A LOCAL STATIONARY LONG-MEMORY MODEL FOR INTERNET TRAFFIC

Li Song and Pascal Bondon

CNRS UMR 8506, Université Paris-Sud
91192, Gif-sur-Yvette, France
phone: +33 1 69 85 17 31, fax: +33 1 69 85 17 65, email: bondon@lss.supelec.fr

ABSTRACT

We present in this paper a piecewise fractional autoregressive integrated moving average (FARIMA) model and a procedure to fit this model to local-stationary traffic data. The procedure consists in finding the number as well as the locations of structural break points in the series and estimating the orders and the parameters of each segment. The effectiveness of the procedure is illustrated by Monte Carlo simulations. An application to real internet traffic data is considered and shows that the piecewise FARIMA model is able to capture the non-stationarity and the long-memory of these data.

1. INTRODUCTION

During the last years, many researchers in signal processing, computer science, and applied probability have focused on models for internet traffic. Previous works have shown that the classical models for telephone traffic cannot be applied to internet traffic [8, 10]. After the seminal study of [8], increasing evidence has been put up for the failure of traditional (Poisson-based) models to account for the long-range dependence (LRD) present at the large time scales in network traffic. At the same time, due to the many different network mechanisms and various source characteristics, short-range dependence also exists and plays a central role [2, 4, 6]. Therefore a model like the famous fractional autoregressive integrated moving average (FARIMA) model [7], is required to describe both short and long memories simultaneously. But accurate estimation of a FARIMA model often requires a large sample of data taken over a long period of time, which in turn increases the chance of structural breaks over time. From the real data test [14, 12], the assumption that the data can be modeled by a stationary process with constant parameters may be unrealistic. Previous models, like fractional Brownian motion [9] and FARIMA model fail to analyze structural changes that are ubiquitous in network traffic. Furthermore, less attention has been paid to model LRD in presence of several structural breaks, most common methodologies being designed to detect a single break point (BP). Some studies discuss multiple BPs but they consider only changes in the long-memory parameter and the mean level. For example, [11] assumes that the ARMA parameters are constant which may be unrealistic in practice. Therefore, there is a need for a model with more flexibility in modeling non-stationarity.

The objective of this paper is to propose a parametric model which inherits the advantages of FARIMA processes, to model non-stationary long-memory internet traffic. We choose a piecewise FARIMA process assuming that the signal can be divided into sub-segments that are essentially stationary. The remainder of this paper is organized as follows. In Section 2, the model is presented, and in Section 3, the fitting procedure is described. Numerical simulation results are presented and discussed in Section 4. A real traffic data modeling is considered in Section 5 and concluding remarks can be found in Section 6.

2. MODEL DESCRIPTION

We suppose that the non-stationary process $\{Y_t\}$, $t = 1,\dots,n$, can be segmented into $m+1$ blocks of stationary FARIMA processes. For $j = 1,\dots,m$, denote the BP between the jth and $(j+1)$th FARIMA processes as τ_j, and set $\tau_0 = 1$ and $\tau_{m+1} = n+1$. Then the jth block of $\{Y_t\}$ is modeled by

$$Y_t = X_{t,j}, \qquad \tau_{j-1} \le t < \tau_j, \tag{1}$$

where $\{X_{t,j}\}$ is the FARIMA(p_j, d_j, q_j) process defined by the difference equation

$$\Phi_j(B)X_{t,j} = \Theta_j(B)(1-B)^{-d_j}\epsilon_t, \tag{2}$$

B is the backward operator $BX_t = X_{t-1}$, $\{\varepsilon_t\}$ is a sequence of zero-mean iid random variables with finite variance σ_ε^2, $d_j \in (0, 1/2)$, and the polynomials $\Phi_j(z) = 1-\phi_{j,1}z-\cdots-\phi_{j,p_j}z^p$ and $\Theta_j(z) = 1+\theta_{j,1}z+\cdots+\theta_{j,q_j}z^q$ with real coefficients have no common zeros and neither $\Phi_j(z)$ nor $\Theta_j(z)$ has zeros in the closed unit disk $\{z \in \mathbb{C} : |z| \le 1\}$. The process $(1-B)^{-d}\epsilon_t$ is defined by

$$(1-B)^{-d}\epsilon_t = \textstyle\sum_{k=0}^{\infty} \varphi_k(d)\epsilon_{t-k}, \tag{3}$$

where $\varphi_0(d) = 1$ and $\varphi_k(d) = \prod_{s=1}^{k} \frac{d+s-1}{s}$ for $k \ge 1$. Since $d < 1/2$, $\sum_{k=0}^{\infty} \varphi_k(d)^2 < \infty$ and the series in (3) converges in the mean square sense.

Let $p \ge \max(p_j)$, $q \ge \max(q_j)$, $\alpha_j = (d_j, \phi_{j,1},\dots,\phi_{j,p}, \theta_{j,1},\dots,\theta_{j,q})$ where $\phi_{j,k} = 0$ for $k > p_j$ and $\theta_{j,k} = 0$ for $k > q_j$, and $\beta_j = (p_j, q_j, \alpha_j)$. Vector β_j contains the parameters of the jth model and β_j is constant on each interval $[\tau_{j-1}, \tau_j)$.

3. ESTIMATION PROCEDURE

The problem of fitting model (1)–(2) to data consists in finding $(\tau_1, \ldots, \tau_m, \beta_1, \ldots, \beta_{m+1})$. The first problem is to estimate the BPs accurately, which can be realized by detecting the structural parameters changes. It was shown in [13] that some of the best available techniques to estimate the parameters may be misled by non-stationary characters of the observed time series, and some of these non-stationarity effects can often be alleviated by estimating the parameters using data locally. That is to say, it is better to divide the original time series into a set of elementary sub-series of length E and use the data in the same sub-series to get a local parameter estimation. Predefining a suitable length E for the elementary sub-series is not always an easy task: on the one hand, due to LRD properties, a reasonable number of observations are needed to obtain adequate parameter estimates, and then E can't be too short; on the other, the probability of meeting a BP increases as E grows. E should vary with the hidden true model order and the estimate method, e.g. the Whittle estimates ([5]) and the Whittle estimates with autoregressive truncation ([1]). Hence some restriction should be put on E, and E is chosen by empirical experience.

In the following, we consider the truncated series formed by the $K = [n/E]$ elementary sub-series defined on the intervals $I_k = ((k-1)E, kE]$ for $k = 1, \ldots, K$. We make the following assumptions:

1. the number m of BPs of the truncated series is known and $m \leq K/2$ (see Remark 1 when m is unknown),
2. there is at most one BP in each interval,
3. at least one interval separates two consecutive BPs.

Then, the m BPs τ_j, $j = 1, \ldots, m$, are dispersed into a few elementary intervals.

The following four steps procedure is proposed to fitting model (1)–(2) to a local stationary time series.

Step 1 : Local estimation. For each interval I_k, a pair $(\hat{p}_k, \hat{q}_k)$ is selected by employing the Bayes Information Criterion (BIC) as suggested in [15], and the model's parameters α_k are estimated by the Gaussian maximum-likelihood estimates (MLE) based on autoregressive approximations $\hat{\alpha}_k$, see e.g. [7]. Therefore, Step 1 gives the local estimates $\hat{\beta}_k = (\hat{p}_k, \hat{q}_k, \hat{\alpha}_k)$ for $k = 1, \ldots, K$.

Step 2 : Selection of intervals with a BP. If model (1)–(2) is suitable for the data, one expects that $\hat{\beta}_k$ is close to the true values of the parameters when there is no BP in the interval I_k. Now, if there is a BP in I_k and no BP in I_{k-1} and I_{k+1}, $\hat{\beta}_k$ should be significantly different from both $\hat{\beta}_{k-1}$ and $\hat{\beta}_{k+1}$. Then, let $k_0 = 0$, $k_{m+1} = K$, and

$$(\hat{k}_1, \ldots, \hat{k}_m) = \underset{1 \leq k_1 < \cdots < k_m < K}{\operatorname{argmin}} \sum_{j=1}^{m+1} \sum_{k=k_{j-1}+1}^{k_j} \Big(\|\hat{\alpha}_k - \bar{\alpha}_j\|^2 + \psi(|\hat{p}_k - \bar{p}_j|) + \psi(|\hat{q}_k - \bar{q}_j|) \Big), \quad (4)$$

where $\bar{\alpha}_j = \dfrac{1}{k_j - k_{j-1}} \displaystyle\sum_{k=k_{j-1}+1}^{k_j} \hat{\alpha}_k$, $\bar{p}_j$ (resp. $\bar{q}_j$) is the order which is the most frequently selected among the orders $\hat{p}_k$ (resp. $\hat{q}_k$) for $k = k_{j-1}+1, \ldots, k_j$. In the case where $\bar{p}_j$ (resp. $\bar{q}_j$) is not unique, the lowest order is chosen. Function $\psi(\cdot)$ is positive and strictly increasing. Let $J_k = ((k-0.5)E, (k+0.5)E]$ for $k = 1, \ldots, K-1$. We select the intervals $(J_{\hat{k}_1}, \ldots, J_{\hat{k}_m})$ as being those containing a BP.

Step 3 : Estimation of the BPs. Suppose that all the intervals $J_{\hat{k}_j}$ are selected properly, i.e., $\tau_j \in J_{\hat{k}_j}$. Therefore, for any fixed j, there is no BP in the "previous" block between $J_{\hat{k}_{j-1}}$ and $J_{\hat{k}_j}$, viz. $((\hat{k}_{j-1}+0.5)E, (\hat{k}_j - 0.5)E]$ where we set $\hat{k}_0 + 0.5 = 0$, and we define $\hat{\beta}_p$ as the MLE of β_j based on the data in this block. In the same way, let $\hat{\beta}_n$ be the MLE of β_{j+1} based on the data in the "next" block between $J_{\hat{k}_j}$ and $J_{\hat{k}_{j+1}}$, viz. $((\hat{k}_j + 0.5)E, (\hat{k}_{j+1} - 0.5)E]$ where we set $\hat{k}_{m+1} - 0.5 = K$. We treat $\hat{\beta}_p$ and $\hat{\beta}_n$ as two benchmarks. These estimates are more precise than any local estimate calculated in Step 1 since they involve more data. Suppose that $l \in J_{\hat{k}_j}$ is the BP τ_j. Then we can calculate the MLE $\hat{\beta}_{l_p}$ of β_j and $\hat{\beta}_{l_n}$ of β_{j+1} based respectively on $((\hat{k}_{j-1}+0.5)E, l]$ and $(l, (\hat{k}_{j+1} - 0.5)E]$. These estimates should be close to benchmarks $\hat{\beta}_p$ and $\hat{\beta}_n$, respectively. Hence, our choice of the BP estimate $\hat{\tau}_j$ is based on the following criterion

$$\hat{\tau}_j = \underset{l \in J_{\hat{k}_j}}{\operatorname{argmin}} \Big(\|\hat{\alpha}_{l_p} - \hat{\alpha}_p\|^2 + \psi(|\hat{p}_{l_p} - \hat{p}_p|) + \psi(|\hat{q}_{l_p} - \hat{q}_p|) + \|\hat{\alpha}_{l_n} - \hat{\alpha}_n\|^2 + \psi(|\hat{p}_{l_n} - \hat{p}_n|) + \psi(|\hat{q}_{l_n} - \hat{q}_n|) \Big). \quad (5)$$

To reduce the complexity, $\hat{\beta}_{l_p}$ and $\hat{\beta}_{l_n}$ are calculated using the data in $(l-E, l)$ and $(l, l+E)$, respectively, and this gives good results in practice as shown in Section 4.

Step 4 : Estimation of the parameters of each stationary block. Once $(\hat{\tau}_1, \ldots, \hat{\tau}_m)$ are obtained, the parameters β_j of the stationary sequence $X_{t,j}$ for $j = 1, \ldots, m+1$, can be estimated on the basis on the data in $(\hat{\tau}_{j-1}, \hat{\tau}_j]$, where $\hat{\tau}_0 = 1$ and $\hat{\tau}_{m+1} = KE$.

Remark 1. The procedure assumes that the number m of BPs is known, but of course m is unknown in practice. One way to estimate m consists in increasing sequentially one by one the number of BPs in the procedure. Indeed, when estimating a single BP model in the presence of multiple BPs, the estimate of the interval which contains the BP will be typically one of the true intervals with a BP, namely the one which is dominant in the sense that selecting this interval allows to minimize the sum of squared (4) where $m = 1$. Next, we minimize (4) with $m = 2$ which gives two dominating intervals with a BP. When iterating this process beyond the true number of BPs, two intervals are founded which are very close to each other and correspond to the same BP. This allows to determine the true number of BPs when the BPs are not too close to each others. We show in Section 4 that this method for finding m works well when at least $2E$ data separate each BP.

4. SIMULATIONS

In this section, we illustrate the implementation procedure by a simple Monte Carlo experiment and show its effectiveness. This experiment is based on 100 replications of a piecewise FARIMA process $\{Y_t\}$, $t = 1, \ldots, n$ where $n = 40000$, $\{\epsilon_t\}$ is Gaussian with $\sigma_\epsilon^2 = 1$ in (2), and we have 4 BPs at $\tau_1 = 7800$, $\tau_2 = 16350$, $\tau_3 = 23550$ and $\tau_4 = 32100$. We set $E = 2000$, and therefore $K = 20$. The BPs fall within J_4, J_8, J_{12} and J_{16}. The orders and the parameters of the different segments $X_{t,j}$ in (2) are given in table 1 and figure 1 displays a typical realization of this model.

Parameters	Segment $X_{t,j}$				
β_j	1	2	3	4	5
p_j	1	0	1	1	0
q_j	2	0	0	1	1
d_j	0.20	0.40	0.10	0.30	0.15
ϕ_j	-0.7	-	-0.80	0.30	-
$\theta_{j,1}$	0.60	-	-	-0.70	0.40
$\theta_{j,2}$	-0.20	-	-	-	-

Table 1: Time-varying orders and parameters.

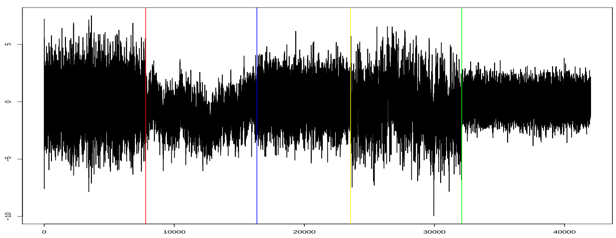

Figure 1: A realization of the piecewise FARIMA process with parameters in table 1. The vertical lines indicate the true BPs locations.

Table 2 exhibits the most frequently selected orders for each elementary sub-series in Step 1. Note that the four BPs locate in I_4, I_9, I_{12} and I_{17}, respectively. For most I_k's, these orders are the true ones. Observe that BIC performs well for detecting the low orders zero and one, and works less definitely for finding order two. Indeed, the true order $(1, 2)$ of the first three elementary sub-series is found in only 60% of the cases, while the orders zero and one of the others elementary sub-series are found in at least 75% of the cases. Of course, when there is a BP in the elementary sub-series and the order changes, the order given by BIC is not reliable (see the columns in bold in table 2).

Sub-series	1	2	3	**4**	5
Order $(\hat{p}_k, \hat{q}_k)$	(1,2)	(1,2)	(1,2)	**(2,2)**	(0,0)
Frequency	58	55	62	**38**	92
Sub-series	6	7	8	**9**	10
Order $(\hat{p}_k, \hat{q}_k)$	(0,0)	(0,0)	(0,0)	**(1,0)**	(1,0)
Frequency	89	88	83	**58**	84
Sub-series	11	**12**	13	14	15
Order $(\hat{p}_k, \hat{q}_k)$	(1,0)	**(2,0)**	(1,1)	(1,1)	(1,1)
Frequency	86	**49**	76	81	82
Sub-series	16	**17**	18	19	20
Order $(\hat{p}_k, \hat{q}_k)$	(1,1)	**(1,2)**	(0,1)	(0,1)	(0,1)
Frequency	78	**15**	90	88	91

Table 2: Selected orders in Step 1.

After the local estimation (Step 1), we calculate the intervals with a BP (Step 2) corresponding to the BPs numbers $m = 1, \ldots, 6$. For $m = 1, \ldots, 4$, all the selected intervals in the Monte Carlo experiments were well separated. For $m = 5, 6$, we founded intervals close to each others in all the experiments, which hints that the number of BPs should be 4. Table 3 gives the selected intervals in the 100 simulations for $m = 5$ (the last column indicates the number of times the quintuplet $(\hat{k}_1, \ldots, \hat{k}_5)$ were selected). We see that when the BPs number used in Step 2 is the true BPs number plus one, the additional interval chosen by the procedure is close to an interval containing a BP.

$\hat{k}_1$	$\hat{k}_2$	$\hat{k}_3$	$\hat{k}_4$	$\hat{k}_5$	%
3	4	8	12	16	2
4	5	8	12	16	13
4	7	8	12	16	11
4	7	12	16	17	1
4	8	11	12	16	48
4	8	12	16	17	25

Table 3: Selected intervals in Step 2 for $m = 5$.

Figure 2 shows the estimations of the intervals with a BP in Step 2 where $m = 4$. These estimations are very good since the intervals with highest frequencies are the right ones. Some selected intervals spread to a few intervals near the right ones. The estimation of the first interval diffuses from J_2 to J_6, that of the second interval from J_6 to J_{10}, that of the third interval from J_{10} to J_{14}, and the one of the fourth interval from J_{15} to J_{18}. For the fourth interval for instance, the true interval J_{16} is selected 83 times, and J_{15}, J_{17} and J_{18} are chosen 1, 13 and 3 times, respectively. This is partly caused by the impreciseness of the order selection in Step 1.

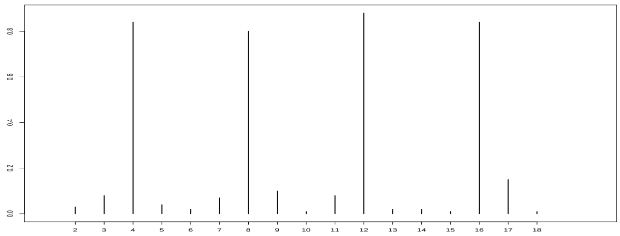

Figure 2: Selected intervals in Step 2: $m = 4$.

Table 4 presents the sample means $\hat{\mu}(\hat{\lambda}_j)$ and standard errors $\hat{\sigma}(\hat{\lambda}_j)$ of the BPs estimations in Step 3 over the 100 realizations. Following [3], we use the standardized parameter $\lambda_j = \tau_j/n$. We see that the estimations are close to the true values and the standard errors are quite small.

λ_j	0.1950	0.4088	0.5888	0.8025
$\hat{\mu}(\hat{\lambda}_j)$	0.1947	0.4074	0.5865	0.8091
$\hat{\sigma}(\hat{\lambda}_j)$	0.0042	0.0209	0.0184	0.0163

Table 4: Estimated BPs in Step 3.

Table 5 gives the most frequently selected orders and the corresponding sample means and standard errors of the parameters estimates in Step 4 for each stationary segment identified in Step 3. We see that the true orders are well identified and the estimated parameters are quite near to the true values given in table 1.

Estimate	Segment $X_{t,j}$				
$\hat{\beta}_j$	1	2	3	4	5
$(\hat{p}_j, \hat{q}_j)$	(1,2)	(0,0)	(1,0)	(1,1)	(0,1)
Frequency	62	98	87	93	88
$\hat{\mu}(\hat{d}_j)$	0.21	0.39	0.10	0.30	0.13
$\hat{\sigma}(\hat{d}_j)$	0.02	0.08	0.09	0.05	0.06
$\hat{\mu}(\hat{\phi}_{j,1})$	-0.73	-	-0.80	0.33	-
$\hat{\sigma}(\hat{\phi}_{j,1})$	0.06	-	0.02	0.07	-
$\hat{\mu}(\hat{\theta}_{j,1})$	0.58	-	-	-0.63	0.39
$\hat{\sigma}(\hat{\theta}_{j,1})$	0.08	-	-	0.07	0.09
$\hat{\mu}(\hat{\theta}_{j,2})$	-0.19	-	-	-	-
$\hat{\sigma}(\hat{\theta}_{j,2})$	0.02	-	-	-	-

Table 5: Selected orders and estimated parameters in Step 4.

5. APPLICATION TO TRAFFIC DATA

We use our procedure to obtain a model for the first OC48c Packet-over-SONET data set published by the NLANR MNA team. These 28000 data are the numbers of IP bytes collected at the Indianapolis router node on Sunday, April 6, 2003, per 30 millisecond time intervals during 14 minutes. Figure 3 plots the series and figure 4 shows that its auto-correlations decay slowly and are significant for large lags (more than 500), which is a strong evidence of LRD.

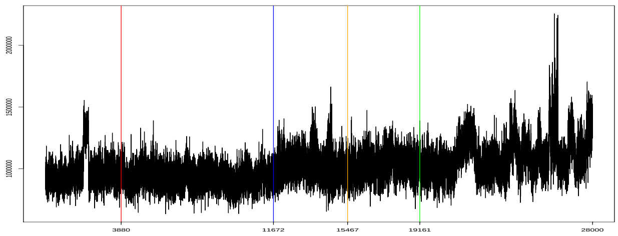

Figure 3: Internet traffic data. The vertical lines indicate the estimated BPs locations in Step 4.

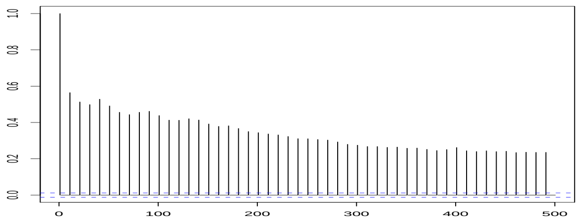

Figure 4: Auto-correlations of the internet traffic data.

We use $E = 2000$, which corresponds to an observation time of one minute for each elementary sub-series, and then $K = 14$. Using BIC in Step 1, all elementary sub-series are modelled by fractional noise models. Thus, we choose a piecewise FARIMA$(0, d_j, 0)$ model to represent the traffic data. After local estimation, we use different BPs numbers $m = 1, \ldots, 6$ to calculate the intervals with a BP (Step 2) and the results are summarized in table 6. When $m = 5$, we see that $\hat{k}_4 = 10$ is close to $\hat{k}_5 = 11$, and when $m = 6$, $\hat{k}_2 = 4$, $\hat{k}_3 = 5$ and $\hat{k}_4 = 6$ are close. Consequently, we retain in Step 2 a model with only four BPs, namely $\hat{k}_1 = 2$, $\hat{k}_2 = 6$, $\hat{k}_3 = 8$ and $\hat{k}_4 = 10$. The corresponding estimations of the BPs locations obtained in Step 3 are $\hat{\tau}_1 = 3880$, $\hat{\tau}_2 = 11672$, $\hat{\tau}_3 = 15467$ and $\hat{\tau}_4 = 19161$. These BPs are indicated by vertical lines in figure 3. Finding these BPs by a simple inspection of figure 3 is hard, and this, especially because here $p_j = q_j = 0$ and only the LRD parameter d_j changes with time.

Finally, figure 5 displays the LRD parameter estimate of each segment obtained in Step 4 and the local

m	$\hat{k}_j$					
1	15					
2	2	10				
3	2	6	10			
4	2	6	8	10		
5	2	6	8	10	11	
6	2	4	5	6	8	10

Table 6: BPs number selection.

estimate obtained in Step 1. The dot line fluctuates around the solid line which represents the structural changes of the LRD parameter.

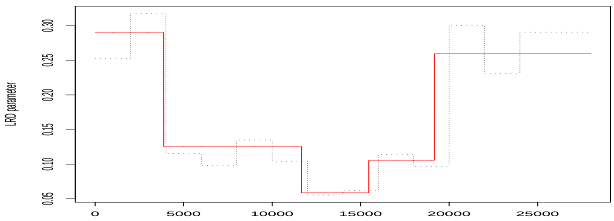

Figure 5: Estimated LRD parameter for the internet traffic data. Parameter estimate in Step 4 (solid line); local parameter estimate in Step 1 (dot lines).

6. CONCLUSION

In this paper, we have proposed a piecewise FARIMA model and the methodology for fitting it to a local stationary long-memory signal. This model is able to capture the structural break properties of the signal, it is flexible and allows to model simultaneously long and short range dependence. The model fitting consists in a four steps procedure designed to estimate both the BPs and the parameters. Simulations have shown good performances of the method. When applying our methodology to internet traffic data, a piecewise fractional noise model was selected. Future work includes applying piecewise FARIMA models in network design, management, traffic prediction and measurement-based network control.

REFERENCES

[1] J. Beran. On a class of M-estimators for Gaussian long-memory models. *Biometrika*, 81(4):755–766, 1994.

[2] H. Chang, M. Devetsikiotis, I. Lambadaris, and A. R. Kaye. Modeling and simulation of self-similar variable bit rate compressed video: a unified approach. *SIGCOMM Comput. Commun. Rev.*, 25(4):114–125, 1995.

[3] R. A. Davis, T. C. M. Lee, and G. A. Rodriguez-Yam. Structural break estimation for nonstationary time series models. *J. Amer. Statist. Assoc.*, 101(473):223–239, 2006.

[4] J. Eran, R. Sherman, M. S. Taqqu, and W. Willinger. Long-Range Dependence in Variable-Bit-Rate Video Traffic. *IEEE Trans. Commun*, 43:1566–1579, 1995.

[5] R. Fox and M. S. Taqqu. Large-sample properties of parameter estimates for strongly dependent stationary Gaussian time series. *Ann. Statist.*, 14(2):517–532, 1986.

[6] M. W. Garrett and W. Willinger. Analysis, Modeling and Generation of Self-Similar VBR Video Traffic. *ACM SIGCOMM Computer Communication*, 24(4):269–280, 1994.

[7] C. W. J. Granger and R. Joyeux. An introduction to long-memory time series models and fractional differencing. *J. Time Ser. Anal.*, 1(1):15–29, 1980.

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

[9] I. Norros. A storage model with self-similar input. *Queueing Systems Theory Appl.*, 16(3-4):387–396, 1994.

[10] V. Paxson and S. Floyd. Wide-area traffic: the failure of Poisson modeling. In *SIGCOMM '94: Proceedings of the conference on Communications architectures, protocols and applications*, pages 257–268, New York, NY, USA, 1994. ACM Press.

[11] B. K. Ray and R. S. Tsay. Bayesian methods for change-point detection in long-range dependent processes. *Journal of Time Series Analysis*, 23(6):687–705, 2002.

[12] L. Song, P. Bondon, Y. Cao, and Q. Cheng. A Time-Varying FARIMA Model for Internet Traffic. In *Image and Signal Processing, 2008. CISP'08*, volume 5, pages 83–87, 2008.

[13] S. Stoev and M. S. Taqqu. Asymptotic self-similarity and wavelet estimation for long-range dependent fractional autoregressive integrated moving average time series with stable innovations. *J. Time Ser. Anal.*, 26(2):211–249, 2005.

[14] S. Stoev, M. S. Taqqu, C. Park, G. Michailidis, and J. S. Marron. LASS: a tool for the local analysis of self-similarity. *Comput. Statist. Data Anal.*, 50(9):2447–2471, 2006.

[15] K. Torre, D. Delignières, and L. Lemoine. Detection of long-range dependence and estimation of fractal exponents through ARFIMA modelling. *British Journal of Mathematical and Statistical Psychology*, 60(1):85–106, 2007.