Heavy-Tailed ON/OFF Source Behavior and Self-Similar Traffic

Parag Pruthi
Royal Institute of Technology, Stockholm, Sweden
Ashok Erramilli
Bellcore, New Jersey, USA

ABSTRACT: Recent traffic measurement studies suggest that the self-similarity observed in packet traffic arises from aggregating individual sources which behave in an ON/OFF manner with heavy-tailed sojourn times in one or both of the states. In this paper, we investigate the connection between general ON/OFF behavior, self-similarity and queueing performance. We use chaotic maps to model general ON/OFF behavior with combinations of heavy tailed and light tailed sojourn time behavior. We present results which show that chaotic maps which capture the heavy-tailed sojourn time behavior in the OFF and/or ON states generate traffic that is asymptotically self-similar. However, the resulting queue length distribution decays as a power law with the heavy ON source, and as an exponential with the light ON source, even though both processes exhibit identical *1/f* noise behavior. To resolve this apparent paradox, we consider aggregates of ON and OFF sources, and show that the nature of the ON period is less consequential, and in both instances the aggregate appears to converge to Fractional Brownian Motion (FBM). The queueing behavior is heavy in both cases, corresponding to the "stretched exponential" form predicted by FBM models. This indicates that differences in the single source case arise due to the impacts of higher order statistics, which become less significant as sources are aggregated. Convergence to FBM is observed to be slower with light ON sources. Our results indicate that in assessing the impact of long range dependence on performance, the potential impacts of other factors, such as higher-order statistics, must also be considered. Further, our analysis indicates conditions under which long-range dependence can dominate queueing performance in fast packet and SS7 networks, and with Variable Bit Rate (VBR) video applications.

KEYWORDS: Packet Traffic Modeling, Deterministic Chaotic Maps, Self-Similar, Long Range Dependence, Fractals, Chaos, Intermittency, Performance Analysis.

1 Introduction

Conventional teletraffic theory is largely based on Markovian assumptions of the traffic arrival process and of service time distributions. Although the theory permits other models for the arrival and holding time distributions, most current research is based on Markovian assumptions. Thus traditional models of packet arrival process are characterized by interarrival times that decay exponentially, by variances that decay inversely with sample size, by a power spectrum which is convergent near the origin, etc. and lead to traffic models which are short-range dependent. These models have been used with tremendous success in the design and operation of telephone networks.

In contrast, traffic arrival processes in packet based networks are much more bursty and intermittent. A number of recent measurement studies from the full range of packet based networks and services (ISDN packet, Ethernet, SS7, VBR Video) [2][4][9][11][12][13] indicate that packet traffic is characterized by interarrival times that decay with heavy tails, by variances that decay as a fractional power of the sample size, by a power spectrum that is divergent near the origin, and by correlations that are long range dependent. A number of analytical and experimental studies have established the performance significance of these features [4][8][14].

More recent measurement work has focused on the physical basis of the self-similarity observed in the full range of packet based networks. Based on a preliminary analysis of individual sources on an Ethernet, Willinger [17] observes that individual sources can be represented by the familiar ON/OFF abstraction: the source is either transmitting at a peak rate when it is in the ON state, or it is completely idle when it is in the OFF state. In conventional traffic models, the sojourn times in the two states are characterized by distributions with light-tailed behavior e.g., those that decay exponentially. In contrast, Willinger [17] observes that in actual traffic these sojourn time distributions decay far more slowly, as power laws, such that the variance of the sojourn times is infinite. Analogous conclusions are made by Meier-Hellstern *et al.* [13] in studies of individual ISDN data traffic sources. These statements are consistent with the observation that in self-similar traffic, bursts occur over all durations, and there is no characteristic length or time scale for traffic bursts. Based on theoretical results [16] which suggest that aggregating a large number of ON/OFF sources with the same heavy-tailed distribution in the two states results in a self-similar process, Willinger makes the conjecture that heavy-tailed ON/OFF behavior provides the physical basis for the self-similarity observed in packet data traffic.

These measurement results motivate the investigation of more general ON/OFF source behavior, in which the distribution of sojourn times in one of the two states is heavy-tailed, in the sense that its variance is infinite. In principle, one can approximate the heavy-tailed sojourn time behavior by a mixture of exponentials, but the number of parameters required to match observed data increases as the sample size (and the range of the observed sojourn times) increases. In the spirit of parsimoniously modeling the many time scales inherent in self similar traffic, we instead use chaotic maps to model such ON/OFF behavior. We are motivated in part by experiences in a number of other disciplines in which chaotic maps have been used as efficient generators of fractal processes. Our approach is based on earlier work on the application of deterministic chaotic maps to model traffic flows [5][6][7][8]. Specifically, we consider the simplest class of chaotic systems, known as one dimensional (1-D) chaotic maps, in which the evolution of a state variable x over discrete time n is described by a *deterministic* nonlinear transformation $x_{n+1} = f(x_n)$. We can model packet traffic sources using such maps by setting up a correspondence between the state variable x_n and activity of the source. For example, we model ON/OFF source behavior by stipulating that the source is generating traffic at a peak rate if x_n exceeds a threshold, and is idle otherwise. By suitable choice of the function $f(\cdot)$ we are able to model and analyze a wide range of ON/OFF behavior. In this setting, we can also cast simple queueing systems in terms of two dimensional deterministic transformations.

In section 2 we describe the basic model and introduce some of the fundamental concepts. We then look in detail at two different maps, the *Single Intermittency* map, and the *Double Intermittency*

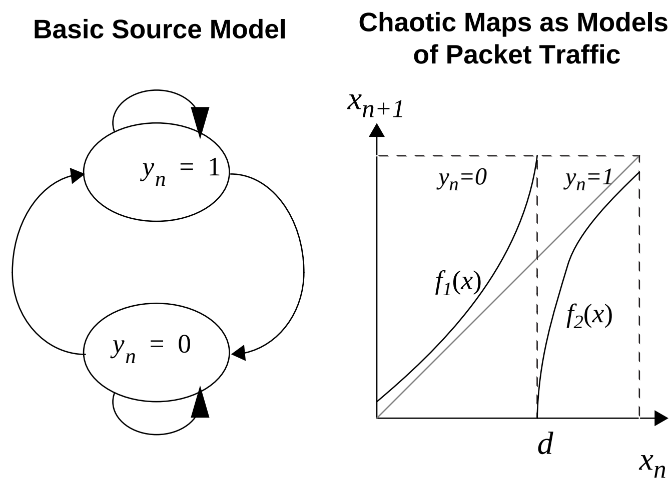

Figure 1 Simple nonlinear map as packet generator. When $x_n \geq d$ the source is taken to be active, generating packet(s) with every iteration (y_n=1)

map as models of general ON/OFF behavior in sections 3 and 4 respectively. We then consider the performance implications of queues driven by such traffic processes and show the importance of considering the physical basis of self similarity in section 5. Section 6 concludes this paper and discusses some open issues.

2 Chaotic Maps as Traffic Source Models

In this section we review the chaotic map formulation to model traffic sources. Consider a one-dimensional map in which the state variable x_n evolves over time according to the nonlinear map [6][7][8]:

$$x_{n+1} = \begin{cases} f_1(x_n), & 0 \leq x_n < d \\ f_2(x_n), & d \leq x_n \leq 1 \end{cases} \qquad \text{(EQ 1)}$$

Note that the formulation is completely deterministic, and a given initial condition fully defines a trajectory in the phase space. This is analogous to a "realization" of a stochastic process. We can now model a packet generation process by assuming that the source is in a passive or active state at time n depending on whether x_n is below or above a threshold (Figure 1). Every iteration of the map in the **active** state is taken to generate a batch of $k \geq 1$ packets. The packet arrival process is then described by the evolution of an associate indicator variable y_n:

$$y_n = \begin{cases} 0, & (0 \leq x_n < d) \\ 1, & (d \leq x_n \leq 1) \end{cases} \qquad \text{(EQ 2)}$$

Other source interpretations are also possible [7][8].

In practice, the y_n are observed, while the x_n are hidden, and the challenge is to find suitable $f_1(\cdot)$ and $f_2(\cdot)$ such that y_n match those properties of actual packet traffic that are relevant for queueing performance. The following results from [8] relate the choice of $f_1(\cdot)$ and $f_2(\cdot)$ to the sojourn time behavior in the ON and OFF states. If $f_1(0)=0$ and $f_2(1)=1$,

- linear segments will generate geometric sojourn times in the ON & OFF states i.e., $P(I > i) \propto a^i$
- nonlinear segments that have a Taylor series expansion of the form $f(x) \approx \varepsilon + x + cx^m$ in the vicinity of 0 or 1 will generate heavy-tailed sojourn times i.e.,

$$P(I > i) \propto i^{-\frac{1}{m-1}}, \; i \to \infty \qquad \text{(EQ 3)}$$

The deterministic formulation of a chaotic map also permits performance analysis. Queueing systems can also be modeled as a two dimensional deterministic transformation [6][7][8]. As an illustration, consider a server with a deterministic service time equal to the iteration interval of the underlying chaotic map (*e.g.* cell processing in ATM networks). The performance of this server can be analyzed in the framework of a discrete queueing system whose state at time n is given by the queue length l_n. If every iteration of the map in the active state produces a batch of k packets[1], the server can be completely described by the following coupled system of nonlinear equations:

$$l_{n+1} = max(l_n - 1, 0) + ky_n(x_n) \qquad \text{(EQ 4)}$$

$$x_{n+1} = f(x_n) \qquad \text{(EQ 5)}$$

where the $max(\cdot)$ function accounts for the boundary conditions at $l_n = 0$. Even though (EQ 4) and (EQ 5) are fully deterministic, the chaotic map formulation nevertheless captures randomness in the arrival process. The state of the source and queue is represented by the point (x_n, l_n), and the evolution of the queue from any initial condition, which is fully governed by (EQ 4) and (EQ 5), can be represented by motion in the plane with $x_n \in (0, 1)$ and integer l_n. The marginal distribution of iterates along the l axis corresponds to the queue length distribution and can be numerically computed in principle.

In particular we are interested in modeling traffic sources with a range of ON/OFF behavior, and relating it to self-similar traffic. In the next section we show that a simple chaotic map can be a generator of *1/f noise*, which is the frequency domain manifestation of long range dependence.

3 Single Intermittency Map

Intermittency is a phenomenon which has been used to study systems (especially turbulence) which are characterized by alternating periods of long "regular" phases and relatively short irregular "bursts". We use this map to model sources which are characterized by heavy tailed OFF (corresponding to long "regular" phases), and light-tailed ON distributions (corresponding to the short irregular "bursts"). Using the results stated in the previous section, we construct the following map [6]:

1. For any backlogs to exist k must be greater than 1.

$$x_{n+1} = \begin{cases} \varepsilon + x_n + c x_n^m & 0 \le x_n < d \\ \dfrac{x_n - d}{1-d} & d \le x_n \le 1 \end{cases} \quad \text{(EQ 6)}$$

where

$$c = \frac{1-\varepsilon-d}{d^m} \quad \text{and } \varepsilon \ll d \quad \text{(EQ 7)}$$

The inactive state $0 \le x_n < d$ is represented by a nonlinear segment. If the initial state x_o at the beginning of any passive period is not near the origin, the resulting passive period is relatively short. However with finite probability the initial state can be close to the origin (*i.e.* $x_n \approx O(\varepsilon) \gg x_n^m$), and here the state variable evolves very slowly, and it takes many iterations of the map to emerge from this region. This generates the very long idle periods. The active period $d \le x_n \le 1$, on the other hand, is represented by a linear segment. Thus the sojourn times in the active state are essentially geometric [8] and sources of this type generate limited bursts of packets interspersed with (occasionally) long idle intervals. A source of messages in existing SS7 networks can be loosely described in such terms: a burst of messages at the initiation of a call, followed by an idle period corresponding to the holding time of the call (which can span many time scales), followed by another burst corresponding to the termination of the call. It is a reasonable conjecture that the sustained inactivity modeled by the single intermittency map is at the root of the self-similarity observed in traffic from existing SS7 networks [9]. Note that the limited burst scenario may not be valid when signaling data from new services such as PCS, mobile computing, etc. make up a substantial portion of signaling traffic.

The sojourn times in the inactive state are heavy tailed with infinite variance when ($3/2<m<2$). The effect of ε is to limit the maximum duration of the idle period. We show in [15] that the traffic process generated by the *Single Intermittency* map displays *1/f*-noise, and is asymptotically second-order self-similar. The power spectrum $S(\omega)$ of the generated traffic process is asymptotically found to be,

$$S(\omega) \propto \omega^{-\frac{2m-3}{m-1}} \quad \text{for } 3/2<m<2. \quad \text{(EQ 8)}$$

Equivalently, the output process has a correlation structure that decays as $\rho_h \propto h^{-\beta}$ asymptotically with $\beta = (2-m)/(m-1)$. It is relatively straight-forward to show that processes with such a power law correlation structure are asymptotically second-order self-similar [3], with the Hurst parameter H, given by, $H = (3m-4)/(2m-2)$ in the range $(1/2,1)$, when m is in the range [$3/2 < m < 2$]. The effect of ε is to place an upper-cutoff on the correlations; in this paper we assume that there are no upper cut-offs on the power law correlation behavior, and set ε=0.

4 Double Intermittency Map

The sojourn time in the ON period can be made heavy-tailed as well by replacing the linear segment in the single intermittency map with an appropriate nonlinear segment. Consider the map [7][8],

$$x_{n+1} = \begin{cases} \varepsilon_1 + x_n + c_1 x_n^m & 0 \le x_n < d \\ -\varepsilon_2 + x_n - c_2 (1-x_n)^m & d \le x_n \le 1 \end{cases} \quad \text{(EQ 9)}$$

$$c_1 = \frac{1-\varepsilon_1-d}{d^m} \qquad c_2 = -\frac{\varepsilon_2 - d}{(1-d)^m} \quad \text{(EQ 10)}$$

As discussed in Section 2, the sojourn time distributions of the active and inactive periods behave as [8] $P(I>i) \propto i^{-\frac{1}{m-1}}$, $(i \to \infty)$. To this extent, this map is a representation of the ON/OFF behavior noted by Willinger [17] in preliminary studies of single Ethernet sources. Taqqu and Levy [16] derive results which indicate that aggregating a large number of heavy tailed ON-OFF sources will in the limit lead to Fractional Brownian Motion (FBM), and the Hurst index H is given by $H = (3m-4)/(2m-2)$, [$3/2 < m < 2$]. Considerable numerical evidence also shows that the output is long-range dependent with (as before) $H = (3m-4)/(2m-2)$ [$3/2 < m < 2$]. This also provides a connection between our approach, and the Fractional Brownian Motion (FBM) model proposed by Norros [14] to model self-similar packet traffic. Thus, both sustained inactivity (exemplified by a heavy tailed OFF, light tailed ON source) and sustained activity (heavy tailed ON/OFF sources) gives rise to traffic that is asymptotically self-similar in the sense that it exhibits *1/f*-noise. Next we consider the performance implications of the two self-similar traffic processes discussed above.

5 Performance Analysis

5.1 Single Intermittency Map

For the queueing system of (EQ 4) and (EQ 5) driven by the single intermittency map it can be shown that the tail of the queue length distributions decay geometrically as a function of the generalized occupancy, *i.e.* $P(L>l) \propto \gamma^{l+1}$ [7][8][1]. This result can be established by simulating the queueing system, or by numerically solving for the queue-length distribution from the dynamical system description. The exponential decay characteristic of this form of the queue length distribution is a consequence of limited burst lengths even though the input traffic process exhibits long-range dependence. Figure 2 illustrates this via simulation of (EQ 4) and (EQ 5) for various values of *m*. As can be seen, the complementary distribution function *P(L>l)* is clearly linear on a semi-log plot, indicating an exponential decay. In particular, changing *m*, which changes the target Hurst parameter value, has *no* effect on the exponential nature of the decay, which must arise from the limited burst durations.

5.2 Double Intermittency Map

In contrast, analysis of the double intermittency map shows that the effect of heavy-tailed sojourn times for the active state leads

1. This result is reminiscent of GI/M/. theory, and by analogy, we can refer to γ as the *generalized occupancy.*

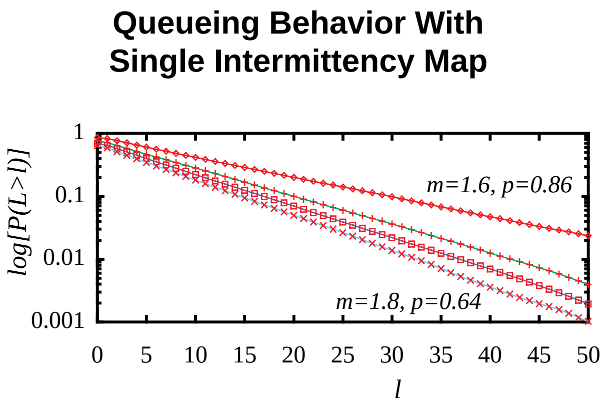

Figure 2 The queue length distribution has a geometric decay $(P(L>l) \propto \gamma^{l+1})$ for a deterministic server driven by the traffic process generated from the *Single Intermittency* map. Shown here for various values of the Hurst parameter and utilizations (p) (both ranging from about 2/3 to 7/8). *Note the semi-log plot.*

to a power law decay in the tail of the queue length distribution function [8]:

$$P(L>l) \propto l^{-\left(\frac{2-m}{m-1}\right)}, \quad l \to \infty . \qquad \text{(EQ 11)}$$

This power-law behavior is apparent in Figure 3 which shows the plot of the queue length distribution on a log-log plot. Note that

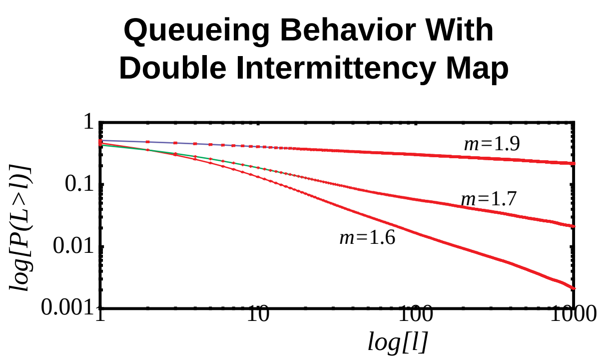

Figure 3 The queue length distribution is heavy-tailed, $\left(P(L>l) \propto l^{-\frac{2-m}{m-1}}\right)$; a result due to the power-law burst size distribution. *Note the log-log plot*.

the decay is linear in the log-log plot. Such heavy-tailed queueing behavior has been studied in M/G/1 models with heavy-tailed service distributions (see [1] for an analogous result with Pareto service times) but is unprecedented in conventional arrival processes with constant service times. The queue length distribution is so heavy that in the range *(3/2<m<2)*, the average queue length is unbounded. From a physical basis, this result is not surprising, given that the burst lengths have an infinite variance. This suggests that for single sources with heavy-tailed ON behavior, peak rate allocation must be used to avoid severe performance degradation. Note that the form of the queue length distribution here is heavier than the Weibullian form predicted by the FBM model of [14]. Aggregates of ON/OFF sources are discussed next.

5.3 Multiple Source Aggregation

The enormous contrast in the queueing behavior with two processes that asymptotically have the same second order behavior is striking. These differences can be attributed to the impacts of other factors, such as marginal distributions, and higher order statistics. It is shown in [10] that a form of the central limit theorem applies to aggregates of long range dependent traces, and aggregating a large number of sources results in marginals that are Gaussian, and in which second-order statistics dominate. Thus, by considering aggregates of ON/OFF sources, one would expect the effect of other factors, such as differences in the marginals, and higher order statistics, to be reduced.

Figure 4 shows the time series obtained by aggregating 50 sources

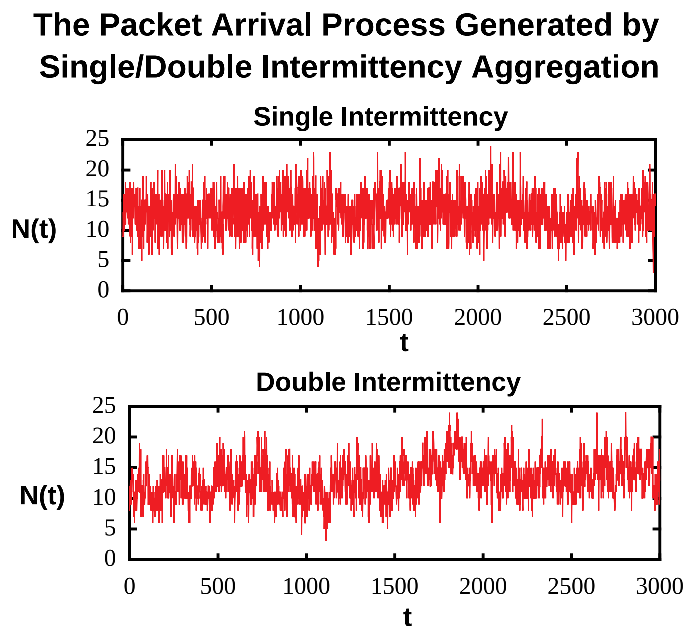

Figure 4 The packet arrival process generated by aggregating 50 sources of the Single Intermittency type (top) and Double Intermittency type (bottom). Shown here is the number of packets generated *N(t)* at discrete time *t*. Note the low frequency variation in these processes indicating *1/f*-noise.

of the *Single and Double Intermittency* type. Note the presence of low frequency variations that indicate the presence of $1/f$ -noise or long-range dependence discussed earlier. Also note the broader band in the plot of the aggregates of the single intermittency map as compared to the aggregates of the double intermittency map indicating the presence of more relative power in the high frequency range (which is to be expected, because limited bursts will correspond to high frequency variations).

Figure 5 shows the variance-time plot [11] obtained from a multiplexing of 50 sources of the *Double Intermittency* type. The slowly decaying variances are an evidence of long-range dependence, and the estimate of H from this plot (0.87) is very close to the target value (=0.875) (the more rigorous Whittle [11] estimate also gives similar results). As Figure 6 shows, the queue length distributions appear to be Weibullian, or “stretched exponential”. This is consistent with several theoretical results: *(i)* Taqqu and Levy [16] indicate that aggregating a large number of heavy tailed active/inactive sources will in the limit lead to Fractional Brownian Motion (FBM) *(ii)* Norros [14] shows that FBM arrival pro-

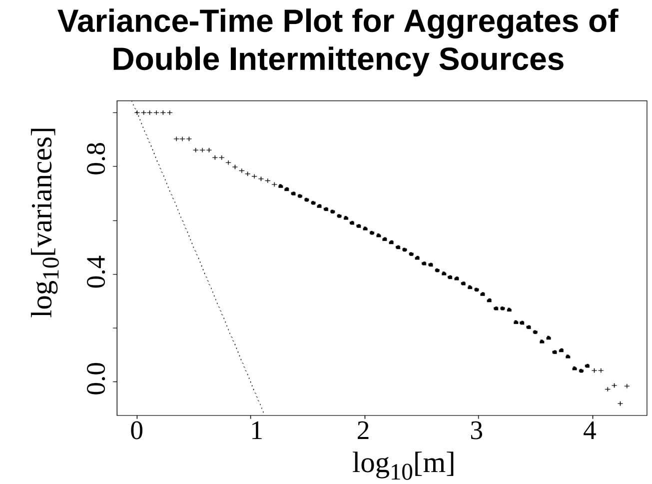

Figure 5 The variance time plot of trace obtained by aggregating 50 sources of the double intermittency type: in absence of correlations, variances should decay inversely with sample size (dotted line); decay slower than this rate indicates long-range dependence in trace

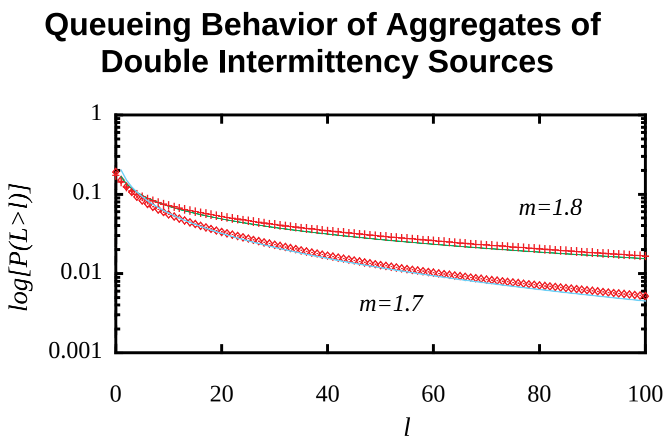

Figure 6 The queue length distribution has a Stretched Exponential decay $\left(P(L > l) \propto e^{-\gamma x^{\beta}}\right)$ (with $\beta \leq 1$) for a deterministic server driven by the traffic process generated from the aggregation of 50 *Double Intermittency* type sources. The fit is given by the solid lines which fall almost on the curves.

cesses generate "stretched exponential" heavy-tailed queueing. This suggests that aggregating as few as 50 sources leads to a good approximation of FBM, and aggregating the output of the double intermittency map may be an efficient way of generating FBM. For the purposes of traffic analysis, this technique has the advantage that the number of arrivals in a time interval is non-negative, in contrast to other techniques of generating FBM. However, the quality of the generated traces must be investigated further, along with computational issues.

Figure 7 shows the variance-time plot of the time series obtained by aggregating 50 sources of the single intermittency type. As before, the slowly decaying variances indicate the presence of long-range dependence, and the estimate of the Hurst parameter from this plot (0.85) is close to the target value (for m=1.8, H =0.875) (the Whittle estimate also gives similar results).

Simulation studies of the queueing behavior obtained with aggre-

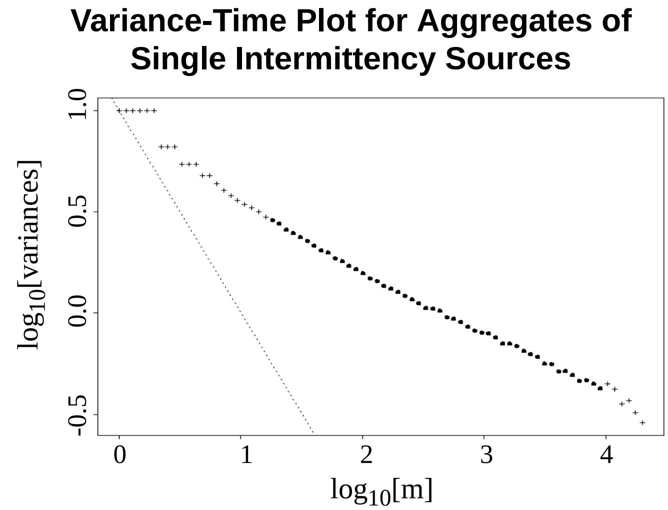

Figure 7 Variance time plot of trace obtained by aggregating 50 sources of the Single Intermittency type: decay rate slower indicating presence of long-range dependence

gates of sources of the *Single Intermittency* type also show convergence to the stretched exponential form (Figure 8), though the quantitative agreement is less precise. This can be attributed to slower convergence to FBM. In particular, note the difference in the decay rate of the queue length distributions of the aggregates of the single intermittency and double intermittency, though the target FBM parameter values (mean, peakedness and the Hurst parameter) are the same. In practice, the excess high-frequency power in the output of the single intermittency map presents difficulties in matching the peakedness. This high-frequency component is relatively inconsequential in determining queueing performance, which is dominated by the long range dependence. Thus for the same total power (related to the peakedness) the single intermittency map generates lesser power in the lower frequencies, and consequently, queueing backlogs are less heavy. Such potential pitfalls should be taken into account while matching FBM parameters to asymtptoically self-similar traffic. For traffic generation using aggregations of the single intermittency map, the high-frequency variations can be eliminated by aggregat-

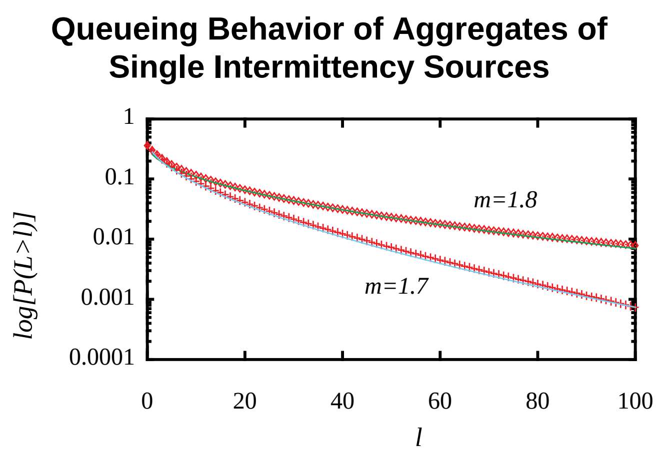

Figure 8 The queue length distribution has a stretched exponential decay for a deterministic server driven by the traffic process generated from the aggregation of 50 *Single Intermittency* type sources. The fit is given by the solid lines which fall almost on the curves.

ing in time.

6 Discussion and Conclusions

We have related self-similar traffic to a variety of ON/OFF source behavior. Using chaotic maps to model a range of ON/OFF source behavior, we show that long-range dependent traces can be generated by sources which have heavy-tailed OFF behavior. Both single sources as well as aggregates of sources with heavy-tailed OFF and light-or heavy-tailed ON behavior generate long-range dependence, as indicated by variance-time plots, power spectra and other statistical indicators. There are nevertheless differences in queueing behavior, based on the nature of the ON period. For single sources, light-tailed ON periods produce queue length distributions that decay exponentially; whereas sources which have heavy-tailed ON periods generate queue length distributions that decay as power laws. Neither of these sources generates the stretched exponential queue length distributions predicted by FBM models.

This paradox - of traffic traces with the similar second-order statistics generating dramatically different queueing behavior - can be explained on the basis of differences in higher-order statistics. This is supported by considering aggregates of heavy-tailed ON/OFF sources, which appear to converge to FBM, and lead to queueing behavior consistent with FBM models, regardless of the nature of the ON period. The impact of higher order statistics is diminished in aggregating independent sources [8]. In general, in assessing impacts of long range dependence on performance, the potential impacts of other factors, such as higher order statistics should not be discounted. In particular the physical basis of the self-similar behavior (sustained inactivity, or activity) can provide insights into the performance impacts of long range dependence. The differences in single source queueing behavior can be readily understood on a physical basis by considering the durations of the ON period. Ethernet sources appear to fit into the category of heavy ON/OFF source behavior [17]; for sources of this type, peak rate allocation may be required, and aggregates appear to converge fairly quickly to FBM (“exactly self-similar”). Loosely speaking, SS7 sources appear to fit into the category of heavy OFF, but light ON sources. For single sources of this type, queueing behavior is rather tame, but aggregates across sources and in time once again lead to FBM (“asymptotically self-similar”). Physically, a traffic stream consisting of many limited bursts will, at a high enough utilization level, be indistinguishable from a stream that is generated by extended bursts. As SS7 traffic levels increase, the impacts of long range dependence may be more significant. The advent of newer SS7 services may accelerate this trend.

There is considerable scope for further work: more in-depth theoretical work to establish the emprirical results presented here; use of general ON/OFF sources to generate exact and asymptotically self-similar traffic; analysis of source aggregations on the basis of flows in the invariant measure of the chaotic maps; the physical basis of self-similarity in various applications etc.

ACKNOWLEDGEMENTS

This work was partly supported by a grant from the Swedish Research Council. We would like to acknowledge the many useful discussions we have had with W. Willinger of Bellcore, USA, and with O. Narayan of Harvard University.

REFERENCES

[1] J. Abate, G.L. Choudhury, and W. Whitt, “Waiting Time Probability in Queues with Long-tail Service Distributions,” Queueing Systems and their Applications, 1994.

[2] J. Beran, R. Sherman, M.S. Taqqu and W. Willinger, “Variable Bit Rate Video Traffic and Long Range Dependence,” accepted for publication in IEEE/ACM Trans on Networking, subject to revisions, 1992.

[3] D. Cox, “Long-Range Dependence: A Review," in Statistics: An Appraisal, David and David (eds.), The Iowa State University Press, 1984.

[4] A. Erramilli, D. Gosby, W.Willinger, “Engineering for Realistic Traffic: A Fractal Analysis of Burstiness”, Proc of Special ITC Seminar, Bangalore, India, 1993.

[5] A. Erramilli and R.P. Singh, “Application of Deterministic Chaotic Maps to Characterize Broadband Traffic”, 7th ITC Specialists Seminar, Morristown, 1990.

[6] A. Erramilli, R.P. Singh, and P. Pruthi, “Chaotic Maps as Models of Packet Traffic,” The Fundamental Role of Teletraffic in the Evolution of Telecommunications Networks, Proc. of the 14th ITC, pp. 329-338, 6-10 June, 1994.

[7] A. Erramilli, R.P. Singh, and P. Pruthi, “Modeling Packet Traffic with Chaotic Maps,” Royal Institute of Technology, ISRN KTH/IT/R-94/18--SE, Stockholm-Kista, Sweden, Aug. 1994.

[8] A. Erramilli, R.P. Singh, and P. Pruthi, “An Application of Deterministic Chaotic Maps to Model Packet Traffic,” to appear Queueing Systems and their Applications, 1995.

[9] A. Erramilli and W. Willinger, “Fractal Properties in Packet Traffic Measurements”, *Proc. of ITC Regional Seminar*, St. Petersburg 1993.

[10] A. Erramilli, W. Willinger and O. Narayan, “Experimental Queueing Analysis with Long Range Dependent Packet Traffic”, submitted to IEEE Transactions on Networking, Sept 1994.

[11] W.E. Leland, M.S. Taqqu, W. Willinger and D.V. Wilson, “On the Self-Similar Nature of Ethernet Traffic”, IEEE/ACM Transactions on Networking, Vol. 2, No.1, 1994.

[12] W.E. Leland and D.V. Wilson, “High Time Resolution measurements and Analysis of LAN Traffic: Implications for LAN Interconnection,” *Proc. of IEEE INFOCOM ‘91*, Bal Harbor, FL, 1991.

[13] K. Meier-Hellstern, P. Wirth, Y-L. Yan and D. Hoeflin, “Traffic Models for ISDN Data Users: Office Automation Application," Proc. ITC-13, Copenhagen, 1991.

[14] I. Norros, “A storage model with self-similar input,” Queueing Systems and their Applications, 1994.

[15] P. Pruthi and A. Erramilli, “Analysis of ON/OFF source behavior with Chaotic Maps”, TR., Bellcore, NJ, 1995.

[16] M.S. Taqqu and J.B. Levy, “Using Renewal Processes to Generate Long Range Dependence and High Variability”, Dependence in Probability and Statistics, Progress in Prob. and Vol 11, Birkhauser; Boston (1986), pp. 73-89.

[17] W. Willinger, “Traffic Modeling of High-Speed Networks: Theory and Practice”, Stochastic Networks, Kelly and Williams eds, Springer-Verlag 1994.