

Modeling Multifractal IP Traffic: Characterization of Packet Arrivals and Packet Sizes using Stochastic L-Systems

Paulo Salvador, António Nogueira, Rui Valadas
Institute of Telecommunications, University of Aveiro
Campus Universitário, 3810-193 Aveiro, Portugal
phone: (+351) 234377900, fax: (+351) 234377901
emails: salvador@av.it.pt, nogueira@av.it.pt, rv@det.ua.pt

Abstract

We consider a novel traffic model, and an associated parameter fitting procedure, that jointly describes the packet arrival and the packet size processes, allowing for the incorporation of multifractal scaling behavior in both characteristics. The model is based on so-called stochastic L-Systems, which were introduced by biologist A. Lindenmayer as a method to model plant growth. We provide a detailed comparison with a related multifractal model based on conservative cascades. Our results, that include applying the fitting procedure to real observed data with multifractal scaling behavior on both the packet arrival and packet size processes, show that L-System based models can achieve excellent fitting performance in terms of first and second order statistics and queuing behavior.

Keywords: Traffic modeling, Multifractal, L-System.

1 Introduction

Accurate modeling of IP traffic requires the joint characterization of both the packet arrival and the packet size processes. In particular, this is required for providing a close fit of the queuing behavior (i.e., the packet loss ratio or average packet delay suffered in a network node). The queuing behavior addresses the effect of traffic on network performance, and is one of the most important

criteria to assess the suitability of traffic models (and associated parameter fitting procedures). Here, the analysis consists in comparing the curves of packet loss ratio (or average packet delay) versus buffer size, obtained with the measured traces (through trace-driven simulation) and with the inferred traffic model (using again trace-driven simulation or numerical computation of the performance measures whenever possible). When dealing with models that characterize only the arrival process, it is common practice to assume that the packet size is fixed and equal to the average packet size of the measured trace. This may lead to large errors when the packets have variable size, such as in IP traffic, and more so, when the packet size process exhibits long-range dependence and/or multifractality.

Recent analysis of measured Internet WAN traffic has revealed that multifractal structures, such as random cascades, can help explaining the scaling behavior typically associated to networking mechanisms operating on small time scales (e.g. TCP flow control). A cascade (or multiplicative process) is a process that fragments a set into smaller and smaller components according to a fixed rule, and at the same time fragments the measure of the components by another (possibly random) rule. Random cascades were introduced by Mandelbrot as a physical model for turbulence [1]. In the traffic modeling context, the set can be interpreted as a time interval and the measure as the number of arrivals or number of bytes (in that interval).

The multifractal nature of network traffic was first noticed by Riedi and Lévy Véhel [2]. Subsequently various studies have addressed the characterization and modeling of multifractal traffic, essentially within the framework of random cascades [3] [4] [5] [6] [7] [8] [9] [10].

In [11] we have introduced the use of stochastic Lindenmayer-Systems (hereafter referred to as L-Systems), as a model for multifractal network traffic. In this paper, we propose a new modeling approach (including a parameter fitting procedure) based on stochastic L-Systems, that characterizes both the packet arrival and packet size processes. L-Systems are string rewriting techniques which were introduced by biologist A. Lindenmayer in 1968 as a method to model plant growth [12]. They are characterized by an alphabet, an axiom and a set of production rules. The alphabet is a set of symbols; the production rules define transformations of symbols into strings of symbols; starting from an initial string (the axiom), an L-System constructs iteratively sequences of symbols replacing each symbol by the corresponding string according to the production rules. Stochastic L-Systems are a method to construct recursively random sequences with multifractal behavior [13]. The proposed traffic model uses two stochastic L-Systems, one for the packet arrivals and the other for packet sizes. Here the

packet arrival process is described by the sequence of arrival rates (number of arrivals per time interval, in pkts/sec) and the packet size process by the sequence of mean sizes (where the mean is computed over the same time interval, in bytes). The traffic model also includes the ability of modeling multiple scaling behavior. A multifractal traffic model also characterizing the packet arrival and packet size processes was introduced in [10]; this is based on conservative cascades, which are a special case of random cascades. When compared with random cascades, stochastic L-Systems introduce a dependence on the construction process (due to the production rules), that has a meaningful physical explanation, and can help understanding the joint impact of network mechanisms and resource limitations on observed traffic.

This paper is organized as follows. In section 2 we give some background on L-Systems. In sections 3 and 4 we present the traffic model and describe the associated fitting procedure. Section 5 provides a detailed comparison with random cascades and includes a physical explanation supporting L-System based traffic models. In section 6 we discuss the results of applying the proposed fitting procedure to measured traffic traces. Finally, section 7 presents the conclusions.

2 L-Systems Background

The basic idea behind L-Systems is to define complex objects by successively replacing parts of a simple object using a set rules. The L-System is a feedback machine that operates on strings of symbols. The set of symbols is called the alphabet. Starting from an initial state (called axiom), an L-System operates, at each iteration, by applying the set of production (or rewriting) rules simultaneously to all symbols of an input string to give an output string. For a comprehensive introduction to L-Systems see [13].

Consider a simple example of an organism growing through cell subdivisions. There are two types of cells represented by letters A and B. Cell subdivisions are modeled by replacing these symbols with strings of symbols: cell A subdivides into two cells represented by string AB; cell B subdivides into two cells represented by string AA. The ordering of the symbols is relevant in an L-System. The organism modeled by this L-System grows by repeated cell subdivisions. At birth the organism is the single cell A. After one subdivision the organism is two cells represented by string AB. After two subdivisions, the organism has four cells given by string ABAA, and after three subdivisions the organism has eight cells represented by string ABAAABAB. Using the formalism of L-Systems this

growth process can be described as:

Alphabet:	{A,B}
Axiom:	A
Rules:	A \rightarrow AB
	B \rightarrow AA

which results in:

Start	A
Iteration 1	AB
Iteration 2	ABAA
Iteration 3	ABAAABAB
Iteration 4	ABAAABABABAAABAA
...	...

The production rules can be stochastic. In stochastic L-Systems there may be several production rules for one symbol, and the specific rule is selected according to a probability distribution. Taking previous example, one production rule could be to convert A into AB with probability 0.4 or into BB with probability 0.6 (instead of converting always A into AB). In this case, after 3 iterations several strings are possible, e.g., ABAAABAB, BBAAABAB, or BBBB BBBB. Stochastic L-Systems are a method to construct recursively random sequences with multifractal behavior [13].

3 Traffic Model

In this section we explain the proposed multifractal traffic model, which characterizes both the packet arrivals and the packet sizes. In particular, the traffic model will be a time series describing the arrival rate and the mean packet size over (small) non-overlapping time intervals. The time series construction is based on two independent L-Systems, one for packet arrivals and the other for packet sizes. For the packet arrivals, we work on an alphabet of arrival rates defined by

$$\vec{\lambda} = \{\lambda_1, \lambda_2, \dots, \lambda_L\}, \lambda_i \in \mathbb{R}_0^+, i = 1, \dots, L. \quad (1)$$

and with production rules that randomly generates two arrival rates from a previous one. For the packet sizes, we work on an alphabet of mean packet sizes defined by

$$\vec{\gamma} = \{\gamma_1, \gamma_2, \dots, \gamma_G\}, \gamma_i \in \mathbb{R}_0^+, i = 1, \dots, G. \quad (2)$$

and with production rules that randomly generates two mean packet sizes from a previous one. Without loss of generality, we assume $\lambda_1 < \lambda_2 < \dots < \lambda_L$ and $\gamma_1 < \gamma_2 < \dots < \gamma_G$.

The time series of arrival rate and mean packet size pairs, is constructed progressively, governed by two L-System machines, where each iteration produces a new time scale. Starting with the coarsest time scale, where traffic is characterized by a single arrival rate and a single mean packet size over a single time interval, each iteration generates a finer time scale by (i) division of each (parent) time interval in two new equal length (child) subintervals and (ii) association of arrival rates and mean packet sizes to each new subinterval, according to the production rules of each stochastic L-System. For both L-Systems we allow the grouping of time scales in time scale ranges (not necessarily the same for the two L-Systems) and the definition of different sets of production rules for each time scale range. This is motivated by the fact that each set of production rules maps into a distinct scaling behavior [14]. Classically, there will be scaling if the log-log plot of the q^{th} order energies (usual energy is $q = 2$) as a function of scale behaves linearly; if the plot is (globally) non-linear, different time scale ranges can be detected where linearity is observed (see [6], for example). The traffic process construction is illustrated in Figure 1.

To characterize the traffic process we define the pair $(X_{(j,r_a)}^{(i)}, Y_{(j,r_b)}^{(i)})$, where $X_{(j,r_a)}^{(i)} \in \vec{\lambda}$ and $Y_{(j,r_b)}^{(i)} \in \vec{\gamma}$ are the arrival rate and the mean packet size at time interval i of time scale j and time scale ranges r_a for arrivals and r_b for sizes, respectively. Let the number of scales be S and the number of ranges of scales be R_a for arrivals and R_b for sizes. For convenience, we let j decrease from $j = S - 1$ (at the coarsest time scale) to $j = 0$ (at the finest time scale). Also, we let r_a (r_b) decrease from R_a (R_b) (the ranges of coarsest time scales), to 1 (the range of finest time scales). Thus, the number of time intervals at time scale j , which we will denote by N_j , is 2^{S-j-1} . Moreover, assuming a unitary width for the intervals of the finest time scale, $j = 0$, the width at scale j will be 2^j . To relate time scales and time scale ranges we define j_{r_a} and j_{r_b} as the coarsest scales j in ranges r_a and r_b , respectively. Thus, in Figure 1, $S = 4$, $R_a = 2$, $R_b = 2$, $j_2 = 3$ and $j_1 = 1$ for arrivals, and $j_2 = 3$ and $j_1 = 2$ for sizes.

In order to assure that the mean packet arrival rate and the mean packet size (computed over all time intervals) are the same in all time scales, so as to maintain




Figure 1: Construction of L-System based traffic model.

physical meaning, we will impose the following conditions to the production rules:

$$X_{(j,r_a)}^{(i)} = \frac{1}{2}X_{(j-1,r'_a)}^{(2i-1)} + \frac{1}{2}X_{(j-1,r'_a)}^{(2i)} \quad (3)$$

$$Y_{(j,r_b)}^{(i)} = \frac{1}{2}Y_{(j-1,r'_b)}^{(2i-1)} + \frac{1}{2}Y_{(j-1,r'_b)}^{(2i)} \quad (4)$$

i.e., the mapping of arrival rates (mean sizes) is such that the arrival rate (mean size) averaged over the left and right child subintervals will be equal to the parent arrival rate (mean size).

The generation of arrival rates and mean sizes can be described by axioms $X_{(S-1,R_a)}^{(1)}$ and $Y_{(S-1,R_b)}^{(1)}$, the arrival rate and mean size at the coarsest time scale, and by production rules defined by

$$X_{(j,r_a)}^{(i)} = \lambda_l \xrightarrow{p_{lq}^{(r_a)}} \begin{cases} X_{(j-1,r'_a)}^{(2i-1)} = \lambda_q \\ X_{(j-1,r'_a)}^{(2i)} = 2\lambda_l - \lambda_q \end{cases} \quad (5)$$

and

$$Y_{(j,r_b)}^{(i)} = \gamma_m \xrightarrow{v_{mn}^{(r_b)}} \begin{cases} Y_{(j-1,r'_b)}^{(2i-1)} = \gamma_n \\ Y_{(j-1,r'_b)}^{(2i)} = 2\gamma_m - \gamma_n \end{cases} \quad (6)$$

where $\sum_{l=1}^L p_{lq}^{(r_a)} = 1, \forall q$, and $\sum_{m=1}^G v_{mn}^{(r_b)} = 1, \forall n$. Thus, an arrival rate λ_l (mean size γ_m) in interval i , scale j and range r_a (r_b) produces, with probability $p_{lq}^{(r_a)} (v_{mn}^{(r_b)})$, arrival rate λ_q (mean size γ_n) at the left subinterval $2i-1$ and arrival rate $2\lambda_l - \lambda_q$ (mean size $2\gamma_m - \gamma_n$) at the right subinterval $2i$, of next scale $j-1$ and range r'_a (r'_b). The production rules of packet arrival rates can be totally described by $R_a L \times L$ matrices

$$\mathbf{P}^{(r_a)} = \left(p_{lq}^{(r_a)} \right), l, q = 1, \dots, L, r_a = 1, \dots, R_a \quad (7)$$

and the production rules of mean packet sizes can be totally described by $R_b G \times G$ matrices

$$\mathbf{V}^{(r_b)} = \left(v_{mn}^{(r_b)} \right), m, n = 1, \dots, G, r_b = 1, \dots, R_b \quad (8)$$

In order to guarantee that the alphabet of packet arrival rates is closed with respect to the production rules we impose the following conditions: (i) $\lambda_i - \lambda_{i-1} = \frac{\lambda_L - \lambda_1}{L-1}, i = 2, 3, \dots, L$, i.e., the λ_i values will be equidistant; (ii) $p_{lq}^{(r_a)} = 0$ if $q > l + \min(L-l, l-1)$ or $q < l - \min(L-l, l-1)$. Similar conditions have to be imposed on mean packet sizes. Thus, (i) $\gamma_i - \gamma_{i-1} = \frac{\gamma_G - \gamma_1}{G-1}, i = 2, 3, \dots, G$; (ii) $v_{mn}^{(r_b)} = 0$ if $n > m + \min(G-m, m-1)$ or $n < m - \min(G-m, m-1)$.

Finally, the traffic process at scale j can be defined by the time series

$$Z_{(j)} = \{(X_{(j,r_a)}^{(i)}, Y_{(j,r_b)}^{(i)}), i = 1, \dots, N_j\} \quad (9)$$

4 Fitting Procedure

The fitting procedure determines the parameters of the two L-Systems from real data observations. It starts by fixing a sampling interval Δ and considering the time series representing the number of packet arrivals and the mean packet size in each non-overlapping sampling interval. Let this (empirical) time series be $\{(A_k, B_k), k = 1, 2, \dots, K\}$, where A_k and B_k represent the number of arrivals and the mean packet size in sampling interval k . For convenience, we take the length of the time series K to be a power of 2. The inference procedure can then be divided in three steps: (i) determination of alphabet and axiom of arrival rates and mean sizes, (ii) identification of the time scale ranges for each L-System and, (iii) inference of the production rules of arrival rates and mean sizes. The flow diagram of each L-System fitting procedure is represented in figure 2.

The alphabets of both L-Systems will consist in L and G equidistant values, respectively, ranging from the minimum to the maximum ones present in data. Both L and G are a compromise between accuracy and complexity. Due to the mass preservation property of L-Systems defined in (3) and (4), the axiom of packet arrival rates is inferred as the mean arrival rate of $\{A_k\}$, rounded to the closest alphabet element, i.e.,

$$X_{(S-1, R_a)}^{(1)} = \Lambda_{\vec{\lambda}} \left((1/K\Delta) \sum_{k=1}^K A_k \right) \quad (10)$$

and the axiom of mean packet sizes is inferred as the mean of $\{B_k\}$, rounded to the closest alphabet element, i.e.,

$$Y_{(S-1, R_b)}^{(1)} = \Lambda_{\vec{\gamma}} \left((1/K) \sum_{k=1}^K B_k \right) \quad (11)$$

where $\Lambda_{\vec{\alpha}}(x)$ represents a function that rounds x towards the nearest element of $\vec{\alpha}$.

The identification of time scale ranges is based on wavelet scaling analysis. We use the method described in [6], which resorts to the (second-order) logscale diagram. A (second-order) logscale diagram is a plot of y_j against j , together with confidence intervals about the y_j , where y_j is a function of the wavelet discrete transform coefficients at scale j . The time scale ranges correspond to the set of time scales for which, within the limits of the confidence intervals, the y_j fall on a straight line, i.e., the scaling behavior is linear in a time scale range. This procedure is applied to both the packet arrival and the packet size processes. Figure 7 shows the logscale diagram of the packet arrival process of a trace measured at the University of Aveiro (which is described in section 6). It can be seen that there are 4 time scale ranges (within a total of 18 time scales) defined by $j_1 = 3, j_2 = 7, j_3 = 10$ and $j_4 = 17$.

The next step is the inference of the production rules of the two L-Systems, which are fully characterized by the $\mathbf{P}^{(r_a)}$ and $\mathbf{V}^{(r_b)}$ matrices. First, data is rounded in order to define the sequences of packet arrival rates and mean packet sizes at each time scale. This comprises obtaining the arrival rates $X_{(j,r_a)}^{(i)}$ from $\{A_k\}$ through

$$X_{(j,r_a)}^{(i)} = \Lambda_{\vec{\lambda}} \left((N_j/K\Delta) \sum_{k=K(i-1)/N_j+1}^{Ki/N_j} A_k \right) \quad (12)$$

and the mean sizes $Y_{(j,r_b)}^{(i)}$ from $\{B_k\}$ through

$$Y_{(j,r)}^{(i)} = \Lambda_{\vec{\gamma}} \left((N_j/K) \sum_{k=K(i-1)/N_j+1}^{Ki/N_j} B_k \right) \quad (13)$$

with $i = 1, \dots, N_j$, for each j .

Letting $c_{lq}^{(r_a)}$ represent the number of times that, over all scales of range r_a , the parent $X_{(j,r_a)}^{(i)} = \lambda_l$ produced the left child $X_{(j-1,r'_a)}^{(2i-1)} = \lambda_q$, the production rule probabilities for arrival rates can be inferred as

$$p_{lq}^{(r_a)} = c_{lq}^{(r_a)} / \sum_{q=1}^L c_{lq}^{(r_a)}, \quad l = 1, \dots, L, \quad r_a = 1, \dots, R_a \quad (14)$$

Similarly, letting $d_{mn}^{(r_b)}$ represent the number of times that, over all scales of range r_b , the parent $Y_{(j,r_b)}^{(i)} = \gamma_m$ produced the left child $Y_{(j-1,r'_b)}^{(2i-1)} = \gamma_n$, the production rule probabilities for mean sizes can be inferred as

$$v_{mn}^{(r_b)} = d_{mn}^{(r_b)} / \sum_{n=1}^G d_{mn}^{(r_b)}, \quad m = 1, \dots, G, \quad r_b = 1, \dots, R_b \quad (15)$$

5 Relation with Random Cascades

A random cascade is characterized by an initial mass, uniformly distributed over a single interval, which is subdivided along the different stages of the cascade construction. Each interval is divided in two (or more) identical subintervals and the mass is randomly assigned to each subinterval, according to a random variable W called the generator. In the case of a traffic process the mass can be interpreted as the number of arrivals or bytes within a time interval. Let $W_{i,j}$, $i = 1, \dots, N_j$, $j = S-1, \dots, 0$, denote an (independent) random variable, having the same distribution as W , that redistributes mass from time scale j into time interval i (belonging to subsequent time scale $j-1$). The construction of the traffic process, starts at the coarsest time scale $S-1$ with an initial mass M distributed uniformly over a unit time interval. We will restrict our discussion to the case where a (parent) interval is subdivided in only two (child) intervals. In the first iteration, a finer time scale is produced, by dividing the time interval in two new subintervals of length $1/2$, and assigning mass $MW_{1,S-1}$ to the left subinterval and $MW_{2,S-1}$ to the right subinterval. At the second iteration,

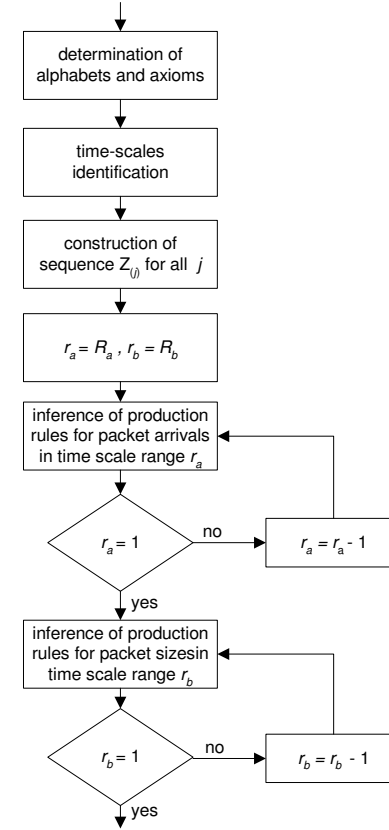


Figure 2: Flow diagram of inference procedure.

producing time scale $S-2$, each of these intervals generates new subintervals, one at the left and another at the right, giving rise to four subintervals of length $1/4$ with masses $MW_{1,S-1}W_{1,S-2}$, $MW_{1,S-1}W_{2,S-2}$, $MW_{2,S-1}W_{3,S-2}$ and $MW_{2,S-1}W_{4,S-2}$, respectively. Iterating this construction process, the mass of parent interval i from scale j is redistributed into child subintervals $2i-1$ and $2i$ of scale $j-1$ with probabilities $W_{2i-1,j}$ and $W_{2i,j}$, respectively. Note that the mass at each time scale is only preserved in expectation. Note also that the way mass gets redistributed is independent from mass itself, since the $W_{i,j}$ are independent of mass.

Feldmann et al. [3] have proposed the use of conservative cascades, a special case of random cascades, as a model for IP traffic. In conservative cascades, the generator W takes on values in $(0, 1)$, has mean $1/2$ and is symmetric about its mean. Furthermore, mass is redistributed such that the total mass assigned to left and right child subintervals remains equal to that of the parent interval. Thus, the mass is preserved throughout the splitting process: if the mass of the i^{th} parent interval is Q , the mass of the left child interval will be $QW_{2i-1,j}$ and that of the corresponding right child interval will be $Q(1 - W_{2i-1,j})$. As a result, the mass at all stages will be (exactly) M , if the initial mass is M . This is the mass preservation property which also occurs in our stochastic L-System constructions. Note that a different generator can be used for each timescale, leading to multiscaling behavior. Feldmann et al. supported the adoption of conservative cascade models in the networking context by observing that the transmitted traffic is constructed through fragmentation at successive network layers, and that the total number of bytes is roughly preserved during this fragmentation process. A typical example is the dynamics of a Web session, where user clicks results in requests, requests give rise to connections, connections are made up of flows, and flows consist of individual packets. Gao and Rubin [10] considered a traffic model with two conservative cascades, which characterize the packet arrival and the packet size processes, respectively.

In L-Systems, the mass in interval i of scale j is $2^j X_{(j,r_a)}^{(i)}$ for packet arrivals and $2^j Y_{(j,r_b)}^{(i)}$ for packet sizes. In addition to the mass preservation property, L-Systems include another important feature that is not present in conservative cascades (nor in random cascades) and has a meaningful physical explanation. In L-Systems the way mass is redistributed to left and right child subintervals can be made dependent on the mass of the parent interval, whereas in the conservative cascade construction this dependence is not allowed. Recall that, for example in the case of packet arrivals, a parent arrival rate λ_l gives rise to a left child arrival rate λ_q with probability $p_{lq}^{(r_a)}$, which depends explicitly on l . Take again the example of Feldmann et al. and consider the way Web requests are scheduled over time. The number of requests per time interval that a Web server can handle is necessarily limited, due to resource availability constraints. Therefore, the way user clicks produce requests depends on the overall number of clicks. If a Web server has more requests to process it will distribute them more sparsely over time. Thus, the mass itself (the number of clicks, in this case) influences the way distribution takes place. At a lower level, consider the way flows produce packets. If the network is more congested, the feedback control exercised by TCP imposes that packets will be more sparsely distributed over time. Thus, once more,

the mass (the number of flows, in this case) influences the way distribution takes place.

To bring further insight into this issue, we next evaluate whether or not the independence property of the generator, subjacent to conservative cascades, is present in data, using the pOct.TL Bellcore trace. The generator, as defined in the context of conservative cascades, represents the fraction of mass distributed to a left child subinterval. In L-Systems this fraction is $2^{j-1} X_{(j-1,r'_a)}^{(2i-1)} / 2^j X_{(j,r_a)}^{(i)} = \lambda_l / 2\lambda_q$ for packet arrivals and $2^{j-1} Y_{(j-1,r'_b)}^{(2i-1)} / 2^j Y_{(j,r_b)}^{(i)} = \gamma_m / 2\gamma_n$ for packet sizes. We define $W_{a_{j|l}}$ as the generators for the sequence of arrival rates, at time scale j , conditioned on the arrival rate λ_l and $W_{b_{j|m}}$ as the generators for the sequence of mean sizes, at time scale j , conditioned on the mean size γ_m . Thus, the $W_{a_{j|l}}$ and $W_{b_{j|m}}$ are discrete random variables; for each l $W_{a_{j|l}}$ can take values $\lambda_l / 2\lambda_q$ and for each m $W_{b_{j|m}}$ can take values $\gamma_m / 2\gamma_n$, and both can be easily inferred from data. Clearly, if the independence assumption is true then, at each j , the $W_{a_{j|l}}$, $l = 1, \dots, L$ should have the same distribution, and the same argument applies to $W_{b_{j|m}}$, $m = 1, \dots, G$. In Figure 3 and Figure 4 we plot the mean and variance of the inferred conditional $W_{a_{j|l}}$, for all λ_l and for $j = 1, 2, 3, 4$ and 5 . In Figure 5 and Figure 6 we plot the mean and variance of the inferred conditional $W_{b_{j|m}}$, for all γ_m and for $j = 1, 2, 3, 4$ and 5 . Results show that, within the same time scale, the generators can have very different variances and mean values, both in packet arrivals and packet sizes. For example, at $j = 0$ the $W_{a_{j|l}}$ variance is 0.04 for $\lambda_l = 21$ pkts/sec and is 0.001 for $\lambda_l = 100$ pkts/sec.

Also, at $j = 0$ the $W_{b_{j|m}}$ variance is 0.0516 for $\gamma_m = 341$ bytes and is 3.51×10^{-3} for $\gamma_m = 701$ bytes. Therefore, they have dissimilar distributions which indicates that the independence assumption is not valid in the case of the pOct.TL Bellcore trace. We have observed this same behavior in many other traces. To conclude, the effect of resource availability limitations is an important factor with strong impact in the traffic generation process, which can be captured through an L-System based construction (but not through a conservative cascade).

6 Numerical Results

We have applied our fitting procedure to two traces of IP traffic: (i) the well known pOct.TL Bellcore trace [15] and (ii) one trace measured at the University of Aveiro (UA) which exhibits non-trivial multifractal scaling behavior. The sampling interval was 0.1 seconds in both traces. The UA trace is representative




Figure 3: $W_{a_{j|l}}$ mean versus λ_l and j , arrival rates of trace pOct.TL.




Figure 4: $W_{a_{j|l}}$ variance versus λ_l and j , arrival rates of trace pOct.TL.

of Internet access traffic produced within a University campus environment. The University of Aveiro is connected to the Internet through a 10 Mb/s ATM link and the measurements were carried out in a 100 Mb/s Ethernet link connecting the border router to the firewall, which only transports Internet access traffic. The UA trace consists of 20 millions packets captured on July 3rd 2001 from 8.00 pm to 3.15 am. The traffic analyzer was a 1.2 GHz AMD Athlon PC, with 1.5 Gbytes of RAM, running WinDump. The measurements recorded the arrival instant and the IP header of each packet. The mean arrival rate of the UA trace is 766 pkts/sec.




Figure 5: $W_{b_{j|m}}$ mean versus γ_m and j , mean sizes of trace pOct.TL.




Figure 6: $W_{b_{j|m}}$ variance versus γ_m and j , mean sizes process of trace pOct.TL.

The packet arrival process of the pOct.TL trace was fitted to a stochastic L-System with an alphabet of $L = 243$ arrival rates, from the minimum to the maximum present in data, in steps of 10 pkts/sec (the minimum and maximum were 10 pkts/sec and 2430 pkts/sec, respectively). The logscale diagram identified 5 time scale ranges (within a total of 14 time scales) defined by $j_1 = 3, j_2 = 6, j_3 = 8, j_4 = 9$ and $j_5 = 13$. The packet size process was fitted to an alphabet of $G = 50$ packet sizes from 64 bytes to 1421 bytes, in steps of 27 bytes. The logscale diagram identified 3 time scale ranges defined by $j_1 = 3, j_2 = 8$ and




Figure 7: Logscale diagram, arrival rates of trace UA.

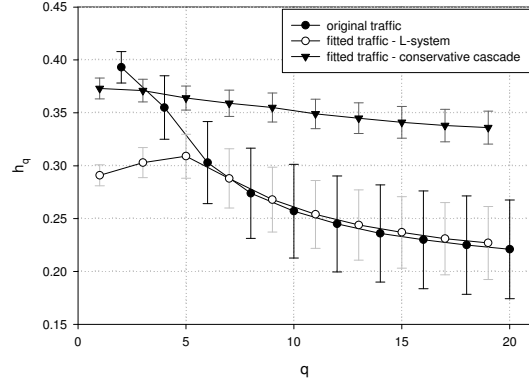


Figure 8: Linear multiscale diagram, arrival rates of trace UA.

$j_3 = 13$.

In the case of the UA trace the alphabet fitted for packet arrivals was 170 pkts/sec to 4850 pkts/sec, in steps of 10 pkts/sec ($L = 469$). As referred before, the UA trace has 4 time scale ranges (Figure 7). The alphabet of packet sizes was 66 bytes to 1214 bytes, in steps of 23 bytes ($G = 50$). In this case, the logscale diagram identified 3 time scale ranges defined by $j_1 = 3$, $j_2 = 10$ and $j_3 = 17$.

The parameter estimation took less than 1 minute, using a MATLAB implementation running in the PC described above. This shows that the fitting

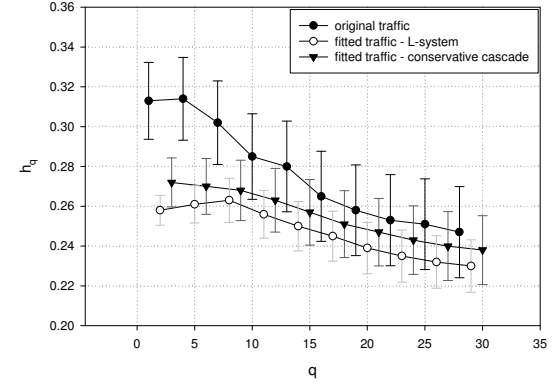


Figure 9: Linear multiscale diagram, mean sizes of trace UA.




Figure 10: Probability function of byte counts, trace pOct.TL.

procedure is computationally very efficient (note that the size of the alphabet, the number of ranges and the size of the trace, which determine the computational time, are all relatively large).

We assess the suitability of the traffic model and the accuracy of the fitting procedure using several criteria. We compare both the probability and autocovariance functions of the byte counts (number of byte arrivals in sampling interval) obtained with the fitted stochastic L-System model and with the original data trace. Note that the sequence of byte counts is an aggregation of the




Figure 11: Autocovariance of byte counts, trace pOct.TL.

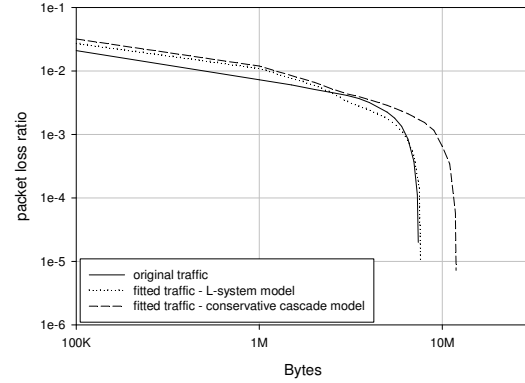


Figure 12: Packet loss ratio versus buffer size, trace pOct.TL.

sequences of packet arrivals and packet sizes. We also analyze the queuing behavior by comparing the packet loss ratio, obtained through trace-driven simulation, using two types of input traffic: (i) the original trace and (ii) the trace generated according to the fitted stochastic L-System. Our comparisons are extended to a double conservative cascade model. The model was inferred using the procedure presented in [7], using one cascade to model the packet arrivals and another to model the packet sizes. Here, the generators at each time scale are fitted to a truncated normal distribution, where the mean is always 1/2 but the variance

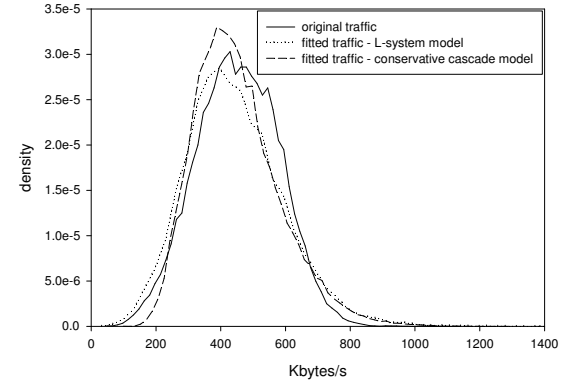


Figure 13: Probability function of byte counts, trace UA.

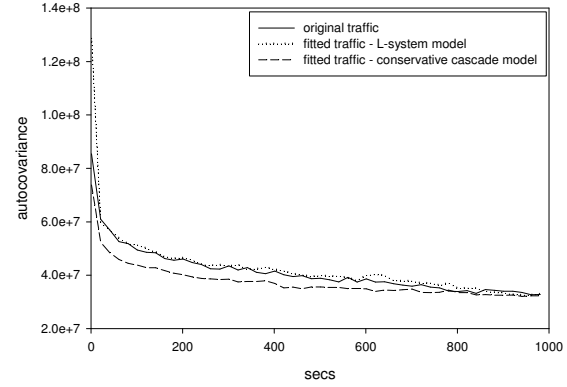


Figure 14: Autocovariance of byte counts, trace UA.

is adjusted individually at each time scale.

Multifractality is assessed using a linear multiscale diagram [6], which plots $h_q = \alpha_q/q$ against q ; α_q represents the q^{th} order scaling exponent, estimated in the q^{th} order logscale diagram. Multifractal scaling behavior is detected when there is no horizontal alignment (within the limits of confidence intervals). Figure 8 and Figure 9 show that the packet arrival and packet size processes of the UA trace, and of the traces generated using both the L-System and conservative cascade models, all have multifractal scaling behavior. A similar analysis carried

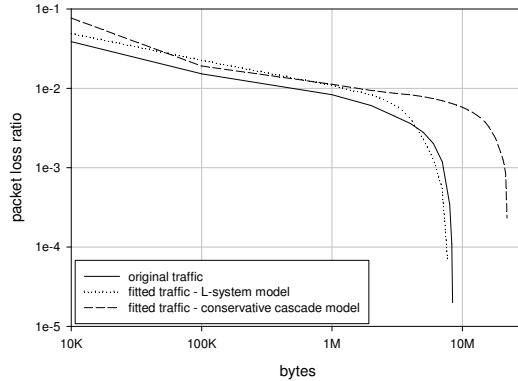


Figure 15: Packet loss ratio versus buffer size, trace UA.

out on the pOct.TL trace revealed that this trace is not multifractal.

In the case of the probability function (Figure 13), both the conservative cascade and the L-System models fitted very well the UA trace. The performance of the conservative cascade model was not so good in the case of the pOct.TL trace (Figure 10). This can be explained by the symmetry imposed on the generator W , which makes it difficult for conservative cascades to fit asymmetric probability distributions. In the case of the autocovariance function (Figure 11 and Figure 14), the fitting performance achieved by both models was again very good in the case of the UA trace, but clearly inferior for conservative cascades in the case of the pOct.TL trace.

To assess the queuing behavior the buffer size was varied from 100 Kbytes to 20 Mbytes. The service rate was 518 Kbytes/s for the pOct.TL trace (corresponding to an utilization of 0.7) and 512 Kbytes/s for the UA trace (corresponding to an utilization of 0.9). Figure 12 and Figure 15 show that, for both traces, the fitting of the queuing behavior was very good in the case of L-Systems but significant differences occurred with conservative cascades. It is interesting to note that, in the case of the UA trace, these differences occurred despite of a very good fitting in the first and second order statistics. This clearly illustrates that first and second order statistics are not sufficient to characterize multifractal traffic.

We also have analyzed a number of other traces. Our results show that, in general, L-Systems achieves better performance than conservative cascades. This can be explained as follows. First, L-Systems allow the mass redistribution to depend on the mass itself, a feature that is clearly present in real observed data.

Second, L-Systems provide a higher number of parameters, which are meaningful from the point of view of physical reality. Third, the generator in conservative cascades is assumed to be symmetric which restricts the fitting of the probability function.

7 Conclusions

In this paper we proposed a novel traffic model, and an associated parameter fitting procedure, that jointly describes the packet arrival and the packet size processes, allowing for the incorporation of multifractal scaling behavior in both characteristics. The model is based on so-called stochastic L-Systems, which were introduced by biologist A. Lindenmayer as a method to model plant growth. We provided a detailed comparison with a related multifractal model based on conservative cascades. Our results, that include applying the fitting procedure to real observed data with multifractal scaling behavior on both the packet arrival and packet size processes, showed that L-System based models can achieve excellent fitting performance in terms of first and second order statistics and queuing behavior.

Acknowledgments

This work was part of project POSI/CPS/42069/2001 "Traffic Modeling and Performance Evaluation of Multiservice Networks", funded by Fundação para a Ciência e Tecnologia, Portugal. P. Salvador wishes to thank Fundação para a Ciência e a Tecnologia, Portugal, for support under grant BD/19781/99.

References

- [1] B. Mandelbrot, "Intermittant turbulence in self-similar cascades: Divergence of high moments and dimensions of the carrier," *Journal of Fluid Mechanics*, vol. 62, pp. 331–358, 1974.
- [2] R. Riedi and J. L. Véhel, "Multifractal properties of TCP traffic: a numerical study," *Technical Report No 3129, INRIA Rocquencourt, France*, Feb. 1997, Available at www.dsp.rice.edu/~riedi.

- [3] A. Feldmann, A.C. Gilbert, and W. Willinger, "Data networks as cascades: Investigating the multifractal nature of internet WAN traffic," in *Proceedings of SIGCOMM*, 1998, pp. 42–55.
- [4] R. Riedi, M. Crouse, V. Ribeiro, and R. Baraniuk, "A multifractal wavelet model with application to network traffic," *IEEE Transactions on Information Theory*, vol. 45, no. 4, pp. 992–1018, April 1999.
- [5] A. Feldmann, A.C. Gilbert, P. Huang, and W. Willinger, "Dynamics of IP traffic: A study of the role of variability and the impact of control," in *SIGCOMM*, 1999, pp. 301–313.
- [6] P. Abry, P. Flandrin, M. Taqqu, and D. Veitch, "Wavelets for the analysis, estimation and synthesis of scaling data," in *Self-Similar Network Traffic Analysis and Performance Evaluation*, K. Park and W. Willinger, Eds., 1999.
- [7] A. Gilbert, W. Willinger, and A. Feldmann, "Scaling analysis of conservative cascades, with applications to network traffic," *IEEE Transactions on Information Theory*, vol. 45, no. 3, pp. 971–992, April 1999.
- [8] A. Erramilli, O. Narayan, A. Neidhardt, and I. Saniee, "Performance impacts of multi-scaling in wide area TCP/IP traffic," in *Proceedings of INFOCOM'2000*, 2000.
- [9] D. Veitch, P. Abry, P. Flandrin, and P. Chainais, "Infinitely divisible cascade analysis of network traffic data," in *Proceedings of the Int. Conf. on Acoust. Speech and Sig. Proc.*, June 2000.
- [10] J. Gao and I. Rubin, "Multifractal modeling of counting processes of long-range-dependent network traffic," *Computer Communications*, vol. 24, pp. 1400–1410, 2001.
- [11] P. Salvador, A. Nogueira, and R. Valadas, "Modeling multifractal traffic with Stochastic L-Systems," in *Proceeding of GLOBECOM 2002*, Nov. 2002.
- [12] A. Lindenmayer, "Mathematical models for cellular interactions in development II. Simple and branching filaments with two-sided inputs," *Journal of Theoretical Biology*, vol. 18, pp. 300–315, 1968.
- [13] H. Peitgen, H. Jurgens, and D. Saupe, *Chaos and Fractals: New Frontiers of Science*, Springer-Verlag, 1992.

- [14] P. Salvador, A. Nogueira, and R. Valadas, "Scaling behavior analysis of multifractal traffic models based on stochastic L-systems," Tech. Rep., University of Aveiro, 2002.
- [15] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of ethernet traffic (extended version)," *IEEE/ACM Transactions on Networking*, vol. 2, no. 1, pp. 1–15, Feb. 1994.

Biographies

Paulo Salvador was born in Coimbra, Portugal, on February 17, 1975. He graduated in Electronics and Telecommunications Engineering in 1998 from University of Aveiro, Aveiro, Portugal. In 1998, he joined Institute of Telecommunications, Aveiro, Portugal, as a researcher. He is currently a PhD. student at University of Aveiro. His main research areas of interest are traffic modeling, network planning and Internet security.

António Nogueira was born in Mozambique, on March 12, 1969. He graduated in Electronics and Telecommunications Engineering in 1992 and received the MSc. Degree in Electronics and Telecommunications Engineering in 1997, both from University of Aveiro, Aveiro, Portugal. In 1995, he joined Institute of Telecommunications, Aveiro, Portugal, as a researcher. He is currently an Assistant and a PhD. student at University of Aveiro. His main research areas of interest are traffic modeling and network planning.

Rui Valadas graduated in Electrical and Computers Engineering, in 1986, from the Instituto Superior Técnico, Lisbon, Portugal, and received the Ph.D. degree from the University of Aveiro, Aveiro, Portugal, in 1996. He joined the University of Aveiro in 1986, where he is now an Associate Professor of the Department of Electronics and Telecommunications. He is also leader of the Networks and Multimedia Communications Group at the Institute of Telecommunications, Aveiro, Portugal. His main research interests are in the area of traffic engineering.