Parsimonious Markov Modeling of Processes with Long Range Dependence

Jean-Chrysostome Bolot, Matthias Grossglauser

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Parsimonious Markov Modeling of Processes with Long Range Dependence

Jean-Chrysostome Bolot   Matthias Grossglauser

N° 2835

Mars 1996

THÈME 1

Rapport de recherche

ISSN 0249-6399

Parsimonious Markov Modeling of Processes with Long Range Dependence

Jean-Chrysostome Bolot   Matthias Grossglauser

Thème 1 — Réseaux et systèmes
Projet Rodeo

Rapport de recherche n˚2835 — Mars 1996 — 16 pages

Abstract: Markov models have been widely used to model arrival processes at switches in packet- and cell-switched networks. However, recent experimental evidence suggests that such processes exhibit a long-range dependence (LRD) property which is not captured by these models. Fractal models are attractive because they capture the LRD property while providing parsimonious modeling of processes. Multi-state Markov models can capture the LRD property to some extent. However, they do not follow this principle since every state added to such a model also adds two adjustable parameters and thus increases the complexity of fitting experimental data to these parameters. In this paper, we show that a fractal model can be accurately approximated over a finite range of time scales by parsimonious multi-stage Markov models where the transition rates form a geometric progression along the stages, and each stage models a different time scale.

Key-words: Modeling, parsimonious modeling, long range dependence, Markov modeling

(Résumé : tsvp)

Ce travail a été effectué dans le cadre d'un contrat avec France Telecom/CNET.

Unité de recherche INRIA Sophia-Antipolis
2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex (France)
Téléphone : (33) 93 65 77 77 – Télécopie : (33) 93 65 77 65

Approximation Markovienne de Modèle Fractal pour Processus Avec Dépendance à Long Terme

Résumé : Les modèles markoviens ont été largement utilisés pour modéliser le trafic dans les réseaux. Cependant, des mesures récentes indiquent que les processus de trafic ont des propriétés de dépendance à long terme (DLT) qui ne sont pas être prises en compte par les modèles markoviens. En fait, des modèles markoviens à grand nombre d'états peuvent prendre en compte en partie ce type de dépendence. Malheureusement, à chaque état correspondent en général deux paramètres du modèle, et on obtient donc des modèles qui comportent un très grand nombre de paramètres libres. Se pose alors le problème de l'identification de tous ces paramètres en pratique. C'est ainsi que des modèles dits "fractals" ont été depuis développés. Ces modèles sont intéressants à plus d'un titre. Premièrement, ils prennent en compte de façon exacte la DLT. De plus, ils ne font intervenir qu'un très petit nombre de paramètres. Nous montrons dans ce papier qu'il est en fait possible de développer des modèles markoviens à très faible nombre de paramètres qui approximent le modèle fractal et prennent en compte la DLT sur une gamme d'échelles de temps limitée mais arbitraire.

Mots-clé : Modélisation, parsimonie, dépendence à long terme, modèle Markovien

1 Introduction

Many stochastic models have been proposed in the literature to describe individual sources of traffic as well as traffic generated by the superposition of multiple sources. The more widely used models include the Poisson process and its many extensions such as the Markov-modulated Poisson process, a wide variety of interrupted (i.e. on/off) constant rate processes, etc. Most of these models are based on a Markovian structure essentially for reasons of mathematical tractability. Nevertheless, they have been successfully used to provide insight into many dimensioning and other performance evaluation problems.

However, mouting experimental evidence suggests that the traffic generated by individual sources as well as the traffic resulting from the superposition of multiple sources exhibit a property of long range dependence (LRD). Examples of individual traffic sources with LRD include video coders [10, 2] and ISDN terminals [15]. Examples of network traffic (i.e. resulting from of the superposition of multiple sources) with LRD include Ethernet local area networks [12] and wide area packet networks such as the Internet [17, 11] or the CCS network [7].

It turns out that traffic with LRD from an indivdual source can be modeled using the familiar on/off model widely used in conventional analysis. In this model, the source is active and transmits data at a constant rate (often referred to as the peak rate) during the on period, and it is inactive during the off period. Recent measurements [20] suggest that real traffic sources can be modeled using such on/off models in which the sojourn time distributions in the on and/or off states are heavy-tailed. Heavy tail distributions are characterized by an asymptotic decrease slower than the exponential decreases. Specifically, the probability that a period last longer than some duration d tends to $d^{-\alpha}$. In contrast, the sojourn time distribution for the on and off periods in Markovian models decreases exponentially fast, i.e. the probability that a period last longer than some duration d tends to $\exp(-\alpha\, d)$ with $\alpha > 0$ as d tends to infinity.

On/off sources with heavy-tailed distributions do exhibit the LRD property. Furthermore, the superposition of such sources results in an aggregate traffic which is characterized by a property of self-similarity. This property has been observed in operational networks [12], and thus heavy-tailed on/off source models are attractive candidates for describing real sources of network traffic. Throughout the rest of the paper, we refer to such sources as LRD sources or fractal sources.

The results above motivate the careful investigation of LRD sources and they point to several research directions including in particular i) the characterization of LRD sources and ii) the impact of LRD on network and application performance. An important question in the characterization problem is how to differentiate Markovian and LRD sources. An important question in the performance problem is how much difference the source characteristics impact network and application performance [1, 3, 6, 8, 14, 16].

It was not clear at the outset that Markovian and LRD sources are fundamentally different, and that this difference has an impact on the network. Indeed, it is possible to approximate a heavy-tailed distribution function by a superposition of exponential functions. One might then think of a LRD source as a simple generalization of the Markovian source. This approach, which amounts to mimicking LRD with Markovian models, has been taken for example in [13]. It can be used to obtain accurate approximate performance results since a power law decay can be approximated arbitrarily closely by enough exponential decay functions. However, the resulting Markovian models are complex multi-state models. This presents two problems, namely that of identifying the parameters (states and state transition rates) from experimental data, and that of obtaining closed-form analytic expressions for performance measures. The first problem is the more important one because it is often difficult and time consuming in practice to collect the data required for parameter estimation. Thus, this problem is generally used to promote instead the notion of parsimonious modeling, which essentially is to model processes in as simple a way as possible. We do believe that parsimonious modeling is a very important principle. However, we also believe that a careful comparison of the LRD on/off model and its Markovian approximations can provide additional insights into the characteristics of both models.

The paper is organized as follows. In Section 2, we introduce a simple on/off model for LRD sources. In this model, the transition rates are time dependent. If the rates are chosen appropriately, the sojourn time distributions follow a power law decay. The model is useful because it provides a convenient way to compare fractal and multi-state Markov models. In Section 3, we introduce the notion of effective transition rate and show how it can be used to discriminate between fractal and Markovian models. In Section 4, we show that the fractal model can be accurately approximated over a finite range of time scales by multi-stage Markov models where the transition rates form a geometric progression along the stages, and each stage models a different time scale. Section 5 concludes the paper.

2 An on/off model with time-dependent transition rates

We consider a continuous time two-state model in which transition rates are time dependent. The states are labeled 0 and 1. We also refer to state 0 as the off state, and to state 1 as the on state. We denote by $\lambda_{01}(t)$ the transition rate from state 0 to state 1, and by $\lambda_{10}(t)$ the transition rate from state 1 to state 0. Refer to Figure 1. This model is in a way a continuous time version of the model described in [18]. It

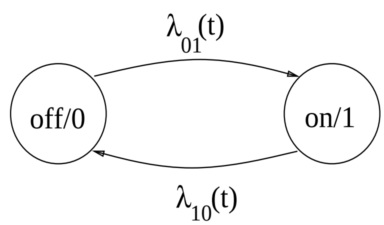

Figure 1: The on/off model with time-dependent transition rates

can be used to represent interarrival processes (one arrival corresponding to a state transition) or fluid flow sources (with a source being active during the on state, and idle during the off state).

We let $f_i(t)$ denote the sojourn time density and $P_i(t)$ denote the sojourn time distribution functions in state i ($i = 0, 1$). If the transition rates $\lambda_{ij}(t)$ from one state to another do not depend on time, then the distribution functions are the usual exponential distributions. For example, $P_i(t) = 1 - \exp(-\lambda t)$ if $\lambda_{ij}(t) = \lambda$ for all values of t. In this paper, we consider transition rates which decrease with time according to a power law. Specifically, we assume

$$\begin{aligned} \lambda_{01}(t) &= \alpha_0 t^{1-\beta_0} \qquad 1 < \beta_0 < 2 \\ \lambda_{10}(t) &= \alpha_1 t^{1-\beta_1} \qquad 1 < \beta_1 < 2 \end{aligned}$$

The singularity for $\lambda_{01}(t)$ and $\lambda_{10}(t)$ at $t = 0$ can be avoided by introducing “setup times” $t_0 > 0$ and $t_1 > 0$ and having for example $\lambda_{01}(t) = \alpha_0(t + t_0)^{1-\beta_0}$. This makes the analysis of the model slightly more tedious without providing additional insight. Therefore, we take $t_0 = t_1 = 0$ thoughout the rest of the paper[1].

Unlike in a Markovian model, the transition rate from the on to the off state decreases with time. Thus, the longer the source stays in the on state, the less likely it is to exist that state. We note that our model falls in the category of models with decreasing failure rate often used in reliability theory.

Let $R_i(t) = 1 - P_i(t)$. The sojourn time distributions corresponding to the above transition rates are given by

$$\lambda_{ij}(t) = -\frac{1}{R_i(t)} \frac{dR_i(t)}{dt}$$

Integrating both sides of the above equation yields

$$\log \frac{R_i(0)}{R_i(t)} = -\frac{\alpha_i}{2 - \beta_i} t^{2-\beta_i}$$

Therefore

$$R_i(t) = R_i(0) \exp\left(-\frac{\alpha_i}{2 - \beta_i} t^{2-\beta_i}\right)$$

Since $R_i(0) = 1$, we obtain

$$P_i(t) = 1 - \exp\left(-\frac{\alpha_i}{2 - \beta_i} t^{2-\beta_i}\right)$$

The corresponding density functions are given by $f_i(t) = dP_i(t)/dt$. After some algebraic manipulations, we find

$$\begin{aligned} f_i(t) &= dP_i(t)/dt \\ &= \alpha_i t^{1-\beta_i} \exp\left(-\frac{\alpha_i}{2 - \beta_i} t^{2-\beta_i}\right) \end{aligned}$$

[1]Mathematically inclined readers can consider the general case and derive the results in the paper by taking the limits $t_0 \to 0$ and $t_1 \to 0$. As an example, the density function for $f_i(t)$ derived later in the paper is found to be $f_i(t) = \alpha(t + t_i)^{1-\beta_i} \exp(\frac{\alpha_i}{2-\beta_i} t_i^{2-\beta_i}) \exp(-\frac{\alpha_i}{2-\beta_i}(t + t_i)^{2-\beta_i})$ in the general case. It is easy to verify that it tends toward $f_i(t) = \alpha_i t^{1-\beta_i} \exp(-\frac{\alpha_i}{2-\beta_i} t^{2-\beta_i})$ in the limit $t_i \to 0$.

Thus, we obtain the Weibull distribution which has the desired power law decay property. Given the above density functions, we can derive expressions for the moments of the sojourn time distribution. For example, let S_i denote the time spent in state i. Then we find the first moment

$$\begin{aligned} E[S_i] &= \int_0^\infty t f_i(t) dt \\ &= \frac{\Gamma[(3-\beta_i)/(2-\beta_i)]}{(\alpha_i/(2-\beta_i))^{1/(2-\beta_i)}} \end{aligned}$$

where $\Gamma[]$ denotes the gamma function.

Of course, we obtain the familiar results of the exponential on/off source if we set $\lambda_{ij}(t) = \lambda$. Then $f(t) = \alpha \exp(-\alpha t)$, $P(t) = \exp(-\alpha t)$, and $E[i] = 1/\alpha$.

3 Instantaneous transition rates and applications

For an on/off model with general time-dependent transition rates, we define the instantaneous transition rate at time t as follows

$$\mu_i(t) = -\frac{d \ln P_i(t)}{dt} \tag{1}$$

This parameter is widely used in reliability theory. There, it is referred to as the age-dependent failure rate, and it represents the probability that a component which has survived until time t fails in the interval $[t, t+dt]$.

For the fractal model, we find from equation 1 that $\mu_i(t) = \alpha_i t^{1-\beta_i}$. For the exponential Markov model, $\mu_i(t) = \lambda$ does not depend on t. Note that the inverse of the instantaneous transition rate is the natural time scale of the process. The Markov model has a natural time scale independent of time, namely $1/\lambda$. However, the fractal model does not have such a unique time scale.

Thus, it is easy to discriminate the fractal and Markov models by observing the variations of $\mu_i(t)$ as a function of t in a log scale plot[2]. Refer to Figure 2. For the

[2]Discriminating both models in practice might not be so easy. Indeed, it is notoriously difficult to reliably estimate the derivative of a continuous curve (here a sojourn time distribution) from measurements. This is because i) the distribution is approximated by a frequency histogram with non-zero bin width, and ii) the derivative estimation is sensitive to noise created by the binning process as well as to noise in the measurement data.

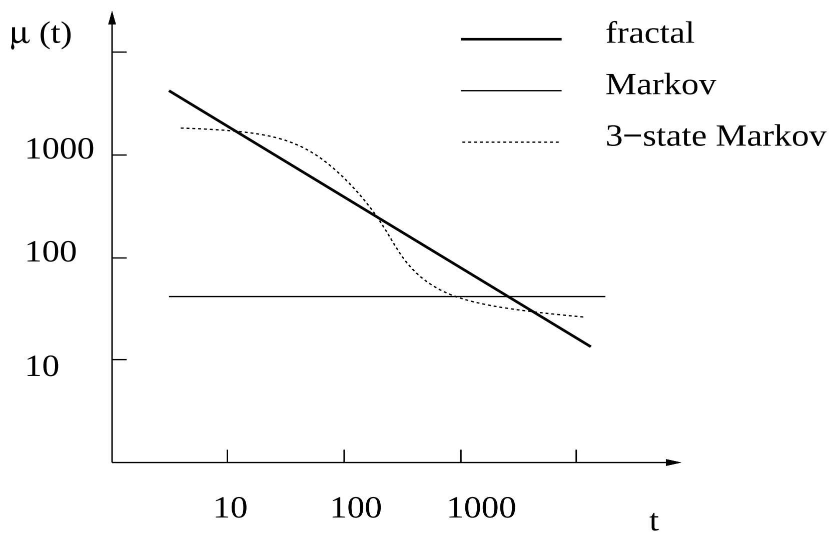

Figure 2: Variations of the instantaneous transition rates $\mu_i(t)$ versus t (log scale) for three models

two state Markov model, the instantaneous transition rate does not depend on time. Thus, it shows on the plot as a horizontal straight line. In contrast, the instantaneous transition rate for the fractal model shows as a straight line with negative slope $1-\beta_i$.

The figure suggests one approach to approximating the fractal model by Markovian models. In this approach, we approximate the line of the fractal model by a staircase function, i.e. by piecewise horizontal lines. By increasing the number of stairs and decreasing the width of individual stairs, it would be possible to approximate the line of the fractal model with arbitrary precision.

The simplest staircase approximation includes only two stairs. Since a Markov model with natural timescale τ is represented by a horizontal line with ordinate equal to $\log \tau$ in Figure 2, we would like to obtain a model which explicitely models two different time scales. A little thought shows that a 3-state Markov model yields precisely a 2 stair approximation. One way to obtain such a model is to add another state to the two-state model and to choose the transition rates so that they produce the appropriate time scale. Specifically, let us consider the model in Figure 3. The on state is now made up of two states, which we refer to as states 1 and 2. The transitions from the on to the off state occur over two timescales the values of which intuitively

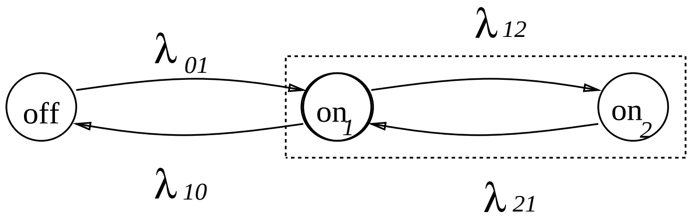

Figure 3: A three-state Markov model

depend on the timescale associated with state 1 (which in turn depends intuitively on the ratio $\lambda_{10}/\lambda_{12}$) and that associated with state 2 (which in turn depends on the value of λ_{21}). The intuition is supported by analytic results, since it is tedious but easy to show that the density function of the sojourn time in the new on state can be expressed as the sum of two exponential functions. Furthermore, the instantaneous transition rate between the on and the off state can be expressed as a rational function of exponentials. Specifically

$$\mu_1(t) = \frac{a\exp(-bt) + c\exp(-dt)}{(a/b)\exp(-bt) + (c/d)\exp(-dt)} \tag{2}$$

where

$$\begin{aligned}
a &= \lambda_{12}(\lambda_{01} - b)/(d - b) \\
c &= \lambda_{12} - a \\
b &= 1/2(\gamma + \sqrt{\gamma^2 - 4\nu}) \\
d &= 1/2(\gamma - \sqrt{\gamma^2 - 4\nu}) \\
\gamma &= \lambda_{12} + \lambda_{01} + \lambda_{10} \\
\nu &= \lambda_{12}\lambda_{01}
\end{aligned}$$

If we now plot $\mu_1(t)$ as a function of t on a log-log scale, we obtain the curve labeled *3-state Markov* in Figure 3. The curve includes two horizontal stretches corresponding to the two timescales mentioned above. The short time scale is obtained by taking the limit of equation 2 for $t \to 0$. We find that $\lim_{t\to 0}\mu_1(t) = a + c = \lambda_{12}$. The long time scale is obtained by taking the limit of equation 2 for $t \to \infty$. We find that $\lim_{t\to\infty}\mu_1(t) = \min(b, d)$.

This model would thus be suitable to capture the characteristics of a traffic source which generates packets in burts. The short time scale characterizes the frequency of transitions (i.e. the packet departure process) within a burst. The long time scale characterizes the duration (i.e. the idle period) between bursts.

In the 3-state model above, we replaced the on state by the two state system (1, 2) with transitions from state 1 to the off state. Another possibility is to replace the on state by another two state system (1, 2) with transitions from both states to the off state. We obtain the model in Figure 4 below. This model can be analyzed in essen-

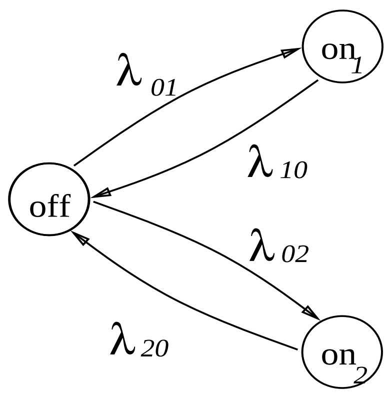

Figure 4: Another three-state Markov model

tially the same way as the model above. The transitions from the on to the off state occur over two timescales the values of which depend on the timescale associated with state 1 (which in turn depends on the value of λ_{10}) and that associated with state 2 (which in turn depends on the value of λ_{20}).

4 Multistate Markov approximations for the fractal model

The idea behind the 3-state models described in Section 3 can be extended to generate $n+1$-state models in which the on state actually include n states. By choosing suitable transition rates between these states, we can model a traffic source with n timescales. In the limit when $n \to \infty$, it appears possible to reach the limit of the fractal model. The downside of this approach, of course, is that the number of parameters required to describe the model increases with n. In this section, we analyze two such $n+1$ state models and formalize the notion of approximating the fractal model by the multi-state Markov models. We also show how the parameter explosion problem can be avoided in part using an intuitively and physically pleasing approach.

We consider three models, namely a fractal model and two $(n+1)$ state Markov. For convenience of exposition, we consider a special case of the fractal model of Section 2 with exponentially distributed idle periods and heavy tailed activity periods. This is done by setting $\beta_0 = 2$, $\alpha_1 = 1$, and $\beta_1 = 1$. We now have to choose a structure for the Markov models. Recall from Section 3 that our approach to approximating the fractal model by Markovian models is to approximate the sloped line in Figure 2 by a staircase function. We found that each stair corresponds to a different timescale of the model. Thus, we would like to model transitions occurring over n different time scales using a $n+1$ state model with appropriate transition rates. Desirable properties for the transition rates are i) they should be relatively easy to determine from experimental data and ii) they should cover as wide a range as possible. In our models, we choose these rates so that they form a geometric progression with ratio τ. Specifically, we consider two $(n+1)$ state Markov models which we refer to as the series and the parallel models. They are shown in Figures 5 and 6, respectively.

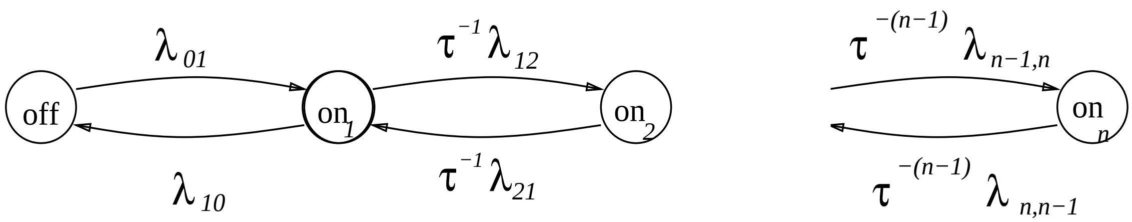

Figure 5: The series model

The parameter τ can be thought of as the largest timescale of the model. Thus, the slowest transitions in the model occur with frequency $\epsilon = 1/\tau$. Without loss of generality, we assume $\epsilon \ll 1$. Furthermore, we assume that the forward and backward transition rates between two states are equal, i.e. $\lambda_{ij} = \lambda_{ji}$ for all values of i and j.

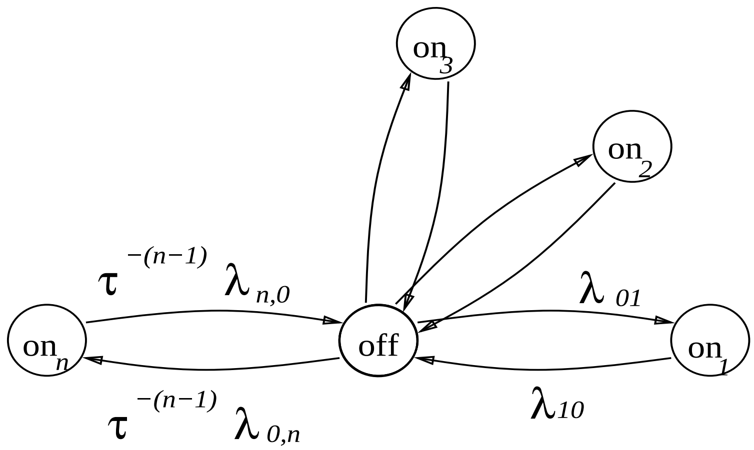

Figure 6: The parallel model

We next analyze and compare the models. We characterize each model by the asymptotic behavior of the sojourn time density functions $f_0(t)$ and $f_1(t)$.

In the *fractal model* considered in this section, the sojourn time in the off state is exponentially distributed and hence $f_0(t) = \alpha_0 \exp(-\alpha_0 t)$. The asymptotic behavior of $f_0(t)$ is thus $f_0(t) \sim \exp(-\alpha_0 t)$. The sojourn time density in the on state is asymptotically $f_1(t) \sim 1/t^2$.

In the *series model*, it is clear that the sojourn time in the off state is exponentially distributed with time constant $1/\lambda_{01}$. Since $\lambda_{01} \sim \alpha_0$, we have $f_0(t) \sim \exp(-\alpha_0 t)$.

The derivation of the sojourn time distribution in the on state is very complicated in the general case. However, it is relatively straighforward to obtain an approximation of it by taking advantage of properties of the geometric progression of the transition rates along the stages of the model. Specifically, each state in the series of states representing the on state can be considered in isolation because the transitions into and out of state i occur at a frequency of the order of τ^{i-1} which is much faster than that associated with state $i-1$ and much slower than that associatied with state $i+1$. Thus, the sojourn time density function in the on state can be approximated by the sum of the densities of each state $i > 0$, weighted by the frequency of visits to state i. Note that this approximation technique has often been used in queueing theory to decompose complex state spaces into more manageable ones [4]. We then derive the average duration between successive visits in state i using a local balance equation. The inverse of this quantity is the desired visit frequency f_i. We find

$$f_i \sim \epsilon^{i-1} \lambda_{i,i-1} / \lambda_{10}$$

We can now derive the sojourn time density function for the on state. We have

$$f_1(t) = \sum_i f_i \epsilon^{i-1} \lambda_{i,i-1} \exp(-\lambda_{i,i-1} t)$$

Replacing f_i in this sum by its expression above and doing a Taylor expansion in terms of ϵ^i, we obtain after some algebraic manipulations that so $f_1(t) \sim 1/t^2$, which is precisely what we also obtained for the fractal model.

In the *parallel model*, the sojourn time in the off state is exponentially distributed with time constant $1/\lambda$ where $\lambda = \sum_i \epsilon^{i-1} \lambda_{0i}$. Since we assume that $\epsilon \ll 1$, we obtain $\lambda \sim \lambda_{01}$ and hence $f_0(t) \sim \exp(-\alpha_0 t)$.

The on time distribution is the sum of the distributions of each on state weighted by the frequency of visits f_i from state 0 to state i. The value of f_i is given by

$$f_i = \epsilon^{i-1} \lambda_{0i} / \sum_i \lambda_{0i} \epsilon^{i-1}$$

and hence $f_i \sim \epsilon^{i-1} \lambda_{0i} / \lambda_{01}$. Observe that this value for f_i is similar to that derived in the analysis of the series model. Thus, we expect to obtain a similar asymptotic expression for $f_1(t)$. Indeed, using the above expression for f_i we have

$$f_1(t) = \sum_i \frac{\epsilon^{2i} \lambda_{i0} \lambda_{0i}}{\lambda_{01}} \exp(-\epsilon^{i-1} \lambda_{i0} t)$$

We then observe, using a reasoning similar to that above for the series model, that the sum above can be reduced to just one term, which is that with index i where $\epsilon^{i-1} \lambda_{i0} t \sim 1$. Thus we obtain

$$\begin{aligned} f_1(t) &\sim \frac{\epsilon^{2i-2} \lambda_{i0} \lambda_{0i}}{\lambda_{01}} \exp(-\epsilon^{i-1} \lambda_{i0} t) \\ &\sim 1/t^2 \end{aligned}$$

In conclusion, we have found that the distribution densities for the sojourn times in the on and the off states are asymptotically identical in the fractal and the Markov models. We use this as an indication that the Markov models are good approximations of the fractal model. We are currently examining the issue of the convergence (is a sense that remains to be defined) of the Markov models to the fractal model as $n \to \infty$.

5 Conclusion

We have considered the problem of comparing a fractal on/off model with by multi-state Markov approximations of it. We have found that Markov models with appropriately chosen rates may indeed be good approximations in the sense that their sojourn time distributions for both the on and the off states are asymptotically identical to those of the fractal model.

Fractal models have been advocated because they follow the principle of parsimonious modeling. Multi-state Markov models do not follow this principle since every state added to such a model also adds two adjustable parameters. This is a problem in particular when trying to fit experimental data with the parameters. However, we have shown that it may be possible to reduce the impact of this parameter fitting problem with multi-stage Markov models where the transition rates form a geometric progression along the stages, and each stage models a different time scale. We believe that our approach based on the analysis of the timescales at which transitions occur provides interesting insight into the differences between the two types of models. Nevertheless, it is not clear at this point how much it simplifies the fitting problem in practice for Markov models. This makes fractal models very attractive to model real traffic related processes, and it provides added motivation to develop queueing analysis techniques that can handle such processes.

Acknowledgements

This work is supported in part by a grant from France Telecom/CNET.

References

[1] R. Addie, M. Zukerman, T. Neame, “Fractal traffic: measurements, modeling, and performance evaluation”, *Proc. Infocom ’95*, pp. 977-984, Boston, MA, April 1995.

[2] J. Beran, R. Sherman, W. Willinger, “Long range dependence in variable bit rate video traffic” *IEEE Trans. Communications,* vol. 43, no. 3, pp. 1566-1579, Feb. 1995.

[3] F. Brichet, J. Roberts, A. Simonian, D. Veitch, “Heavy traffic analysis of a storage model with long-range dependent on/off sources”, Technical Report, CNET, Paris, 1995.

[4] J. Courtois, *Decomposability – Queueing and Computer System Approximations*, Academic Press, 1977.

[5] D. R. Cox, "Long-range dependence: A review", in *Statistical Appraisal*, H. A. David and H. E. David (eds.), Iowas State University Press, 1984.

[6] N. Duffiled, J. T. Lewis, N. O'Connell, "Predicting quality of service for traffic with long-range fluctuations", *Proc. ICC '95*, Seattle, WA, June 1995.

[7] D. Duffy, A. McIntosh, M. Rosenstein, D. Wilson, "Statistical analysis of CCSN/SS7 traffic data from working CCS subnetworks", *IEEE JSAC*, 1994.

[8] A. Erramilli, W. Willinger, O. Narayan, "Experimental queueing analysis with long range dependent packet traffic", to appear in *IEEE/ACM Trans. Networking*.

[9] A. Erramilli, R. P. Singh, P. Pruthi, "An application of deterministic chaotic maps to model packet traffic", to appear in *QUESTA*, 1995.

[10] M. W. Garrett, W. Willinger, "Analysis, modeling and generation of self-similar VBR video traffic", *Proc. ACM Sigcomm '94*, London, UK, pp. 269-280, Aug. 1994.

[11] S. Klivansky, A. Mukherjee, C. Song, "On long-range dependence in NSFNET traffic", Georgia Institute of Technology technical report, 1995.

[12] W. Leland, M. Taqqu, W. Willinger, D. Wilson, "On the self-similar nature of Ethernet traffic (Extended version)", *IEEE/ACM Trans. Networking*, vol. 2, no. 1, pp. 1-15, Feb. 1994.

[13] S. Q. Li, C. L. Hwang, "Queue response to input correlation functions: continuous spectral analysis", *IEEE/ACM Trans. Networking*, vol. 1, no. 3, pp. 678-692, 1993.

[14] N. Likhanov, B. Tsybakov, N. Georganas, "Analysis of an ATM buffer with self-similar ("fractal") traffic", *Proc. Infocom '95*, pp. 985-992, Boston, MA, April 1995.

[15] K. Meier-Hellstern, P. Wirth, Y-L. Yang, D. Hoeflin, "Traffic models for ISDN data users: Office automation application", *Proc. ITC 13*, Copenhagen, pp. 167-172, June 1991.

[16] I. Norros, "A buffer with self similar input", *QUESTA*, 1994

[17] V. Paxson, S. Floyd, "Wide area traffic: The failure of Poisson modeling", *Proc. ACM Sigcomm '94*, London, UK, pp. 257-268, Sept. 1994.

[18] P. Pruthi, A. Erramilli, “Heavy-tailed on/off source behavior and self-similar traffic”, to appear in *ICC ’95*.

[19] D. Veitch, “Novel models of broadband traffic”, *Proc. Globecom ’93*, Houston, TX, pp. 1057-1061, Dec. 1993

[20] W. Willinger, M. Taqqu, R. Sherman, D. Wilson, “Self-similarity through high variability: Statistical analysis of Ethernet LAN traffic at the source level”, *Proc. ACM Sigcomm ’95*, Cambridge, MA, Sept. 1995.

Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irisa, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
ISSN 0249-6399