, "Toward an improved understanding of network traffic dynamics", SelfSimilar Network Trafic and Performance Evaluation, Wiley, June 1999.

[17] H. Saito, T. Tsuchia, Gy. Marosi, Gy. Horváth, P. Tatai, and S. Asano, "Real–time Simulation with Traffic Monitoring Tool", In *ICCCN99*, Natick-Boston, MA, USA, October 11-13, 1999.

[18] M.S. Taqqu, V. Teverovsky, and W. Willinger, "Is network traffic self-similar or multifractal?", *Fractals*, 5:63–73, 1997.

[19] W. Willinger, M.S. Taqqu, and A. Erramilli, "A Bibliographical Guide to Self-Similar Traffic and Performance Modeling for Modern High-Speed Networks Stochastic Networks: Theory and Applications", In *Royal Statistical Society Lecture Notes Series*, volume 4. Oxford University Press, 1996.

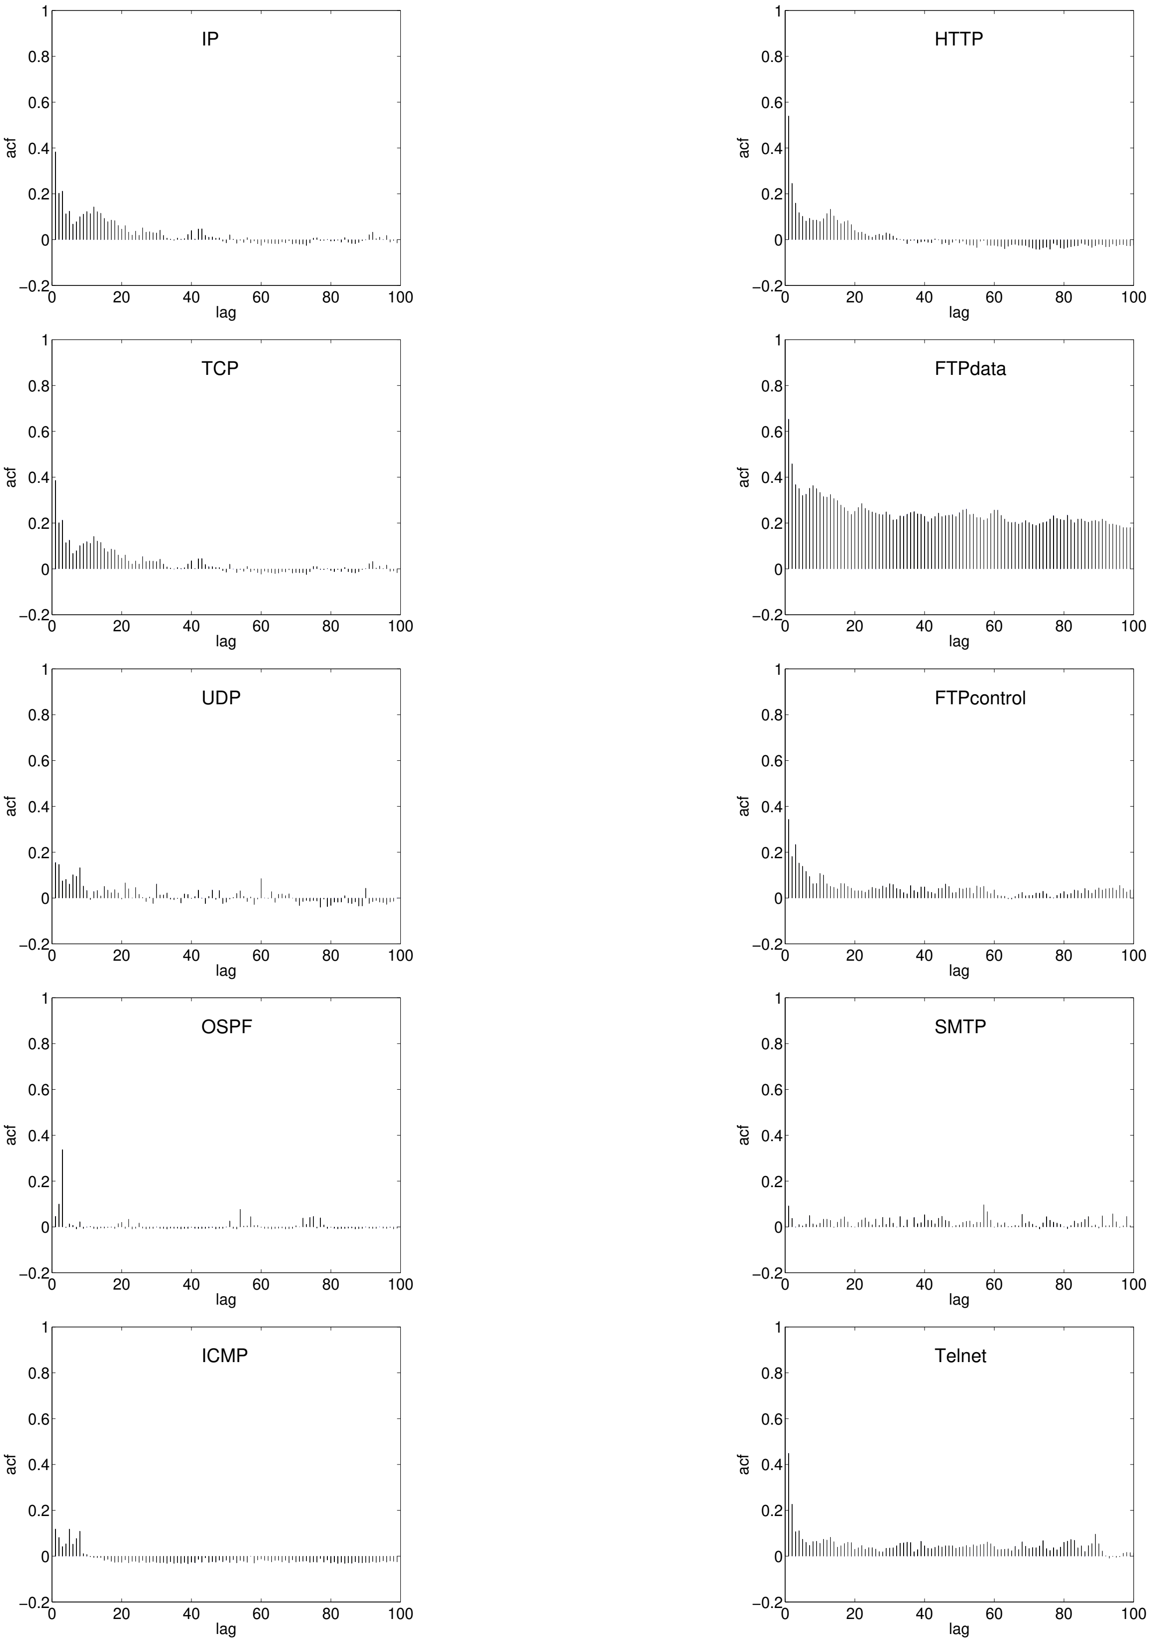

Figure 5: Correlation structure of IP traffic and its components

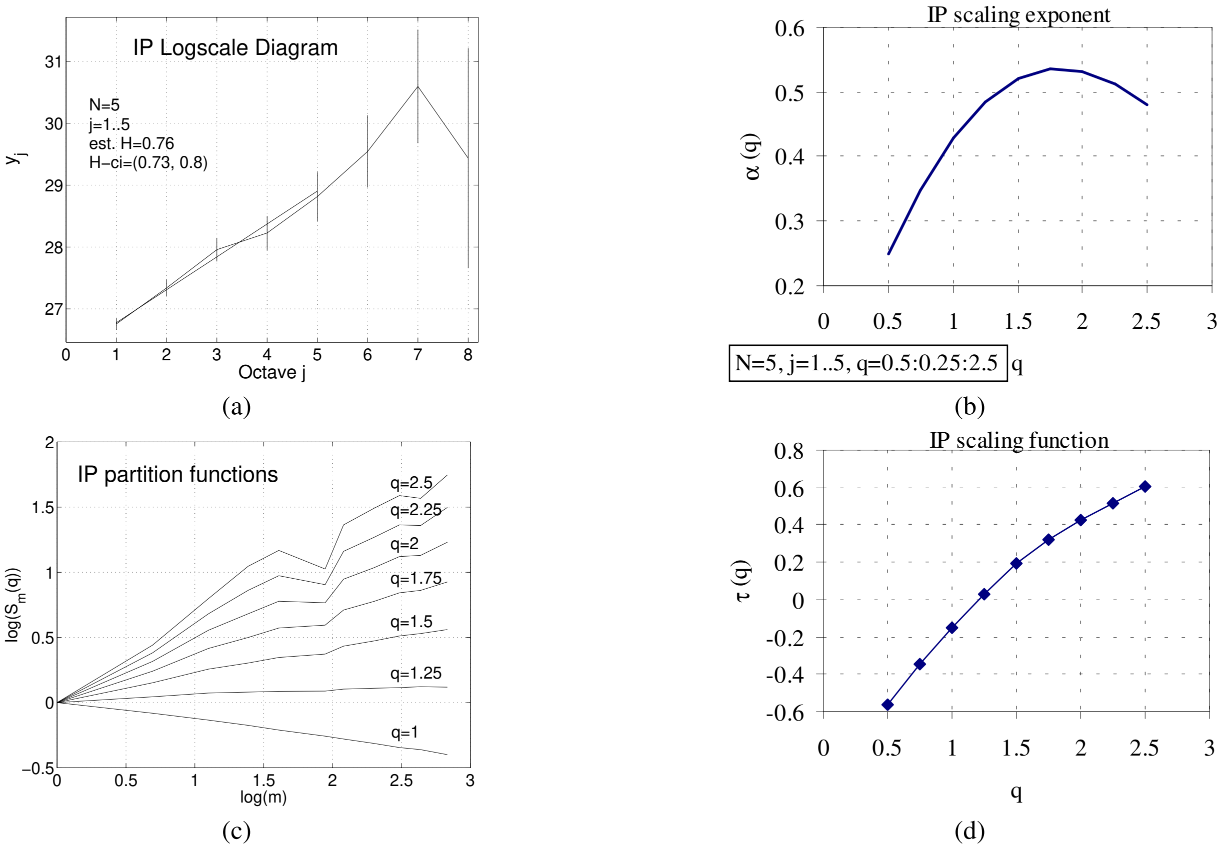

Figure 6: Scaling analysis results of the IP traffic

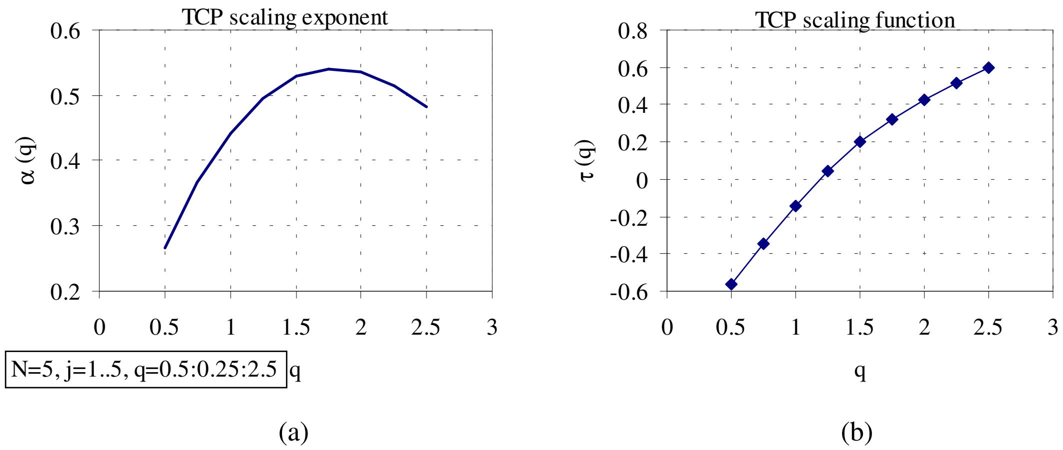

Figure 7: Scaling analysis results of the TCP traffic

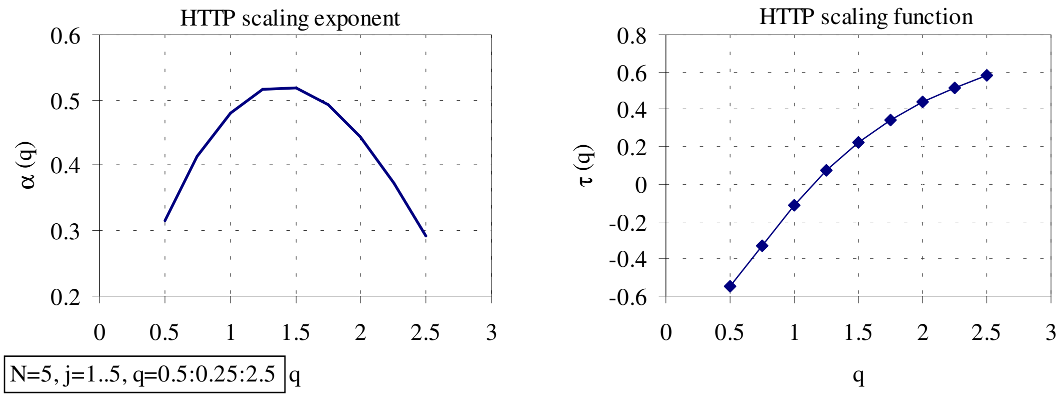

(a) (b)

Figure 8: Scaling analysis results of the HTTP traffic

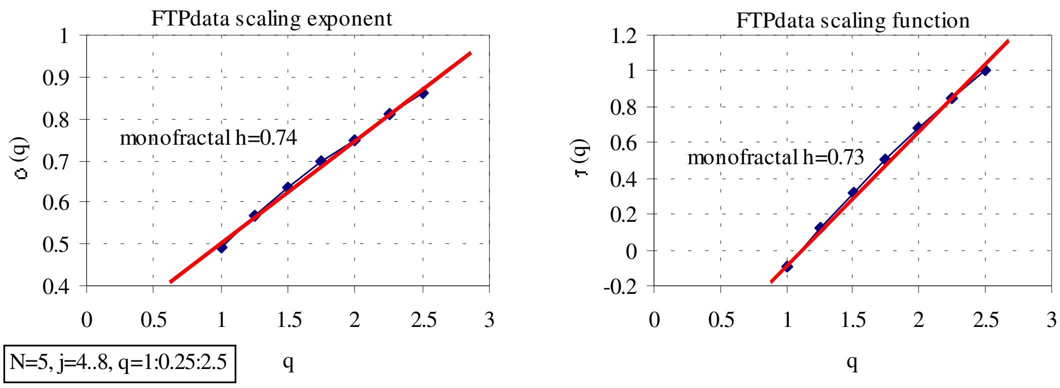

(a) (b)

Figure 9: Scaling analysis results of the FTPdata traffic