On the Statistical Nature of Highly-Aggregated Internet Traffic

Miguel de Vega[1], Salvatore Spadaro[2], Marie-Ange Remiche[1],
Davide Careglio[2], Javier Barrantes[2], and Jürgen Götz[3]

[1] Université Libre de Bruxelles, SMG - CP 210/01
Blvd du Triomphe, 1050 Bruxelles, Belgium
{mdevegar, mremiche}@ulb.ac.be
[2] Advanced Broadband Communications Center (CCABA)
Universitat Politècnica de Catalunya, 08034 Barcelona, Spain
{sspadaro, careglio, jbarranp}@ac.upc.edu
[3] Siemens AG, Corporate Technology
Otto-Hahn-Ring 6, 81739 München, Germany
juergen.goetz@siemens.com

Abstract. We present a complete statistical analysis of highly-aggregated Internet traffic traces captured with a tailor-made measurement platform designed to operate at gigabit speed with ns-precision. Our main result is that the traces are not Poisson but present nontrivial scaling behavior.

1 Introduction

There is an important debate in the literature regarding the statistical nature of highly aggregated traffic in the Internet. On one hand many papers have reported the existence of scaling properties in traffic measured at the network level [1–3]. Similar scaling phenomena has been detected at other levels as well (e.g. transport and application levels) and it implies a deviation from the Poisson traffic model. This deviation is not trivial in the sense that it has a noneligible impact on the network performance like for instance, on the buffer dynamics and on the blocking probability [3, 4].

On the other hand, papers such as [5–7] maintain that the existence of scaling properties depends on the level of aggregation of traffic. As we move from access to core networks this level increases and so traffic should resemble more and more to a Poisson process. Indeed, in the field of optical networks the assumption of Poisson traffic is a very common practice (see for instance [8] and references therein).

Theoretically, in order to find out if the Poisson assumption is correct, one could simply perform a statistical analysis of network traffic in a high-capacity optical backbone link. In the praxis, two main difficulties arise when trying to accomplish such task. First, due to confidentiality issues it is difficult for the research community to access such information, owned in most cases by private operators. Second, the efficient measurement of traffic at Gbps speeds is a very

challenging technical task. Hardware limitations often reduce the precision of the packet time-stamps to μs (e.g. in [7]). Software limitations introduce artifacts in the measured data changing its statistical properties (e.g. TCPdump).

In the framework of the collaboration agreement between Advanced Communications Broadband Center (CCABA) and Supercomputing Center of Catalonia (CESCA), CCABA has created SMARTxAC [9]. It is a tailor-made measurement platform designed to operate at gigabit speed without packet losses and a ns-precision in the packet time-stamp measurements. In this paper, we provide a complete analysis of packet arrival times measured at the optical link interconnecting the Catalan R&D network (about 50 Universities and Research Centers) with the Spanish R&D network and the global Internet.

The paper is structured as follows. We begin in Section 2 with a description of SMARTxAC. In Section 3.1 we use two well-known statistical tests from [10] and apply them in order to test the Poisson hypothesis in the traces. Section 3.2 presents a detailed analysis of the scaling properties of the traffic, in particular, the self-similar and multifractal hypothesis are tested. In Section 4 we compare several traffic models in two different network scenarios in order to gain a quantitative understanding of their ability to capture the essential information from the traces. Section 5 presents the conclusions from this work.

2 Traffic Measurements

In this section we briefly describe the measurement platform shown in Figure 1. The measurement location is the Anella Científica, which is the Catalan R&D network, managed by CESCA and connects about 50 Universities and Research Centers in Catalonia. The point of measurement is a pair of Full-Duplex Gigabit Ethernet links (two per each traffic direction) that connects the Anella Científica to RedIRIS (the Spanish R&D network) and to the global Internet. The measurement is performed at a tapped link using optical splitters.
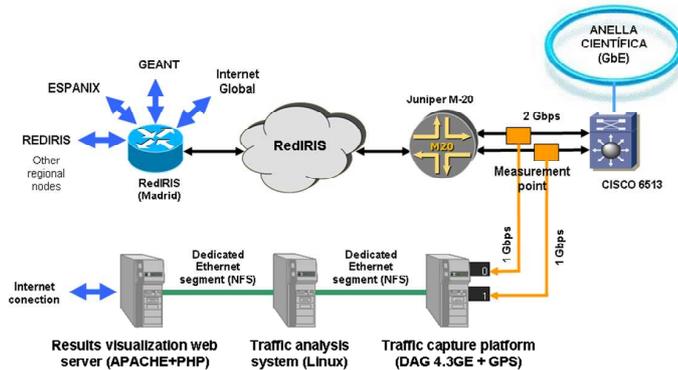

Fig. 1. Measurement platform

SMARTxAC is a passive system able to perform at gigabit speeds without packet losses and ns-precision in the packet time stamp measurements. It integrates the capture engine as well as the real-time flow traffic measurement and analysis in the same software. This software runs in a single machine which is equipped with one or more Endace DAG 4.3GE Gigabit Ethernet measurement cards. In order to analyze all the traffic in real-time, only packet headers are captured and aggregated into flows, which reduces the data volume to be processed and stored by the traffic analysis system. Also, it is possible to configure the platform to capture certain ranges of a frame or to collect high time resolution packet traces, since precision time stamping is performed on the capture device, which is synchronized via GPS.

3 Statistical Analysis of the Traces

In this section we consider the measured packet arrival times to form a sample path of a given stochastic process and investigate its nature by means of different statistical tests.

3.1 Testing the Poisson Hypothesis

The Poisson hypothesis can be tested by means of various approaches, such as the computation of the autocorrelation function of the packet interarrival time sequence (see [3, 5, 7]). Few papers have provided formal statistical tests, such as the Box-Ljung statistic in [7]. In this section we introduce two simple and powerful statistical tests [10] in order to test the Poisson hypothesis.

A Poisson process is a particular case of renewal process in which the interarrival time distribution is exponential. We use this fact in order to indirectly test the Poisson hypothesis using the null hypothesis of a renewal process. For that purpose we use the Lewis-Robinson (LR) and the Pair-wise Comparison Nonparametric Test (PCNT) [10].

Performing the LR and the PCNT tests with an uplink (UL) and downlink (DL) sequence of 2^{23} arrival times (between 2 and 3 minutes of measurement time) and a confidence interval of 5% lead to a rejection of the renewal process hypothesis. This implies a rejection of the Poisson process hypothesis. In order to compare our results with the analysis of the OC48 traces in [7], we rounded our traces from a precision in the packet time-stamps of ns to a precision of μs. Then our tests agreed with those in [7] in accepting the independence assumption for the packet interarrival times.

These LR and PCNT tests are defined as follows. Let $T = \{T_1, \cdots, T_N\}$ be the sequence of N packet arrival times and $X = \{T_2 - T_1, \cdots, T_N - T_{N-1}\}$ the sequence of interarrival times in the trace. We refer to the quantile $z_{\alpha/2}$ as the value for which $P[Z > z_{\alpha/2}] = \frac{\alpha}{2}$, where Z is distributed according to a standard normal distribution, and α represents the confidence interval. The null hypothesis H_0 in both tests is that the observed traces of packet interarrivals form a renewal process.

The Lewis-Robinson Test Under the hypothesis of a Poisson process and conditioning on T_N, the arrival times $\{T_1, \cdots, T_{N-1}\}$, are uniformly distributed on $(0, T_N)$. Let

$$U_L = \frac{\sum_{i=1}^{N-1} T_i - (N-1)\frac{T_N}{2}}{T_N \sqrt{\frac{N-1}{12}}} \;. \tag{1}$$

The Lewis-Robinson statistic U_{LR} is equal to U_L/CV. CV is the coefficient of variation that can be estimated as $CV = \frac{\sqrt{\widehat{\sigma}_X^2}}{\overline{X}}$, where $\widehat{\sigma}_X^2$ and $\overline{X}$ represent the variance and average estimators, respectively. The test criterion is to reject H_0 if $U_{LR} \notin [-z_{\alpha/2}, z_{\alpha/2}]$.

The Pair-wise Comparison Nonparametric Test Let U count the number of times that $X_j > X_i$ for $j > i$ and for all i. Under H_0 the mean value of U is $E[U] = \frac{N(N-1)}{4}$ and its variance can be estimated as $Var[U] = \frac{(2N+5)(N-1)N}{72}$. According to the central limit theorem for large N, U should be approximately distributed as a normal distribution with mean $E[U]$ and variance $Var[U]$. Thus, the statistic:

$$U_{PCNT} = \frac{U - E[U]}{\sqrt{Var[U]}} \;. \tag{2}$$

should be approximately distributed as a standard normal distribution. Therefore, the test criterion is to reject H_0 if $U_{PCNT} \notin [-z_{\alpha/2}, z_{\alpha/2}]$.

3.2 Estimation of the Scaling Properties

We wish now to find a stochastic process that captures better than the Poisson process the statistical nature of the traces. The natural choice at this point [4, 2, 3, 11] is to check if the traces possess scaling behavior. To do this we focus on large and small time scales in order to detect the presence of self-similar and multifractal behavior, respectively.

Self-Similarity A continuous time stochastic process $\{Y(t); t \in \mathbb{R}\}$ is self-similar with Hurst parameter H (in short H-ss), if for all $a > 0$ and $t \geq 0$, and some $0 < H < 1$,

$$Y(t) =_d a^{-H} Y(at) \;. \tag{3}$$

where $=_d$ represents equivalence in the sense of finite dimensional distributions. In the traffic modeling context this definition has the drawback of implying that $Y(t)$ is nonstationary [12]. For this reason it is often desirable to work with the increment process $X(t) = Y(t) - Y(t-1)$ of $Y(t)$. If this process is stationary then $Y(t)$ is referred to as a self similar process with stationary increments (H-sssi). Assume now second-order stationarity and the existence of the second moment of $X(t)$, then the process $\{X(t); t \in \mathbb{Z}\}$ is exactly second-order self-similar with Hurst parameter H ($1/2 < H < 1$) if for all $k \geq 1$:

$$\gamma(k) = \frac{\sigma^2}{2}((k+1)^{2H} - 2k^{2H} + (k-1)^{2H}) \;, \tag{4}$$

where $\gamma(k)$ represents the autocovariance function of $X(t)$. Otherwise $\{X(t)\}$ is asymptotically second-order self-similar if:

$$\lim_{m \to \infty} \gamma^{(m)}(k) = \frac{\sigma^2}{2}((k+1)^{2H} - 2k^{2H} + (k-1)^{2H}) \,, \tag{5}$$

where $\gamma^{(m)}(k)$ is the autocovariance function of the aggregated process $X^{(m)}(i) = \frac{1}{m} \sum_{t=m(i-1)+1}^{mi} X(t)$. In both cases the process $X(t)$ is also said to be long-range dependent (LRD), since the sum of Equation(4) or (5) over k does not converge (see [13, 14] for a more precise definition). This is a direct implication of the hyperbolic decaying form of these equations.

Let us assume that the measured sequence X is part of a sample path of a finite-variance second-order stationary process $X(t)$. From X we wish to estimate the Hurst parameter H, and thus study the possible second-order self-similarity and LRD of $X(t)$. We begin this study with the estimator proposed in [13] based on the discrete wavelet transform (DWT). We now provide an intuitive description of this transform, and refer to [13] (and references therein) for a more formal one.

With the DWT a time series X is successively projected into a series of approximation subspaces V_j. We denote by $aprox_j(t) = (Proj_{V_j} X)(t)$ to the approximation of X in V_j. The subspaces V_j satisfy a series of properties [13], one of them being that $V_j \subset V_{j-1}$. This implies that $aprox_j(t)$ is a coarser (i.e. less detailed) approximation of X than is $approx_{j-1}(t)$, and provides the interpretation of j as a *scale* parameter. As in the map of a geographical region, a time series viewed at a coarser scale j (e.g. scale 1:50000 in the map) contains less information than at a finer scale $j-1$ (e.g. scale 1:25000 in the map). The information which is lost when going from $approx_{j-1}(t)$ of X to $approx_j(t)$ is captured by the so-called detail coefficients $\{d_X(j,k), k \in \mathbb{Z}\}$. They can be computed by comparing with inner products the time series X with the functions $\varphi_{j,k}$, which are shifted (i.e. parameter k) and dilated (i.e. parameter j) templates of a certain function φ_0 called the dual mother wavelet [13]. That is, $\{d_X(j,k) = \langle X, \varphi_{j,k} \rangle, j = 1, \cdots, J, k \in \mathbb{Z}\}$.

The Logscale Diagram (LD) is basically a log-log plot of variance estimates of the series $\{d_X(j,k), k \in \mathbb{Z}\}$ across different scales j, complete with the confidence intervals for these estimates (for more details see [13]). LRD and second-order self-similarity can be detected if from a certain lower cutoff scale j_1 the LD aligns within the confidence intervals up to the largest scale J present in data. If $j_1 = 1$ exact second-order self-similarity is the most reasonable choice. Otherwise, $j_1 \neq 1$ suggests asymptotic second-order self-similarity. Figures 2.a to 2.c present the LD for the DL, UL and BC (Bellcore traces, search for *BC-pAug89* in [15]), in which 24 scales are considered for the DL and UL traces. In DL and UL alignment is detected for large time scales suggesting asymptotic second-order self-similarity. Alignment begins much later than in the BC trace, but it does not seem to degenerate into a flat curve, which would indicate the absence of LRD and self-similarity.

Once detected the presence of asymptotic second-order self-similarity we proceed to the estimation of the corresponding Hurst parameter with a number of well-known estimators from the literature (see [14] and references therein). Table 1 summarizes the results obtained. The estimated Hurst parameter is lower for the DL and UL traces compared to the BC one. However, one has to be careful when interpreting such results. Due to the asymptotic nature of the self-similar phenomenon, not considering a sufficiently high number of scales in the analysis, or considering the lower timescales may lead to biased estimations [16]. In our case, due to computational constraints, the length of the DL and UL traces used with all the estimators was of 2^{22}, except for the $A\&V$ estimator, for which 24 scales were used (and which provides a higher Hurst estimation). Since the asymptotic self-similar behavior begins at scales which are quite high (above scale 15, see Figures 2.a and 2.b), eliminating the last two scales in the rest of the estimators might account for the small values of H observed. This highlights the need for computationally-efficient estimation algorithms in order to estimate the Hurst parameter in highly-aggregated traffic. Another possible interpretation is that the process $X(t)$ is non-stationary. This interpretation drastically reduces the possibilities to model the packet arrival process. Finally, another interpretation is that although the underlying process is (asymptotic second-order) self-similar and LRD, its Hurst parameter is quite close to 0.5 (i.e. absence of LRD), indicating a low degree of LRD.

Table 1. Hurst parameter estimation for the interarrival time series (A&V: Wavelet Abry-Veitch estimator [13], Local W.: Local Whittle, C.I.: confidence interval)

H	Whittle	C.I.	A&V	C.I.	Local W.	Periodogram	Variance	R/S
DL	0.5207	[0.5202, 0.5214]	0.770	[0.713, 0.827]	0.62	0.63	0.59	0.55
UL	0.5501	[0.5495, 0.5507]	0.796	[0.739, 0.853]	0.54	0.65	0.58	0.54
BC	0.6426	[0.6408, 0.6443]	0.752	[0.728, 0.777]	0.77	0.67	0.77	0.73

The Multifractal Behavior In this section we look for scaling in the higher-order moments of the detail DWT coefficient series $\{d_X(j,k), k \in \mathbb{Z}\}$ of the type $S_q(j) = E[|d_X(j,k)|^q] \sim C_q j^{\alpha_q}$, where $S_q(j)$ are referred to as the partition functions [17]. For a H-ss process the function α_q exhibits a simple linear scaling with q: $\alpha_q = Hq + q/2$. The Hurst parameter controls each scaling exponent and thus it constitutes an example of a monofractal process. It is common to define the function ζ_q as [17, 11]:

$$\zeta_q = \alpha_q - q/2 \ . \tag{6}$$

which gives the simple relation $\zeta_q = Hq$ for H-ss processes. If Equation(6) is not linear with q then one says that the process has nontrivial multifractal scaling. In a generalization of the LD one may look for alignment in a log-log plot of estimates of the partition functions $S_q(j)$ across different scales j for different
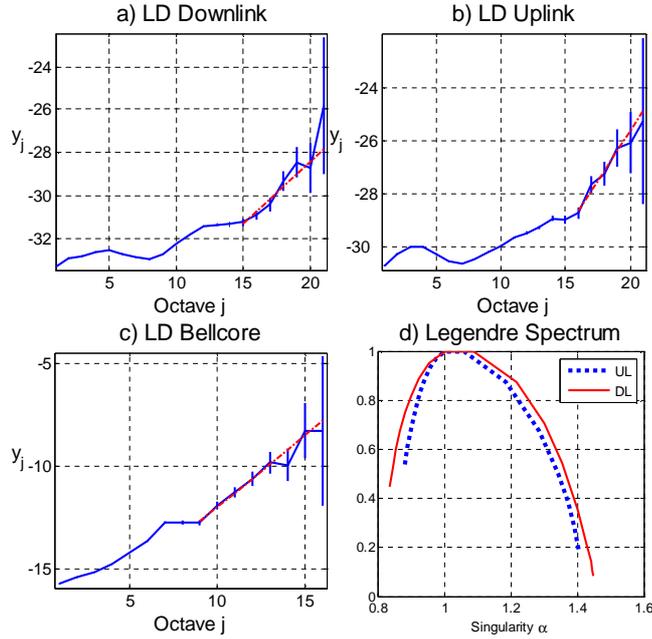
Fig. 2. Logscale Diagram for the: a) DL, b) UL and c) Bellcore traces. d) Estimation of the Legendre spectrum for the DL and UL traces.

values of q. The slope of the linear regression for each q provides an estimation for α_q (i.e. of ζ_q). The Multiscale Diagram (MD) [13] is a plot of the estimation of ζ_q vs. q, together with the corresponding confidence intervals. A lack of alignment within the confidence intervals in the MD suggests nontrivial multifractal scaling (i.e. not linear with q). The Linear Multiscale Diagram (LMD) is a statistically equivalent variant of the MD in which one plots $h_q = \zeta_q/q$ vs. q. The LMD is more convenient for visual inspection: nontrivial multifractal scaling is manifested as a nonalignment (within the confidence intervals) of the LMD curve with a horizontal line. Figures 3.a to 3.c present the LMD for the DL, UL and a realization of a fractional Gaussian noise (fGn) process (see [12]) fitted to the DL trace, respectively. As it can be observed from Figures 3.a and 3.b, the DL and UL trace show no alignment, suggesting nontrivial multifractal scaling. However, the fGn synthesis does show alignment (see horizontal dashed line in Figure 3.c). This suggests the presence of trivial multifractal scaling, as it should be the case for a (second-order) self-similar process.

The study of a multifractal process $\{X(t), t \in \mathbb{R}\}$ reduces to the study of the *erratic* behavior of each sample path $\{X(t, w), t \in \mathbb{R}, w \in \Omega\}$ locally evaluated around t. This erratic behavior is measured with the so called *singularity exponents* [17]. The Haussdorf spectrum can be thought of as a summary of the multifractal properties of a time series. It basically gives the fractal (Haussdorf) dimension of the set of points having the same singularity exponent α, as a

function of α. The Legendre spectrum is less rich but it is usually used as it is far more numerically accessible. Figure 2.d plots an estimation of the Legendre spectrum computed as in [17] for the DL and UL traces. The wide bell-shaped curve is an indicator of the presence of nontrivial multifractal scaling.
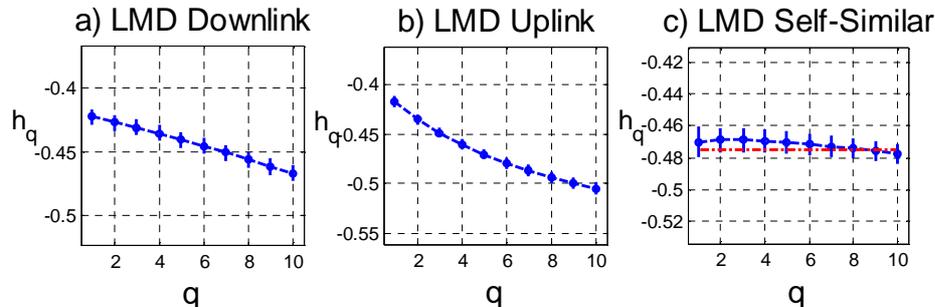
Fig. 3. Linear multiscale diagram for: a) DL, b) UL and c) fGn synthesis fitted to the DL trace.

4   Performance Comparison of Different Traffic Models

In the previous section we have concluded that the traffic traces possess scaling properties. In this section we wish to measure the impact of such deviations from the Poisson assumption from the performance point of view.

We focus on the performance of three traffic models in two different scenarios. We chose the multifractal wavelet model (MWM) [11] as a representative of a traffic model capable of capturing nontrivial multifractal scaling. As the second model we use a variant of the MWM capable of generating (asymptotic second-order) self-similar traffic (SIM). As the third model we consider the Poisson model in order to evaluate its usefulness as an approximate but simple traffic model.

We chose two different scenarios in order to benchmark the different traffic models: a single-server infinite-buffer node, and a multi-server bufferless node. In the first scenario we are interested in the queuing behavior. This scenario may be representative for instance for opaque SDH (SONET) networks. In the second scenario our aim is to measure the blocking probability of the different traces. This scenario may be representative for instance for Optical Burst Switching (OBS) networks.

All three traffic models (MWM, SIM and Poisson) are fitted to the original UL and DL traces, and are used in order to generate synthetic traffic traces representing packet arrival times. These traces are considered in a simulative study together with the original ones (UL and DL) in order to compare their behavior in terms the above mentioned performance parameters. We want to

focus on the ability of the three models to capture the relevant information from the original UL and DL packet arrival time sequences. Thus, we use *iid* exponentially distributed packet sizes in all simulation runs for all traffic traces (including the original DL and UL) in order to eliminate the effect of possible cross-correlations between the packet arrival and packet size sequences. The figures in the following sections present the results averaged for 10 simulation runs of 2^{24} samples each, and with a link load of 50%.

Figures 4.a and 4.b provide the results concerning the blocking probability of the traces from the different models in a bufferless multiserver node. As it can be observed, models which incorporate scaling (MWM and SIM) constitute a good approximation for the blocking probability. The use of the Poisson model implies nonnegligible approximation errors in most practical situations. For instance, for an OBS node with 6 wavelengths (a wavelength can be interpreted as a server) the error is of 1 order of magnitude, and it grows nonlinearly with increasing number of wavelengths.

Figures 4.c and 4.d illustrate the results concerning the buffer occupation level. As it can be observed the Poisson model underestimates this performance measure, while the SIM model overestimates it. The MWM shows a good match. Regarding the SIM curve, similar results have been observed using a different SIM generator based on fGn sample paths modified under a power transformation to make all interarrival values positive. We have observed that for high link loads (i.e. above 40%) the SIM model overestimates the results, while for low loads it underestimates them. Analogous results have been reported in [4].

5   Conclusion

We have addressed the question of whether the Poisson model constitutes a good approximation for highly-aggregated Internet backbone traffic. Thanks to the efficiency of the SMARTxAC platform, we have been able to analyze an unprecedent amount of accurate (packet time-stamps with ns-precision) traffic measurements.

Using simple and powerful statistical tests we conclude that the UL and DL traces are not Poisson with a 5% confidence interval.

We have investigated the presence of scaling behavior in the traces. Evidence suggests nontrivial multifractal scaling and when assuming stationarity, asymptotically second-order self-similarity.

In order to assess the impact on the network performance of the fact that scaling is present in our traces, we conducted a series of simulations to compare different traffic models. The multifractal model resulted the most accurate one. The Poisson model provided a poor approximation of the results. Its approximation error increases rapidly with the size of the buffer and with the number of servers in a node.

According to these results, one should use traffic models which incorporate the concept of scaling in order to analyze highly-aggregated Internet backbone
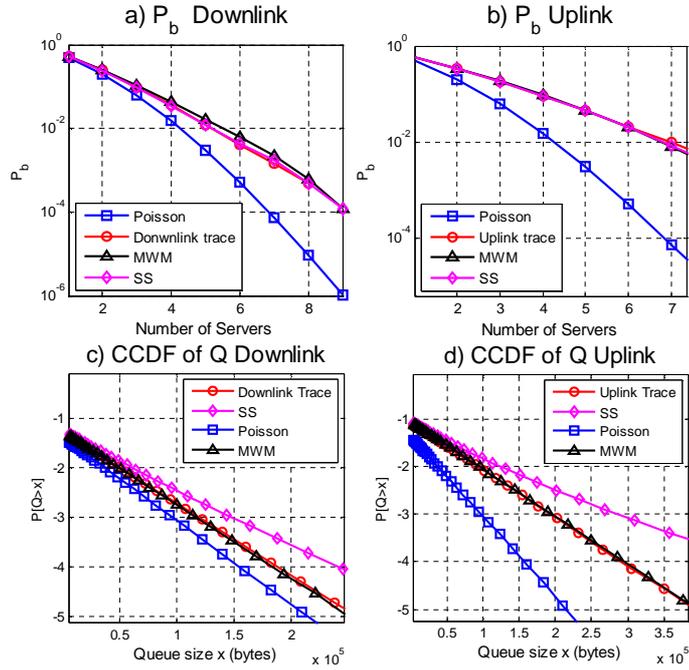
Fig. 4. Blocking and buffer overflow probabilities for the UL and DL traces.

traffic. It remains as an interesting open question the study of the reasons why scaling is still present in highly-aggregated traffic traces.

Acknowledgements

References

1. Leland, W., Taqqu, M., Willinger, W., Wilson, D.:On the self-similar nature of Ethernet traffic, in Proc. of SIGCOMM (1993) 183–193
2. Feldmann, A., Gilbert, A., Willinger, W.: Data Networks as Cascades: Investigating the Multifractal Nature of Internet WAN Traffic, in Proc. of SIGCOMM (1998) 42–55
3. Paxson, V., Floyd, S.: Wide-Area Traffic: The Failure of Poisson Modeling, in IEEE/ACM Transactions on Networking, **3:3** (1995) pp. 226–244
4. Gao, J., Rubin, I.:Multiplicative multifractal modeling of Long-Range-Dependent Network Traffic, Int. J. Commun. Syst., **14** (2001) 783–801
5. Cao, J., Cleveland, W., Lin, D., Sun, D.: Internet traffic tends toward Poisson and independent as the load increases, in Nonlinear Estimation and Classification, Eds. New York, NY: Springer Verlag (2002)

6. Cao, J., Ramanan, K.: A Poisson Limit for Buffer Overflow Probabilities, in Proc. of IEEE INFOCOM (2002)
7. Karagiannis, T., Molle, M., Faloutsos, M., Broido, A.: On the Nonstationarity of Internet Traffic, in Proc. ACM SIGMETRICS **1** (2001) 102–112
8. Hülsermann, R., Bodamer, S., Barry, M., Betker, A., Gauger, C., Jger, M., Khn, M., Spth. J.: A Set of Typical Transport Network Scenarios for Network Modelling. In Proc. of ITG-Fachtagung Photonische Netze (2004)
9. SMARTxAC measurement platform,
   http://www.ccaba.upc.edu/contingut.php?dir=Projects/SMARTxAC
10. Harold, A., Feingold, H.: Repairable Systems Reliability: Misconceptions and Their Causes. Marcel Dekker, Inc., (1984)
11. Riedi, R., Crouse, M., Ribeiro, V., Baraniuk, R.: A Multifractal Wavelet Model with Application to Network Traffic, IEEE Transactions on Information Theory, **45:2** (1999) 992–1018
12. Embrechts, P., Maejima, M.: An introduction to the theory of selfsimilar stochastic processes. International Journal of Modern Physics **14**,(2000) 1399–1420
13. Abry, P., Flandrin, P., Taqqu, M. S., Veitch, D.:Wavelets for the Analysis, Estimation, and Synthesis of Scaling Data, in Self-Similar Network Traffic and Performance Evaluation, John Wiley & Sons, Inc., (2002) 39–88
14. Karagiannis, T., Faloutsos, M., Riedi, R.: Long-Range Dependence: Now you see it now you don't!, in Global Internet (2002)
15. *http://www.ist-mome.org/database/MeasurementData/*
16. Figueiredo, R., Liu, B., Feldmann, A., Misra, V., Towsley, F., Willinger, W.: On TCP and self-similar traffic. Perform. Eval. **61(2-3)**(2005) 129–141
17. Rudolf, R., Vehel, L.:Multifractal Properties of TCP Traffic: A Numerical Study, Technical Report Nr. RR-3129 (1997), *http://citeseer.ist.psu.edu/riedi97multifractal.html*