A Simple Markovian Approach to Model Internet Traffic at Edge Routers

L. Muscariello, M. Mellia, M. Meo, R. Lo Cigno, M. Ajmone Marsan
Dipartimento di Elettronica, Politecnico di Torino,
Corso Duca degli Abruzzi 24, Torino, Italy
email: {<last_name>}@mail.tlc.polito.it

Abstract

In this paper, we propose a simple MMPP (Markov Modulated Poisson Process) traffic model that approximates the LRD (Long Range Dependence) characteristics of traffic traces measured at our institution edge router, at both the flow and packet levels.

The MMPP model mimics the real behavior behind the interaction between users, protocols and the network, using the notion of sessions and flows, therefore resulting in a simple and intuitive model. The queueing behavior of the traffic generated by the model is coherent with the one of the measured traces at several different traffic loads.

While the model is not intended to offer an explanation of the reasons why Internet traffic is LRD, it does offer a simple and manageable tool for dimensioning and planning networks (link and buffer capacities), since the characteristics of the generated traffic are easily controlled through the model input parameters.

I. Introduction

The flows of packets arriving at Internet routers (edge or backbone) cannot be accurately modeled by Poisson processes; this is no news at all, and a vast literature (see [1], [2], [3], [4], [5], to cite just a few) discusses this fact. However, the quest to find convincing explanations for this behavior, as well as simple and accurate models to describe it, is far from concluded, although many different works have explored several possible solutions, as quickly discussed in Sect. II.

One of the main characteristics of Internet traffic, probably the one with a major impact on planning and dimensioning, is the Long Range Dependence (LRD) of several parameters (e.g., packet inter-arrival time, amount of data transferred per unit time, etc.). It is thus unquestionable that Internet traffic does have some sort of *memory*, which is at the basis of LRD, although the origins of this memory are not yet completely clarified. It should however be noted that correlation properties, tail distributions, and all the other characteristics of Internet traffic, are meaningful only over a limited range of time scales. For instance, any correlation property on time scales smaller than the packet transmission time has no physical meaning. Similarly, tails of distributions describing file lengths, become meaningless beyond the limitations imposed by storage media. Therefore, since asymptotic behaviors cannot impact practical results, an Internet traffic model should have a finite scope of application.

In general, traffic models can be useful for two different purposes. The first one, which we do not discuss in this paper, is finding *physical* explanations of the properties displayed by the traffic, and observed in measured data, i.e., explain *why* these properties appear. The second purpose, far less ambitious, but nevertheless useful and interesting, is finding stochastic processes that can be used for the traffic description in analytical models of the network behavior.

This work was supported by the Italian Ministry for Education, University and Research through the TANGO and PlanetIP projects.

Such models are particularly appealing if they are simple, easy to understand, and, above all, if they can be controlled through a small number of parameters, whose influence on the model output is predictable, at least from a qualitative point of view.

One of the simplest and most frequently used approaches to introduce memory into a stochastic process is the use of *Markov Modulated Poisson Processes* (MMPPs). In this paper we propose an MMPP model of Internet traffic, that matches very well the characteristic of the Internet traffic exiting from the Internet through an Edge Router. The MMPP is based on the notions of sessions, flows and packets, typical of real network applications. Compared to other MMPP models [25], our approach tries to mimic the real behavior behind the interaction between users, protocols, and the network. The model has five parameters, that can be mapped directly on measured traffic characteristics, such as the average load, the average packet flow size, etc., and therefore is very simple to tune. Moreover, these same parameters directly influence the output model in a controlled and foreseeable way, so that the MMPP model can be used to predict the network behavior as the traffic characteristics change.

The rest of this paper is organized as follow. Sect. II shortly discusses the literature that influenced our research. Sect. III presents the MMPP model and its ties to measurable traffic characteristics, as well as to some widely accepted Internet features. Sect. IV describes the measured data we use for traffic analysis and for the model validation. Sect. V discusses the performance of the model, highlighting its strengths, but also discussing its limitations. Sect. VI finally concludes the paper.

II. Related Work

In the early 90's, two seminal papers [1], [2] showed that traffic traces captured on both LANs and WANs exhibit long range dependence (LRD) properties, and self-similar characteristics at different time scales. Those discoveries spurred a significant research effort to understand data traffic in packet networks in general, and in the Internet in particular. In addition, the evidence of LRD and of self-similar properties in packet traffic drove many researchers to abandon the usual Markovian assumptions in favor of newer and more complex traffic models. A number of attempts were made to develop models for LRD data traffic. Here we briefly summarize some of the main approaches proposed in the literature.

Looking at packet traffic as a superposition of source-destination traffic flows, simple *ON/OFF models* (or *packet trains models*) were proposed (see for example [3], [4]). If the *ON* (*OFF*, or both) period durations are generated according to heavy-tailed distributions, and the number of multiplexed flows is large, then the resulting aggregate traffic exhibits self-similar properties. Moreover, since the heavy-tailed distribution of file sizes was measured on storage devices, these two properties have often been considered as the origin of the self-similar traffic behavior of data traffic. Other recent studies [6], [7] indicate instead that traffic properties are rooted in the TCP congestion control mechanism, which induces LRD properties in the aggregate traffic resulting from the superposition of independent sources.

Among the several LRD traffic models proposed in the literature, *Fractional Brownian Motion* (FBM) models received a lot of attention, since their Gaussian nature helps in the study of the queuing behavior [3], [4]. However, these models present a restrictive correlation structure, that fails to capture the short-term correlation of real traffic. *Multifractal models* [8] or 'Cascades' [9] are extensions of those models, and can use wavelet decomposition to capture the LRD properties of traffic [10]. While these models are computationally very efficient, they are complex and difficult to tune, due to the lack of a mapping between the traffic parameters and the model coefficients.

Chaotic map models [11], [12] were proposed as a deterministic evolution of systems governed by a set of behavioral rules. The derived models are simple, but it is often difficult to understand the relationship between the model and real traffic parameters.

FARIMA models [13] are widely used in video trace modeling, and can be used to generate LRD sequences. These models are derived by filtering white Gaussian noise, and are able to capture both the short and the long period correlation of traffic. However, the models are quite complex, and their structure makes it very hard to understand the relationship among the filter coefficients and the real traffic data.

In spite of the many proposed traffic models with LRD characteristics, very little work has appeared in the fields of network design and planning, or network performance analysis, based on LRD traffic models. This is mainly due to the difficulty in handling the complex mathematical structure of the stochastic processes on which those traffic models are based. On the other hand, a recent result showed that the long-term correlation of traffic beyond a certain threshold does not influence the performance of a system [14], so that simple models where correlation is limited (such as MMPP models) can be developed.

The results in [15], [16] also provide support to the possibility of using Markovian traffic models in packet networks, showing that the bandwidth sharing in packet networks is insensitive to both the flow size, and the flow arrival process, under the quite commonly accepted assumption (see also [2], [3]) that session arrivals are Poisson.

III. An MMPP Traffic Model

Today Internet traffic is mainly generated by client-server data transfers that use the TCP protocol at the transport layer. We derive our model by keeping in mind that, in layered architectures, the human actions on a terminal interface cause a sequence of events and behaviors of the protocols at the various layers of the protocol stack. For example, a “click” on a web link, causes the generation of a request at the application level (i.e., an HTTP request), which is translated into many transport level connections (TCP flows); each connection generates a sequence of data segments that are transported by the network through IP packets. According to this view, we identify three different time scales. We define a *flow* as a single TCP connection, started by the three-way-handshake procedure, and ended by the closing procedure. Each flow generates a sequence of *packets*, that, according to the TCP congestion control algorithms are injected into the network. It was shown (see for example [2], [17]) that the arrival processes of both packets and flows exhibit LRD properties. In fact, flows do not start independently; rather, they tend to be generated by *sessions*. The much slower time scale of the session dynamics introduces dependencies, which, at the typical time scales of flows and packets, may be perceived as long-range dependencies.

Figure 1 sketches a realization of the model; three sessions arrive, each one generates a given number of flows, and each flow, in its turn, generates a given number of packets, that will be multiplexed on links along the source-destination paths.

While the definitions of packets and flows are rather precise and well accepted, it is more difficult to provide a specific and unique definition of a session. All the web pages downloaded by a user from the same web server in a limited period of time can form a session; a ftp connection from a user that requests many files from a server can form a session; all the

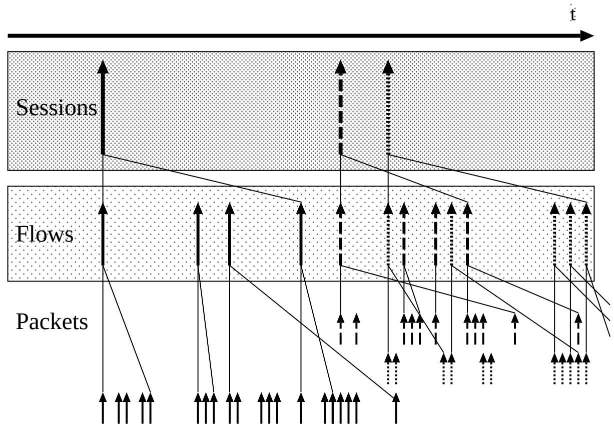

Fig. 1. Sessions, TCP flows, and IP packets

e-mail messages generated by a user that replays to all the previously downloaded e-mails, or even the user activating its connection to the Internet (for example by switching on its computer) are all examples of definitions of a session. Thus, in order to abstract from specific applications, we assume, as a general definition of a session, a set of correlated flows that are submitted to a network interface (not necessarily a *user* interface) in a given time interval.

As was briefly discussed previously, many authors adopt a Poisson arrival process at the session level [2], [15], [16]. This is usually justified by the fact that sessions naturally start after the independent actions taken by a large population of users/applications. At the flow/packet level, more complex models are usually adopted, which take into account think times during which users/network process the received data. Moreover, in order to induce an LRD behavior, think times are often modeled by heavy-tailed distributions, and some kind of correlation is introduced between consecutive flow sizes and/or between inter-arrival times. In this paper, we instead propose a simple and generic model which is based only on Markovian processes, thus requiring no heavy-tailed distribution:

- *Sessions* are generated according to a Poisson process with rate λ_s. Each session starts with the arrival of a new flow. The number of flows generated by a session is a geometrically distributed random variable with mean value equal to N_f. A session closes when it generates the last flow.
- *Flows* belonging to the same session are generated according to a Poisson process with rate λ_f. Each flow starts by generating a packet and ends after having generated a number of packets whose distribution is geometric with mean value N_p.
- *Packets* belonging to the same flow are generated according to a Poisson process with rate λ_p.

Note that, due to the above assumptions, both the packet arrival process and the flow arrival process are MMPP. The Continuous-Time Markov Chain (CTMC) which modulates the packet and flow arrival processes is defined by the state variable $\overline{s} = (n_f, n_s)$, where n_f and n_s denote the number of active flows and the number of active sessions, respectively. A state-transition diagram of the CTMC model of the modulating process is reported in Fig. 2. The transition from state (i, j) to state $(i-1, j)$ corresponds to the termination of a flow which is not the last one of a session; its rate is $i\mu_f$, where $\mu_f = \lambda_p/(N_p - 1)$ and $\beta = 1 - 1/N_f$. Similarly, when the last flow of a session arrives, the session is no more active and the Markov chain moves from state (i, j) to $(i+1, j-1)$; this occurs with rate $j(1-\beta)\lambda_f$.

In state (n_f, n_s), the flow generation rate is equal to $n_s\lambda_f + \lambda_s$; the packet generation rate is equal to $n_f\lambda_p$.

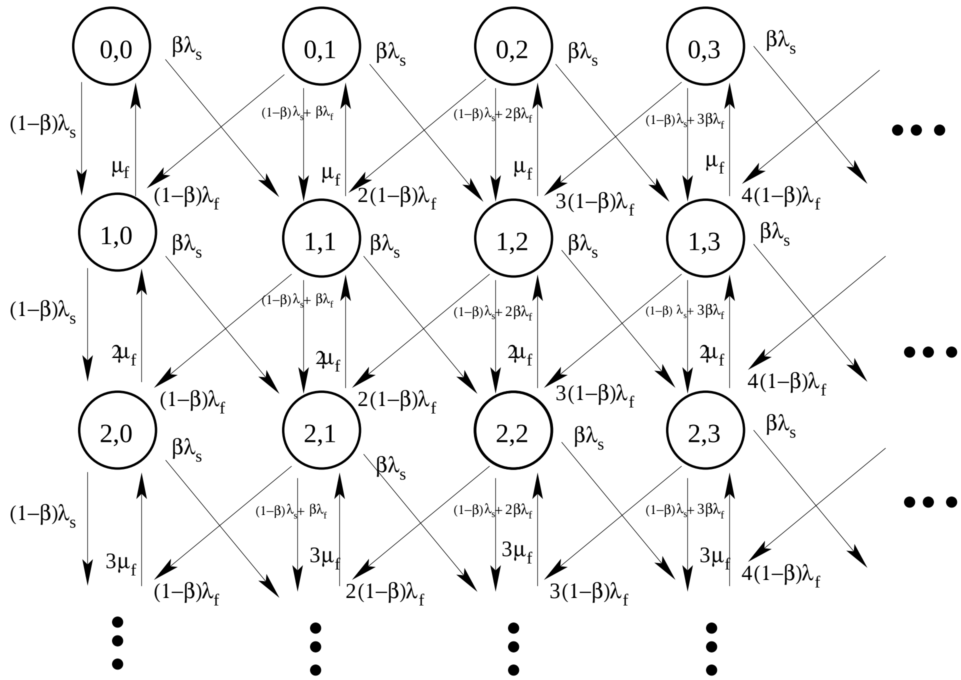

Fig. 2. CTMC of the model

IV. Traffic Measurements and Analysis

We collected and processed traces using `tcpdump` [18] and `Tstat` [19], [20] observing the data flow on the Internet access link of Politecnico di Torino, i.e., between the edge router of Politecnico and the access router of GARR/B-TEN [21], the Italian and European Research network.

`Tstat` is a new tool developed at Politecnico di Torino, which is able to analyze previously collected traces, and produce a number of results at both the IP and TCP level. Indeed, `Tstat` is able to rebuild each TCP connection status by looking at the TCP header in the forward and backward packet flows. `Tstat` requires as input a trace collected on an edge node, such that both data segments and ACK segments can be analyzed. When `Tstat` observes a TCP connection opening and closing, it marks the flow as *complete*, and proceeds analyzing it. Additional information about `Tstat` and statistical analysis performed on collected traces can be found in [19] and [20].

Data were collected on files storing 6 or 3 hours long traces (to avoid exceeding the File System limitation on file size), for a total of more than 100 Gbytes of compressed data. Within the Politecnico Campus LAN, there are approximately 7,000 hosts; most of them are clients, but several servers are regularly accessed from outside institutions.

The data was collected during different time periods, in which the network topology evolved. We selected the following periods:

- Period (A) - June 2000: from 6/1/2000 to 6/11/2000, when the bandwidth of the access link was 4 Mbit/s, and the link between the GARR network and the corresponding US peering was 45 Mbit/s;
- Period (B) - February 2001: from 2/1/2001 to 2/19/2001, when the bandwidth of the access link was 16 Mbit/s, and the link between the GARR network and the corresponding US peering was 622 Mbit/s.

TABLE I

SUMMARY OF THE ANALYZED TRACES

Name	date	start time	stop time	packets $\cdot 10^6$	flows $\cdot 10^3$
Peak'01	2 Feb 01	10:52	13:52	11	540
Night'01	2 Feb 01	04:52	07:52	0.43	30
Peak'00	13 Apr 00	08:10	14:10	12	564
Night'00	13 Apr 00	02:10	08:10	0.92	79

The two periods are characterized by a significant upgrade in network capacity. In particular, the campus access link was a bottleneck during June 2000, while it was not during January 2001. The same consideration applies to the GARR-US peering capacity, which plays a key role, since most of the traffic comes from US research sites.

Among all the traces we collected, we report here results from four traces, which we consider representative of different network scenarios. Table I summarizes the key parameters of the selected traces, reporting in the last two columns the number of samples in each trace. Since our campus network can be mainly considered as a "client" network, i.e., hosts in the network are mainly destinations of information, in the remaining of this paper we will present results considering incoming streams of data, both at the flow level, and at the packet level.

The definition of LRD we use is the divergence of the spectrum at the origin of a stationary stochastic process with finite second moments:

$$f(v) \sim c|v|^{(-\alpha)}, v \to 0 \tag{1}$$

In this definition there are two parameters: (α, c). $\alpha \in [0, 1)$ is the dimensionless scaling exponent, and is the most important parameter describing the qualitative nature of the scaling. For a non-LRD stationary process (hence a process which is not Self-Similar), $\alpha = 0$ at large scales. $c \in \Re^+$ has the same dimension of the variance of the process and describes the quantitative aspects of LRD often referred to as the LRD 'size.' LRD implies that the sum of correlations over all lags is infinite; however, individually their sizes at large lags are proportional to c, and can be arbitrarily small.

In order to estimate the LRD properties of processes we use the wavelet-based approach developed in [22], [23]. We also adopt the tools presented in [22], [23] and usually referred to as the *AV* estimator tool. Other approaches can be pursued to analyze traffic traces, but the wavelet framework has emerged as one of the best estimators, as it offers a very versatile environment, as well as fast and efficient algorithms.

Since traffic traces are finite, and their asymptotic behavior cannot be derived, we always limit the analysis between two scales $(j_{\inf}, j_{\sup})$. In order to evaluate LRD parameters we use the LogScale Diagram, which is essentially a log-log plot of the mean square values estimates of the wavelet detail coefficients x_n^j versus the scale j. j is generally called *octave*, since 2^{-j} has the dimension of a frequency. Through the LogScale Diagram, it is possible to identify the presence of LRD and determine the cutoff scales $(j_{\inf}, j_{\sup})$ at which LRD 'begins' and 'ends.' Within these scales, an LRD process LogScale Diagram is linear with coefficient α. Indeed, for all processes $j_{\sup}$ is limited by the trace length, and $j_{\inf}$ correspond to a scale of a few hundreds ms.

The analysis we report is based on the estimation of the Hurst parameter H. The definition of LRD given in (1) refers to the parameter α. When considering the process of the increments of a Self-Similar process with finite increments,

TABLE II

FLOW LEVEL ANALYSIS OF '01 TRACES

Trace	Peak'01			Night'01		
	$\hat{H}_f$	$\hat{c_f}$	$1/\hat{\Lambda}_f$	$\hat{H}_f$	$\hat{c_f}$	$1/\hat{\Lambda}_f$
I $[ms]$	0.74	82.4	20.01	0.86	8491	358.9
N_{1s}	0.76	59.4	49.9	0.76	1.66	2.79
N_{100ms}	0.75	2.01	4.99	0.73	0.07	0.28
N_{10ms}	0.74	0.07	0.49	0.80	0.001	0.028

TABLE III

FLOW LEVEL ANALYSIS OF '00 TRACES

Trace	Peak'00			Night'00		
	$\hat{H}_f$	$\hat{c_f}$	$1/\hat{\Lambda}_f$	$\hat{H}_f$	$\hat{c_f}$	$1/\hat{\Lambda}_f$
I $[ms]$	0.76	275.1	39.6	0.74	2414	271.6
N_{1s}	0.75	28.4	25.9	0.78	1.54	3.68
N_{100ms}	0.74	0.79	2.53	0.76	0.06	0.37
N_{10ms}	0.75	0.015	0.25	0.78	0.001	0.04

which are the processes considered in early works, the relationship $H = (\alpha + 1)/2$ holds for $H \in [0.5, 1]$; hence it is common practice (though not completely proper) to use the parameter H also discussing LRD, and we stick to this practice; clearly a non-LRD process has $H = 0.5$, while Hurst parameters larger than 0.6–0.8 are normally assumed as an indication of strong LRD.

Among all the possible metrics of traces, we selected as most representative of the traffic characteristics the inter-arrival time process $I(k)$, and the counting processes $N_T(n)$, obtained by counting the number of arrivals in a time interval $[nT, (n+1)T)$. We use three values of T: $1, 0.1, 0.01$ s.

Combining the three tools (`Tstat`, `tcpdump`, and *AV*) we analyzed the metrics defined above at both the flow and packet levels. As explained in Sect. III, the statistical analysis at the session level is not possible, due to the difficulty in defining sessions. When measured data are fed into the model, two free parameters are left and can be used to match the LRD characteristics of traces, as explained in Sect. IV-B.

A. Flow and Packet Level Analysis

The *AV* analysis of the traces produces estimates of three main parameters: the Hurst parameter $\hat{H}$, the $\hat{c}$ factor, and the mean value $1/\hat{\Lambda}$. Each parameter is available for both flow level traces, that we identify with subscript f, and packet traces, that we identify with subscript p. In addition, the traffic measurements yield the following parameters.

$\hat{N}_p$: the average number of packets per flow;

$\hat{\Lambda}_f$: the flow generation rate;

$\hat{\lambda}_p$: the packet generation rate of active flows (obtained as the ratio between the average flow duration and $\hat{N}_p$).

Tables II and III report the results of the flow level analysis for the '01 traces and the '00 traces respectively, while Tables IV and V refer to the packet level analysis for the same traces. Apart from the obvious consideration that during peak hours arrival rates are much higher than during nights, a few considerations are in order.

TABLE IV

PACKET LEVEL ANALYSIS OF '01 TRACES

Trace	Peak'01			Night'01		
	$\hat{H}_p$	$\hat{c_p}$	$1/\hat{\Lambda}_p$	$\hat{H}_p$	$\hat{c_p}$	$1/\hat{\Lambda}_p$
$I\ [ms]$	0.87	0.01	0.63	0.71	$5 \cdot 10^{-4}$	0.02
N_{1s}	0.88	5232	1577	0.73	457.8	54.5
N_{100ms}	0.88	91.4	157.7	0.72	16.6	5.45
N_{10ms}	0.88	1.50	15.8	0.76	0.30	0.54

TABLE V

PACKET LEVEL ANALYSIS OF '00 TRACES

Trace	Peak'00			Night'00		
	$\hat{H}_p$	$\hat{c_p}$	$1/\hat{\Lambda}_p$	$\hat{H}_p$	$\hat{c_p}$	$1/\hat{\Lambda}_p$
$I[ms]$	0.84	0.17	2.25	0.84	40.54	15.74
N_{1s}	0.86	504.19	444.75	0.83	133.93	63.51
N_{100ms}	0.87	14.50	44.49	0.83	3.61	6.35
N_{10ms}	0.88	0.23	4.45	0.87	0.04	0.63

First of all, notice that all processes of all traces show a similar Hurst parameter ranging from 0.71 to 0.88. Indeed, $\hat{H}_f$ is almost independent from the trace or process considered, but it was slightly higher in '01. $\hat{H}_p$, instead is higher during peak hours (around 0.88 for all processes and traces) than during nights, and for the night trace in '00 it was higher for all processes. Globally, these results hint to the fact that LRD in packet networks is probably not due to high load. The network characteristics of the four selected traces are so different (different link speeds, different loads, different patterns between Peak and Night), that this can be taken as a strong indication that LRD is an intrinsic characteristic of the Internet traffic and is not induced by network conditions.

A second consideration concerns c, whose value is extremely variable and clearly connected to the absolute magnitude of the analyzed process (not surprisingly, since it is connected to the mean square value of the process itself).

The last consideration is that the characteristics of the measured traces do not change significantly from '00 to '01, hence in the sequel we will only present results for these latter traces, that are more recent. The same modeling procedure applied to the '00 traces yields similar results.

B. Setting the Model Parameters

The model described in Sect. III is completely described by five parameters:

λ_s : the arrival rate of new sessions;

λ_f : the flow arrival rate per active session;

λ_p : the packet arrival rate per active flow;

N_f : the average number of flows per session;

N_p : the average number of generated packets per active flow.

Three of these parameters can be directly set by measurements on traces, while we use the remaining two, namely those that are harder to measure from traces because of the fuzzy definition of sessions, to match the Hurst parameters

$\hat{H}_f$ and $\hat{H}_p$. Therefore, the model has two "turning-knobs". Sect. V-B will discuss the sensitivity of the model to these parameters.

The matching of model parameters given a set of measured data can be done using the following procedure.

---

1. From the traffic traces estimate $\hat{H}_p$, $\hat{H}_f$, $\hat{\Lambda}_f$, $\hat{\lambda}_p$, $\hat{N}_p$
2. Set $N_p = \hat{N}_p$ and $\lambda_p = \hat{\lambda}_p$
3. Let $C = \lambda_s/\lambda_f$ and set the initial values $N_f = 1$ and $C = 1$
4. Compute $\lambda_s = \frac{\hat{\Lambda}_f}{N_f}$
5. Generate a synthetic sequence with the same number of samples as the real trace
6. Estimate the Hurst parameter at both IP and TCP level of the synthetic trace and compare them with $\hat{H}_p$, $\hat{H}_f$
7. If the fitting is good, the procedure ends else assign new values to N_f and C and go to 4

The new values to be assigned to N_f and C at step 7 of the procedure are chosen following the empirical evidence that the larger N_f is, the larger $\hat{H}_p$ and $\hat{H}_f$ are; and also, the larger C is, the larger $\hat{H}_f$ while C has little influence at packet level (see Sec. V-B for more details).

The selection of λ_f and N_f by means of the fitting procedure at steps 4–7 of the procedure, typically requires only few iterations to provide accurate results (approximately 10 in our tests).

Example:

```
C=1;
Nf=1;
// TCP flow arrivals are Poisson and Hf=0.5
generate a sinthetic sequence and estimate H_f,H_p;
fit = ( |Hf-Hf_est| < eps_f) && ( |Hp-Hp_stim| < eps_p);
// eps_f = eps_p = 0.05
while !(fit)
{
       if ( Hf<Hf_stim )
            Nf = Nf + 5;
       else if ( Hf>Hf_stim )
            Nf = Nf - 1;
       else if ( Hp>Hp_stim )
            C  = 2*C;
       else if ( Hp<Hp_stim )
            C  = C/2;
       generate a sinthetic sequence and estimate H_f,H_p;
       fit = ( |Hf-Hf_est| < eps_f) && ( |Hp-Hp_stim| < eps_p);
}
```

TABLE VI

MEASURED VALUES FROM REAL TRACES USED IN SETTING THE MODEL PARAMETERS

	Peak '01	Night '01
$\hat{N}_p$	22.03	14.35
$\hat{\lambda}_p$	29.57	352.9
$\hat{\Lambda}_f$	49.95	2.78
$\hat{\Lambda}_p$	1267.1	67.34

TABLE VII

MODEL: FLOW LEVEL RESULTS FITTING THE '01 TRACES

Trace	Peak'01			Night'01		
	$\hat{H}_f$	$\hat{c_f}$	$1/\hat{\Lambda}_f$	$\hat{H}_f$	$\hat{c_f}$	$1/\hat{\Lambda}_f$
$I\ [ms]$	0.74	19.2	20.1	0.84	7799	356.8
N_{1s}	0.71	76.7	49.8	0.82	0.51	2.81
N_{100ms}	0.71	2.27	4.98	0.83	$7.6 \cdot 10^{-3}$	0.28
N_{10ms}	0.78	0.013	0.50	0.79	$1 \cdot 10^{-4}$	0.03

Table VI reports the values of the measured parameters used to set the model parameters for the '01 traces. Tables VII and VIII report the synthetic traces characteristics measured by the *AV* tool when the model parameters are fitted to the '01 traces; the values must be compared with those of Tables II and IV. The Hurst parameters are very well matched, while the values for the c parameter are less precise, though the qualitative behavior is the same.

V. PERFORMANCE EVALUATION

The comparison of Tables VII and VIII with Tables II and IV indicates that the simple MMPP model we are considering is indeed capable of capturing some of the LRD characteristics of the traffic we measured. In some sense this conclusion is straightforward, since we tuned the model to obtain them.

The performance metrics we are most interested into, however, are related to the behavior of the synthetic traces when they feed a buffer in front of a transmission link, as well as the predictability and tunability of the model to represent different traffic scenarios. These two aspects are discussed in the following two Sections.

A. Queueing Analysis

We consider a $G/GI/1/B$ queue, where the input process is either the real trace or the synthetic trace generated by the MMPP model. The service time distribution reflects the packet length distribution measured from the real traces, as reported in Fig. 3. The distribution reflects the well known multi-mode distribution, with peaks for very short packets and for the different MTUs (Maximum Transfer Units) in the network. The peak of the Ethernet frame at 1,500 bytes dominates the lot. In order to evaluate the performance of the queue under different values of the

TABLE VIII

MODEL: PACKET LEVEL RESULTS FITTING THE '01 TRACES

Trace	Peak'01			Night'01		
	$\hat{H}_p$	$\hat{c_p}$	$1/\hat{\Lambda}_p$	$\hat{H}_p$	$\hat{c_p}$	$1/\hat{\Lambda}_p$
$I\ [ms]$	0.84	0.038	0.90	0.82	$1 \cdot 10^{-4}$	0.025
N_{1s}	0.87	5943	1113	0.87	35.2	40.06
N_{100ms}	0.82	178.2	111.4	0.79	2.05	4.01
N_{10ms}	0.84	3.13	11.14	0.86	$5 \cdot 10^{-3}$	0.41

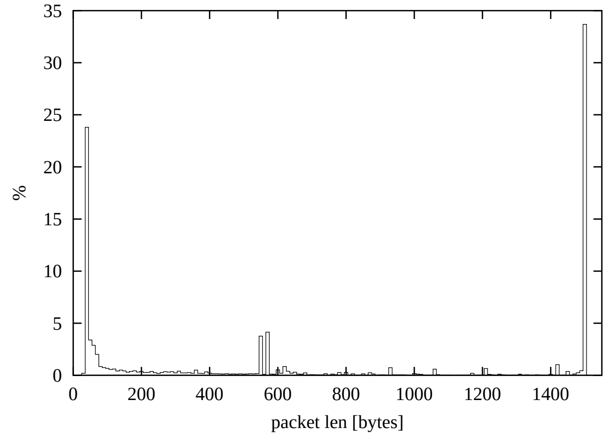

Fig. 3. Packet length distribution used in simulations

load, we change the average service time while keeping the packet length distribution unchanged. Notice that the load of the queue has no relation with the actual load of the link where the traces were collected.

Fig. 4 reports the queue length distribution for the Peak'01 trace for an infinite buffer. The dashed line is obtained using as input the measured trace, the solid line is obtained with the MMPP model and the dot-dashed line with a simple Poisson process whose arrival rate matches the average packet arrival rate measured on the trace. Plots refer to four different loads: 0.9, 0.8, 0.7, and 0.6. For the real traces, the tail of the distributions below 10^{-4} become noisy due to lack of samples. The buffering behavior of the model matches quite well that of the measured traces, while the Poisson model underestimates the buffer level of orders of magnitude (see for instance the probability of having more than 100 packets in the buffer). Notice that the accuracy of the MMPP model predictions tends to increase for large values of the load, which correspond to the most interesting cases. On the contrary, as the load reduces, the tail of the distribution is matched less accurately.

Similar results are shown in Fig. 5 which refers to the Night'01 trace. In this case, the situation is somewhat more complex and less satisfactory, specially when the load is light. It must be noted, however, that this scenario is both less interesting and somewhat more 'artificial' than the previous one. First of all, the number of points in the measured trace is about 50 times smaller than during peak hours (see Table I) and this explains why trace-driven curves are noisier. Second, the real link load at night is extremely low, thus artificially forcing the load to values higher than 0.6 significantly modifies the overall scenario. Yet, despite of this, the MMPP model matches quite well the tail of the distribution, which is typically the most interesting part of the curve.

Fig. 6 reports the queue length distribution for a finite buffer queue driven by the same arrival process as in the previous scenarios for the Peak'01 trace, with ρ=0.9. The considered buffer sizes are B = 32, 64, 128, 256, 512 packets. Notice again the accuracy of the model in evaluated the queue performance.

B. Sensitivity Analysis

We now discuss the impact of the MMPP model parameters on both the Hurst parameters and the queueing behavior.

Fig. 7 reports a sensitivity analysis of the Hurst parameters of the synthetic traces on both N_f and C (we remind that C is the ratio $\frac{\lambda_s}{\lambda_f}$). First of all, notice that the Hurst parameter is always larger than 0.5 at both the flow and packet levels, and it increases as N_f increases, consistently with the intuition that a larger N_f introduces more memory in the system. Moreover, notice that at the packet level, there is a higher degree of memory and correlation since packets are generated by flows which are generated by sessions.

Let us now focus on C, which represents the burstiness of the generation of flows in a session: The larger C is, the more bursty the generation of flows per session is. The influence of C on the Hurst parameter is quite complex. At the flow level, a higher degree of burstiness tends to induce a larger value of the Hurst parameter, while the opposite is true at the packet level.

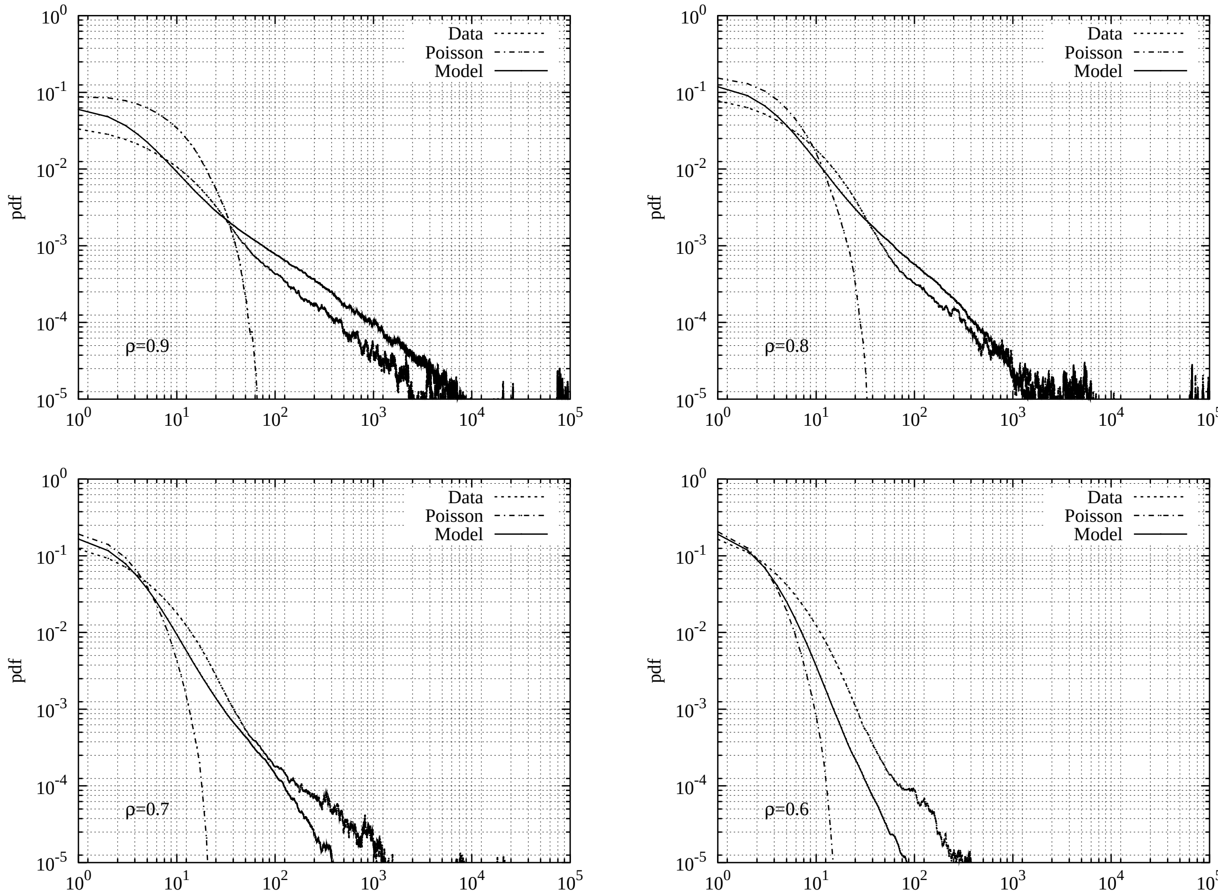

Fig. 4. Buffer occupancy distribution for the model, the Peak'01 trace and Poisson arrivals; load ρ = 0.9, 0.8, 0.7 and 0.6

Finally, Fig. 8 reports the sensitivity of the buffer occupancy to N_f. In the figure, all the parameters except N_f, are set according to the model for the Peak'01 trace. Again, for comparison purposes only, we report also the curve derived with Poisson packet arrivals. The case $N_f = 1$ corresponds to one flow only per session. Increasing the number of flows per session makes the queue tail heavier. Observe that the queue occupancy pdf is the superposition of weighted exponential pdfs, as it must be for an MMPP model. Observe also that, as N_f increases, the range where the queue decay follows roughly a power law becomes longer. Clearly, there is always a value beyond which the queue decay is exponential and this 'cutoff value' is proportional to N_f. Indeed, the behavior of our model confirms the results obtained in [24], where the Hurst parameter of measured traffic, as well as the queueing behavior of the same traffic, are fitted through the use of phase distributions.

VI. Conclusions

In this paper we have analyzed traffic traces at the edge router of our institution and derived a simple MMPP traffic model that matches the traffic characteristics quite well. The model is based on a layered structure of sessions, that generate flows, that finally generate packets.

The characteristics of the synthetic traffic generated with the model match the LRD characteristics observed in the measured traces. One of the interesting features of the MMPP model is that it requires as inputs five parameters only. Three of these parameters can be directly mapped onto easily measurable traffic parameters, such as the average flow arrival rate, the number of packets per flow, and the arrival rate of packets within flows. The other two define the notion of session, and are more difficult to map onto measured parameters, so that a "best-fit" approach is needed to match the traces. However, these two parameters allow the direct and predictable control the Hurst parameter of the synthetic traffic on the considered scaling range.

Most interesting is that the behavior of the synthetic traffic feeding a $G/GI/1/B$ queue matches very well the behavior of the original traces. Thus, the proposed model can be effectively used for dimensioning buffers and link capacities. Indeed, the model provides accurate and reliable

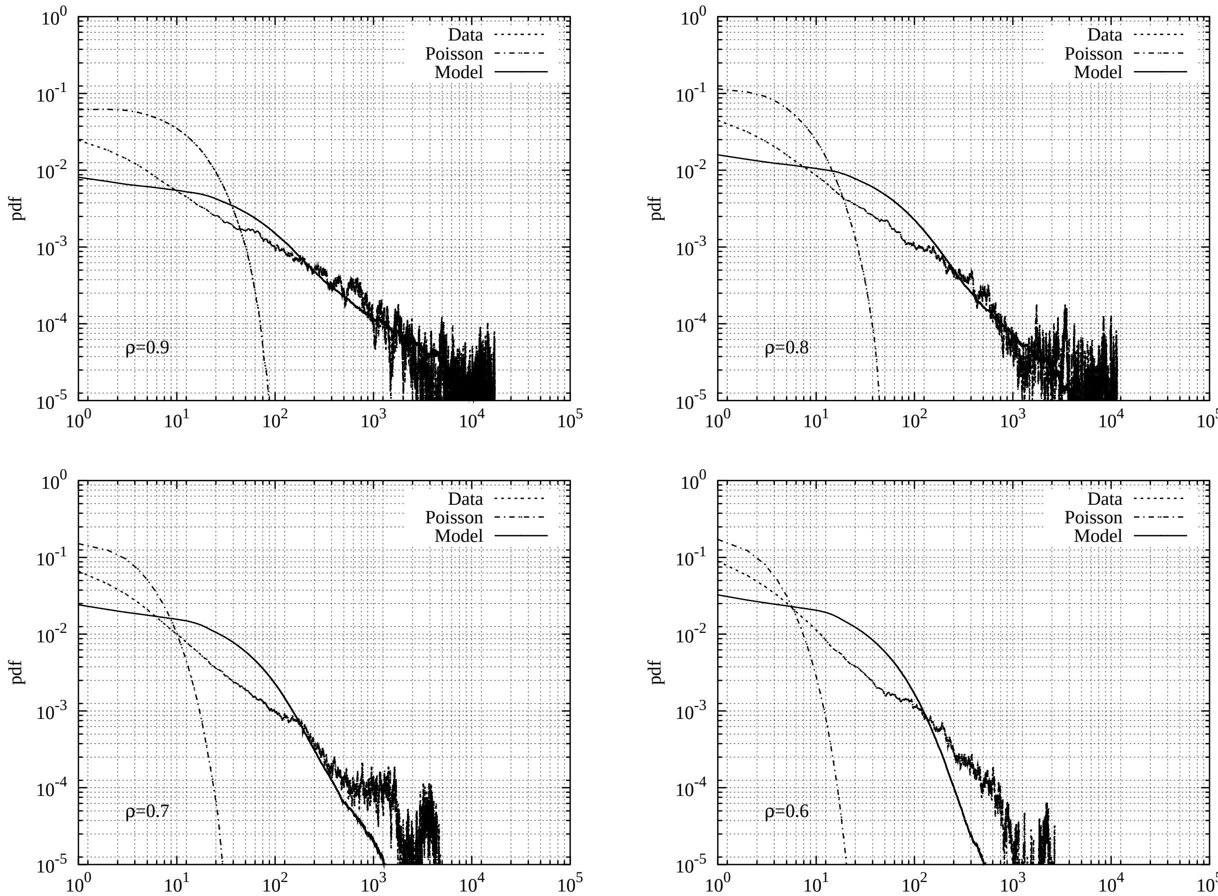

Fig. 5. Buffer occupancy distribution for the model, the Night'01 trace and Poisson arrivals; load ρ = 0.9, 0.8, 0.7 and 0.6

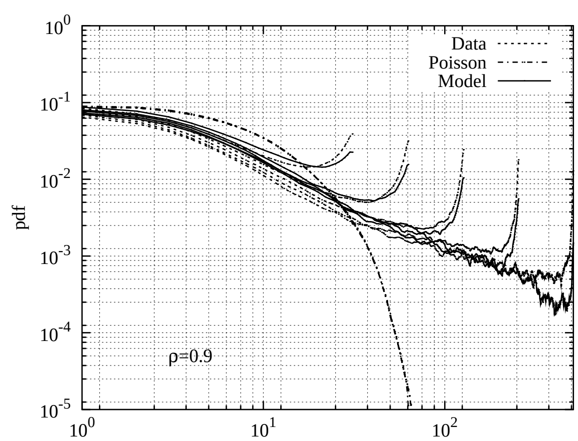

Fig. 6. Buffer occupancy distribution for the model, the Peak'01 trace and Poisson arrivals; finite buffer with capacity equal to 32, 64, 128, 256, 512 packets and load ρ = 0.9

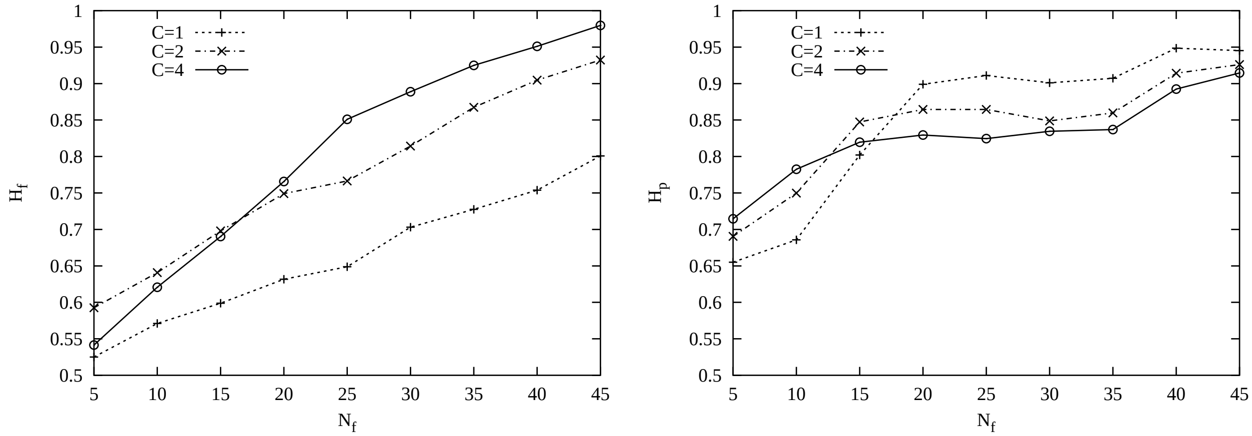

Fig. 7. Impact of N_f on the Hurst parameters of the flow arrival process, for different values of the ratio between λ_s and λ_f

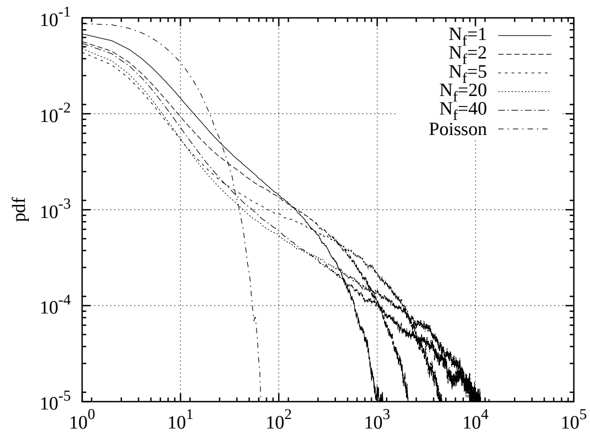

Fig. 8. Buffer occupancy distribution for the model with the parameters used for the Peak'01 trace and different values of N_f

results while being extremely simple. The key features of the proposed MMPP model are its simplicity and its intuitive structure. While, on the one hand, these features present an extremely accurate match of the characteristics of the measured traces, on the other hand, they allow the model to be used by traffic engineers with only a limited knowledge of the sophisticated theoretical aspects of the LRD processes.

REFERENCES

[1] W.E. Leland, M.S. Taqqu, W. Willinger, V. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended version)," *IEEE/ACM Transaction on Networking*, Vol.2, N. 1, pp. 1–15, Jan. 1994.

[2] V. Paxson, S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling," *IEEE/ACM Transactions on Networking*, Vol.3, N.3, pp. 226–244, Jun. 1995.

[3] M.E. Crovella, A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," *IEEE/ACM Transaction on Networking*, Vol.5, N.6, pp. 835–846, Dec. 1997.

[4] W. Willinger, M.S. Taqqu, R. Sherman, D.V. Wilson, "Self-Similarity Through High Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level," *IEEE/ACM Transaction on Networking*, Vol.5, N.1, pp. 71–86, Jan. 1997.

[5] M. E. Crovella, M. S. Taqqu, A. Bestavros, "Heavy-Tailed Probability Distributions in the World Wide Web," In *A Practical Guide To Heavy Tails*, Chapter 1, Chapman & Hall, New York, pp. 3–26, 1998.

[6] A. Veres, M. Boda, "The Chaotic Nature of TCP Congestion Control," *IEEE INFOCOM 2000*, Tel Aviv, Israel, pp. 1715-1723, Mar. 2000.

[7] A. Veres, Z. Kenesi, S. Molnar, G. Vattay, "On the propagation of long-range dependence in the Internet," *ACM Sigcomm 2000*, Stockholm, Sweden, pp. 243-254, Sept. 2000.

[8] H. Riedi, M.S. Crouse, V.J. Ribeiro, R.G. Baraniuk, “A multifractal Wavelet Model with application to data traffic,” *IEEE Transaction on Information Theory*, Vol. 45, No. 4, pp. 992-1018, Apr. 1999.

[9] A. Feldmann, A.C. Gilbert, W. Willinger, “Data Networks as Cascades: Investigating the Multifractsl Nature of Internet WAN Traffic,” *ACM Sigcomm*, Boston, Ma, pp. 42-55, Sept. 1998.

[10] S. Ma, C. Ji, “Modeling Heterogeneous Network Traffic in Wavelet Domain”, *IEEE/ACM Transaction on Networking*, Vol.9, No.5, pp. 634-649, Oct. 2001.

[11] A. Eramilli, R.P. Singh, “An Application of Deterministic Chaotic Maps to Model Packet Traffic,” *Queueing Systems*, Vol.20, pp. 171–206, 1995.

[12] R.J. Mondragon, J.M. Pitts, D.K. Arrowsmith, “Chaotic Intermittency-Sawtooth Map Model of Aggregated Self-Similar Traffic Streams,” *Electronic Letters*, Vol.36, N.2, pp. 184–186, 2000.

[13] M. Krunz, A. Makowski, “A Source Model for VBR Video Traffic Based on M/G/infinity Input Processes,” IEEE Infocom 98, San Francisco, CA, USA, pp. 1441-1448, Apr. 1998.

[14] M. Grosslauser, J. Bolot, “On the Relevance of Long-Range Dependencies in Network Traffic,” *IEEE/ACM Transaction on Networking*, Vol.7, N.5, pp. 629–640, Oct. 1999.

[15] S. Ben Fredj, T. Bonald, A. Proutiere, G. Régnié, J. Roberts, “Statistical Bandwidth Sharing: a Study of Congestion at Flow Level,” *Proc of Sigcomm 2001*, San Diego, CA, pp. 111-122, Aug. 2001.

[16] T. Bonald, A. Proutiere, G. Régnié, J. Roberts, “Insensitivity Results in Statistical Bandwidth Sharing,” *International Teletraffic Conference (ITC) 2001*, San Paolo, Brasil, Nov. 2001.

[17] A. Feldmann, “Characteristics of TCP Connection Arrivals,” In *Park and Willinger (editors): Self-Similar Network Traffic and Performance Evaluation*, Wiley-Interscience, 2000.

[18] S. McCanne, C. Leres, V. Jacobson, “Tcpdump,” http://www.tcpdump.org, 2002.

[19] M. Mellia, A. Carpani, R. Lo Cigno, “Tstat web page,” http://tstat.tlc.polito.it/, 2002.

[20] M. Mellia, A. Carpani, R. Lo Cigno, “Measuring TCP and IP Behaviour on a Edge Node,” *IEEE Globecom 2002*, Taipei, Nov. 2002.

[21] “GARR - Rete dell’Università e della Ricerca Scientifica Italiana,” http://www.garr.it, 2002.

[22] P. Abry, P. Flandrin, M.S. Taqqu, D. Veitch, “Self-Similarity and Long-Range Dependence Through the Wavelet Lens,” In *Doukhan, Oppenheim, Taqqu (editors): Long range dependence: theory and applications*, 2000.

[23] P. Abry, D. Veitch, “Wavelet Analysis of Long Range Dependent Traffic,” it IEEE Transactions on Information Theory, Vol.44, No.1 pp. 2-15, Jan. 1998.

[24] A. Horváth, G.I. Rózsa, M. Telek, “A map fitting method to approximate real traffic behaviour,” *8-th IFIP Workshop on Performance Modelling and Evaluation of ATM and IP Networks*, Ilkley, U.K., July 2000.

[25] A. T. Andersen, B. F. Nielsen, “A markovian approach for modeling packet traffic with long-range dependence,” *IEEE journal on selected areas in communications*, Vol.16, No.5, pp. 719-732, 1998.