Queuing analysis with traffic models based on deterministic dynamical systems*

Venkat Anantharam

Abstract

Measurements of communication traffic from a wide variety of sources have suggested that correlations persist in such sources over fairly long time scales. This has called into question the use of traditional traffic models for the performance analysis of such networks, and has sparked an interest in studying novel kinds of traffic models. In particular, a suggestion has been made to build models based on expanding deterministic dynamical systems. Deterministic dynamical systems can exhibit chaotic behavior, which has many of the features of statistical behavior, and can indeed be studied rigorously using probabilistic techniques. We make some remarks regarding the analysis of queues driven by such traffic models.

Keywords: chaotic dynamics, long-range dependence, queuing systems

1 Introduction

Statistical analyses of measurements of communication traffic from a wide variety of sources have suggested that correlations persist in such sources over fairly long time scales; for examples of such measurements, see [7], [23], and [29]. This has called into question the use of traditional traffic models for the performance analysis of communication networks, and has sparked considerable research on the use of "long-range dependent" traffic models for such performance analysis; for examples see [2],

*Research supported by NSF Grant No. NCR 9422513 and DARPA Grant No. N66001-00-C-8062

[3], [21], [24], [27] and [28]. A feature of most of these works is the observation that the tail behavior of queues with long-range dependent inputs decays much slower than exponentially; it is either as a subexponential Weibull law (i.e. like exp $(-cx^{\gamma})$ for some $0 < \gamma < 1$) or it decays algebraically, i.e. according to a power law (like $cx^{-\gamma}$ for some $\gamma > 0$), depending on the model. The presence of a qualitative difference in queuing behavior with the traffic measured in [23] as compared to that predicted by conventional models is also supported by the experimental analysis of [17].

The persistence of long-range dependence in packet traffic even after it is regulated by simple flow control schemes is demonstrated in [35]. A bibliography of work in this area as of the middle of 1996 is available in [41].

It should be pointed out that this area is quite active, and the jury is still out on the interpretation and significance of these results. There are at least two general grounds for caution. The first has to do with the very existence of "long-range dependence" (a nice discussion of a somewhat analogous earlier controversy in the field of geology is in [20]).

For instance the analysis of the Ethernet measurements in [23] is primarily done using the R/S statistic, variance-time plots, and the periodogram (for an introduction to these statistical techniques see [6]). The R/S statistic can be highly sensitive to trends in the observations; for instance, a special case of the main result of [8] is that the sum of an independent and identically distributed (i.i.d.) sequence of random variables with finite mean and variance and a deterministic trend that is vanishingly small and goes to zero asymptotically in time according to a power law, can result in the appearance of a nontrivial (not equal to 1/2) Hurst exponent when analyzed using the R/S statistic (for a precise formal theorem, see [8]).

On the other hand, the behavior of variance-time plots for the Bellcore measurements of [23] under different kinds of nonstationarity is discussed in [14], and also in [34] (in the latter it is concluded that the measurements support the hypothesis of long-range dependence fairly well). Two interesting works in this direction (which post date the preliminary presentation of our ideas in [4]) are [1] and [36], where wavelet based estimators for the parameters of long-range dependence are developed. [1] pays particular attention to the issue of robustness of such estimators to the presence of trends, argues that the wavelet based estimators are more robust than earlier known estimators in this regard, and also analyses the data of [23] with this issue in view. In this connection one should mention that a very recently proposed alternative to understanding long-range dependence is the so-called "multifractal models". These have the advantage of being loosely analogous to the layered structure by which application layer traffic is handled as it makes its way from source to destination in the modern Internet. See the papers [31] and [18] for this approach. This approach seems to be motivated also by the wavelet based estimators of long-range dependence; similar estimators can be used to determine the necessary parameters for multifractal modeling of network traffic.

Certainly, the variance-time plots for the AUG89.MB measurements of [23] (plotted in Fig.5 of [23]) would seem to support the existence of correlation up to a scale of about 100 seconds, under the hypothesis of stationarity, as a crude back of the envelope calculation of empirical variances for a stationary process model will show (see the appendix). This brings us to the second ground for caution, which is whether traditional traffic models with sufficiently large depth of correlation might already serve adequately for buffer dimensioning and network design purposes. Recent works that sound this note of caution include [19] and [37].

Whatever side the coin may fall on in this controversy, the measurements and statistical analyses of works such as [7], [23], and [29] are fascinating, and beg for some kind of physical explanation, whether it be due to trends, the nature of the underlying protocols, or the existence of "true" "long-range dependence". At least, such a physical understanding would aid in the construction of parsimonious and realistic traffic models, be they "traditional" or "long-range-dependent".

One attempt at such a physical explanation is in [42], where it is argued via statistical analysis of the measurements in [23] that they are consistent with infinite variance of the load at the level of individual sources. It is proved in [33] that the superposition of ON/OFF sources whose ON-periods or OFF-periods have infinite variance results in an aggregate network traffic that is long-range dependent.

In this note our purpose is to make a couple of remarks in connection with another intriguing idea for arriving at physical descriptions of network communication traffic, initially suggested in the networking context in [16]. This idea is related to a number of mostly numerical investigations in the physics community that has found surprisingly long-lived correlations in the evolution of certain chaotic dynamical systems.

Recall that a stationary dynamical system is called strongly mixing if for any two events the correlation between one of them and the time shifted version of the other asymptotically vanishes in the time shift (for a precise definition, see, for instance, page 57 of [30]). While this implies that the autocorrelation function of an observable asymptotically vanishes, it does not say anything about the rate at which it vanishes. It was largely believed that in sufficiently rapidly mixing systems this rate of decay of correlations would be exponential, for instance this was believed to be the case for the so called K-automorphisms (see page 62 of [30] for a precise definition; it is known that every K-automorphism is strongly mixing, see e.g. page 72 of [32]). However simulations of certain billiard type systems (i.e. dynamical systems describing the evolution of a particle bouncing off certain scattering surfaces) in [40] and [5] suggest that such correlation may decay like a subexponential Weibull law; such a law was earlier proved rigorously to be an upper bound to the decay of correlations in a large class of billiard systems in [10]. Even more interesting is the rigorous proof of algebraically decaying correlation's in an example of a K-automorphism in [26], and the rigorous proof in [12] of algebraically decaying correlations for certain (admittedly somewhat pathological) functions on a dynamical system as basic as that given by scaling transformations on a two dimensional torus.

Related to the ideas of this paper are also the works [38] and [39], both of which also post date [4]. [38] studies the evolution of time averages of congestion windows of multiple TCP sessions sharing a link;

the resulting dynamical system is seen to exhibit many of the features of chaos, including apparently having a strange attractor with fractal dimension roughly 1.61 (for details and definitions see that paper). In [39] it is argued that by virtue of the adaptivity of TCP it is possible for a single TCP microflow to pick up long range dependent characteristics from the aggregate flows that it interacts with as it traverses the network; the latter might be long-range dependent by virtue of being the superposition of heavy tailed activity processes.

2 Deterministic dynamical systems

The proposal in [16] is to model a communication traffic source (such as an Ethernet LAN, or a VBR video source) by means of a deterministic nonlinear transformation $x_{n+1} = S(x_n)$ taking values in some state space X. The traffic source is modeled as having an activity level that depends on the current state.

Let us first remark that if X is allowed to be sufficient general, the restriction to deterministic S is not a big one. Indeed, the theory of deterministic chaos, see for instance [9] or [22], tells us that quite complicated statistical behaviour can be expressed by such deterministic maps. For instance, consider the transformation on the unit interval $f : [0,1] \mapsto [0,1]$ given by

$$S(x) = \begin{cases} \frac{x}{p} & if\ 0 \leq x \leq p \\ 1 - \frac{x-p}{1-p} & if\ p \leq x \leq 1 \end{cases}$$

Then, starting from Lebesgue almost any initial condition x_0, the empirical distribution of the sequence of iterates $\{S^n(x_0), n \geq o\}$ will converge to Lebesgue measure on $[0,1]$ (in the weak topology of convergence of measures). Thus, there appears to be a stationary situation in which the state is distributed according to Lebesgue measure.

Suppose the activity level of the source is

$$a(x) = \begin{cases} a & if\ 0 \leq x \leq p \\ 0 & if\ p < x \leq 1 \end{cases}$$

If one were to start with the initial state distributed according to Lebesgue measure, the sequence of iterates $\{a(x_n), n \geq 0\}$ of the source values will be a sequence of i.i.d. $\{0, a\}$ valued Bernoulli random variables with probability of being a equal to p.

The point just made is that deterministic maps can be used to model stochastic processes also. Thus, by choosing X and the transformation S appropriately it should be possible to directly model the traffic generated by fairly complicated protocols and systems, in a directly tangible and physically meaningful way, while at the same time retaining the flexibility to incorporate traditional stochastic processes into the models. Another argument for the potential interest in such deterministic models is that there appears to be a significant amount of determinism in the structure of communication traffic - in Fig 1. of [16] the successive inter-arrival times from the Ethernet measurements of [23] is plotted, and visual inspection shows what appears to be considerable deterministic structure in this plot. It is also valuable to ponder the evidently chaotic nature of the plots of vectors of averages of congestion window sizes of flows running the TCP protocol, provided in [38].

We now proceed to make a couple of remarks about the queuing behavior of queues driven by such traffic models.

3 Large deviations

We first record a special case of a result of [11], see Thm 3.9 (ii) of that paper.

Let $\{a_n, n \geq 0\}$ be a stationary and ergodic sequence of nonnegative real valued random variables, interpreted as the sequence of arrivals into a single server queue, which can serve an amount c per unit time. Thus, the queue size evolves according to the equation

$$q_{n+1} = (q_n + a_n - c)^+$$

We assume that the stability condition $E[a_n] < c$ holds. From the result of [25], we know that the queue has a unique stationary distribution. Let q_∞ be a random variable with this distribution.

We assume that for any θ, $0 < \theta < \infty$, the limit

$$a^*(\theta) \triangleq \lim_{x\to\infty} \frac{1}{n} \log E\left[e^{\theta \sum_{t=0}^{n-1} a_t}\right] \tag{1}$$

exists. Let us define

$$\theta^* \triangleq \sup\{\theta : a^*(\theta) < c\}$$

Then the claim of [11] is that we have

$$\lim_{u\to\infty} \frac{-\log P(q_\infty \geq u)}{u} = \theta^*$$

In this result, note that, since $\theta^* > 0$, the tail probability of the stationary queue size decreases exponentially.

We next note a special case of a result recorded in [13], see Theorem 6.4.4. and pg. 261 of that book.

Given an R^K valued ψ-mixing sequence of random variables $\{Y_n, n \geq 0\}$, the sequence of empirical averages $\{S_n, n \geq 0\}$, where

$$S_0 \triangleq 0, \quad S_n \triangleq \frac{1}{n} \sum_{t=0}^{n-1} Y_t,$$

obeys a large deviations principle with good convex rate function $\Lambda^*(\cdot)$, which is the Legendre transform of the function

$$\Lambda(\lambda) = \lim_{n\to\infty} \frac{1}{n} \log E\left[c^{n\langle \lambda, S_n\rangle}\right], \tag{2}$$

defined for $\lambda \in R^K, \lambda = (\lambda_1, \cdots, \lambda_K)^T$.

Here $\langle \lambda, S_n \rangle$ denotes the inner product of vectors in R^K. In particular, this result says that the limit in equation (2) exists.

For the definition of ψ-mixing, see the remarks following Assumption H2 on pg. 261 of [13]. It should be clear from this definition that ψ-mixing sequences are strongly mixing.

Other terminology used in the statement of this result (Legendre transform; good rate function) is also defined and discussed at length in [13].

4 Main remarks

Consider a deterministic dynamical system, given by a map $S : X \mapsto X$, where (X, F, μ) is a measure space. A priori, the measure μ has nothing to do with the dynamics of S; it serves as a reference measure to permit talking about densities, and, as is customary, we will assume that S is measurable, and *nonsingular* with respect to μ, i.e. the induced measure $S_*(\mu)$ (which is also often written $\mu(S^{-1})$) is absolutely continuous with respect to μ.

We also assume given a partition of X into K measurable subsets $I_1, \ldots, I_K$, real numbers $a_1, \ldots, a_K$, (which need not be distinct), and define the function $a : X \mapsto \{a_1, \ldots, a_K\}$ by

$$a(x) = \sum_{k=1}^{K} a_k 1_{I_k}(x)$$

where $1_{I_k}(x)$ is the indicator function of I_k.

Suppose we start the dynamical system with a distribution that is absolutely continuous with respect to μ and can therefore be written as $f\mu$ for some density $f \in L^1(\mu)$ (i.e. $f \geq 0$ and $\int xf(x)\mu(dx) = 1$). Then the next state will also have a distribution that is absolutely continuous with respect to μ (because S is nonsingular with respect to μ), so there is an operator P on $L^1(\mu)$ such that $P(f)$ is the density of this next state with respect to μ. P is called the *Frobenius – Perron operator* of the transformation.

Under a wide range of conditions it is known that, for arbitrary initial density f, the sequence of iterates $\{P^n(f), n \geq 0\}$ will converge to a limit f^* such that $P(f^*) = f^*$. For details regarding the sense in which convergence takes place, and situations in which it is known to take place, see [9] and [22].

There is a vast literature on this topic, with conditions known that cover several nonlinear maps with quite nontrivial behavior. For maps of the unit interval $[0, 1]$, for instance, a commonly cited example of such a result is recorded in Theorem 6.2.1 of [22].

Thus, the state of the dynamical system appears to converge, in some appropriate sense, to a random state with distribution $f^*\mu$. If one starts the dynamical system with this distribution, then the sequence of arrivals $\{a_n, n \geq 0\}$ will be a stationary ergodic process. It is certainly the case that under a wide range of conditions, this process will be ψ-mixing. Indeed, the underlying dynamical system itself can be shown to be strong mixing in many cases; for instance, exactness of the transformation, which is often quite easy to see, implies mixing.

Combining these observations with the results stated in the preceding section leads to our first remark : If one constructs such an expanding discrete dynamical system model, and an arrival process derived from such a model, then, in the (non-pathological) situation where a stationary distribution exists for the dynamical system, one should typically expect exponential decay of the tails of the stationary queue size in queues driven by the arrivals.

It suffices to apply the second result of section 3 to the sequence of R^K valued random variables

$$\{(1_{I_1}(x_n), \ldots, 1_{I_K}(x_n)), n \geq 0\}$$

to conclude the existence of the limit in equation (1).

We next observe that positive recurrent discrete time finite state Markov chain arrivals are quite easily constructed along the lines of the general scheme we are discussing. To get a stationary sequence $\{a_n, n \geq 0\}$ taking values in $\{a_1, \ldots, a_K\}$ with stationary distribution $(\pi_k, 1 \leq k \leq K)$ and transition probability matrix $(p_{ij}, 1 \leq i, j \leq K)$, take X to be the unit interval $[0, 1]$, F to be the Borel σ-field, and μ to be Lebesgue measure. Take $I_1, \ldots, I_K$ to be intervals of lengths $\pi_1, \ldots, \pi_K$ respectively, and further partition each I_k into intervals $I_{kl}, 1 \leq l \leq K$, of lengths $\pi_k p_{k1}$ respectively.

Consider the deterministic transformation

$$S : [0, 1] \mapsto [0, 1]$$

which is *piecewise linear* and maps the interval I_{kl} onto the interval I_1 (so that it has slope $\frac{\pi_1}{\pi_k p_{kl}}$ on this interval).

Except for the quibble that a state of the Markov chain may lead to another state of the chain with probability 1, which situation can also be easily handled, this is an example of a piecewise linear (expanding) Markov transformation, in the sense of Chapter 9 of [9], because the slopes are all strictly bigger than 1; it is easily seen to leave Lebesgue measure invariant (so the invariant density is 1). The sequence of arrivals $\{a_n, n \geq 0\}$, when we start the dynamical system with Lebesgue measure, will be a Markov chain with initial distribution $(\pi_k, 1 \leq k \leq K)$ and transition probability matrix $(p_{ij}, 1 \leq i, j \leq K)$.

Our second remark is the following: Suppose we are given any piecewise expanding transformation of the interval (for the definition, see page 85 of [9]). These are among the most studied transformations in the literature, and for maps of the interval an expanding condition is required for most results regarding the existence of a stationary distribution. Then one can approximate the transformation by a piecewise linear (expanding) Markov transformation. The arrival process corresponding to this approximation, in stationarity, is a function of a stationary finite state Markov chain. The limit of equation (1) can now be explicitly written down in terms of the slopes of the approximation to the transformation. This is a basic result that goes back to the seminal works of Donsker and Varadhan for the large deviations of the empirical distribution of Markov processes; for a convenient reference for such results, see [15]. Thus, according to the result of [11], the tail probabilities of a single server queue driven by such an arrival process will decay exponentially (assuming stability), and the rate of decay of the tail can be explicitly computed in terms of the parameters of the approximating Markov transformation.

Note that our second remark can be generalized to transformations of spaces other than the interval - all one needs is that it should be possible to construct a finite state symbolic dynamical scheme that approximates the given transformation.

5 Conclusion

Deterministic dynamical systems offer a promising (or at least suggestive!) modeling paradigm for the communication traffic processes generated by systems executing complex protocols, which might also be composed of several interacting components. It is important to realize that one has not lost modeling flexibility by restricting to deterministic transformations, because stochastic processes can also be modeled by this paradigm. What one appears to gain is a direct physical approach to building models based on an explicit description of the underlying system dynamics.

We have pointed out that when we consider a queue fed by arrivals derived from such a system, we should typically expect to see exponentially decaying tails for the stationary queue size. Further, by choosing suitable approximations for the transformation describing the system, we can write explicit formulas for the rate of decay of the tail probabilities of the stationary queue size.

The developments in the physics community studying the rate of decay of correlation's in chaotic dynamical systems, particularly their attempts to come to grips with the empirically observed subexponential decay of correlations in systems of interest to them, prompted the writing of this note, since it seems plausible that the empirically observed long-range dependence of communication network traffic has a similar underlying mechanism, at least in part.

Acknowledgment

I would like to thank Prof. P. S. Krishnaprasad for bring the paper of Crawford and Cary [12] to my attention.

Appendix

Let $X_1, \ldots, X_N$ be a sequence of conditionally independent Poisson random samples, conditioned on means $\mu_1, \ldots, \mu_N$ respectively.

Assume $\mu_n = \lambda + \gamma_n, 1 \leq n \leq N$. This is a simple model for a sequence of packet or byte counts with fluctuating mean.

The samples are aggregated at level m to get

$$\bar{X}_j^{(m)} = \frac{1}{m}(X_{(j-1)m+1} + \ldots + X_{jm}), 1 \leq j \leq M,$$

where $1 \leq n \leq N$, m divides N, and M denotes N/m. The variance time plot is the plot of the logarithm of the empirical variance of the m-aggregated sequence $\bar{X}_1^{(m)}, \ldots, \bar{X}_M^{(m)}$ against $\log m$.

Slightly at variance with the usual definition on Section 4.4. on page 94 of [6], let us define this as

$$\frac{1}{M}\sum_{j=1}^{M}\left(\bar{X}_j^{(m)} - \bar{X}(m)\right)^2,$$

where $\bar{X}(m)$ denotes the empirical mean at granularity m :

$$\bar{X}(m) = \frac{1}{M}\sum_{j=1}^{M}\bar{X}_j^{(m)}.$$

Write Γ_n for $\sum_{t=1}^{n}\gamma_t$. The conditional expectation of the empirical variance of the m-aggregated sequence is then seen to be

$$\frac{\lambda + \Gamma_N/N}{m} - \frac{\lambda + \Gamma_N/N}{N} + \frac{1}{Nm}\sum_{j=1}^{M}(\Gamma_{jm} - \Gamma_{(j-1)m})^2 - \frac{\Gamma_N^2}{N^2}$$

Deriving this is a standard calculation using only the formula $\alpha^2+\alpha$ for the second moment of a Poisson random variable of mean α.

Thus, loosely speaking, if $\Gamma_N = o(N)$, and, for most j, $(\Gamma_{jm} - \Gamma_{(j-1)m})^2 \sim m^{2H}$ over the range $N^\varepsilon \leq m \leq N^{1-\varepsilon}$, then the variance time plot will exhibit a slope of $2 - 2H$ over this range, i.e. over $c \log N \leq \log m \leq (1-c)\log N$.

On the basis of this crude calculation the variance time plot of Fig. 5(b) of [23] for the AUG89.MB measurements described there, which shows a slope of about -0.40 corresponding to roughly $H = 0.80$ over roughly the range $1 \leq \log_{10} m \leq 1$ from a sample of size $N = 360,000$ observations taken every 10 msec, could be taken to indicate correlations on the scale of about 100 seconds, i.e. $10^4 \times 10$ msec, but do not necessarily indicate correlation's on longer time scales than this.

References

[1] Abry, P. and Veitch, D. (1998). Wavelet Analysis of Long-Range Dependent Traffic. *IEEE Transactions on Information Theory*, Vol. 44, No. 1, pp. 2–15

[2] Anantharam, V. (1995). On the Sojourn Time of Sessions at an ATM Buffer with Long-Range Dependent Input Traffic. *Proceedings of the 34rd IEEE Conference on Decision and Control*, New Orleans, Louisiana, December 13 -15, 1995, Volume 1, IEEE Press, pp. 859–864.

[3] Anantharam, V. (1996). Networks of Queues with Long-Range Dependent Traffic Streams, Stochastic Networks: *Stability and Rare Events*, edited by Paul Glasserman, Karl Sigman, and David Yao, Lecture Notes in Statistics, Vol. 117, Springer Verlag, pp. 237–256.

[4] Anantharam, V. (1997). Queuing analysis with traffic models based on deterministic dynamical systems, *Proceedings of the 35th Allerton Conference on Communications, Control, and Computing*, Department of Electrical and Computer Engineering, University of Illinois, Champaign, Illinois, pp. 233–241.

[5] Artuso, R., Casati, G., and Guarneri, I. (1995). Geometric scaling of correlation decay in chaotic billiards. *Physical Review E*, Vol. 51, No. 5, pp. R3807–R3810.

[6] Beran, J. (1994). *Statistics for Long-Memory Processes*. Monographs on Statistics and Applied Probability, Vol. 61, Chapman and Hall, New York.

[7] Beran, J., Sherman, R., Taqqu, M.S., and Willinger, W. (1995). Long-range dependence in variable-bit-rate video traffic. *IEEE Transactions on Communications*, Feb.-April 1995, vol.43, Nos. 2-4, pt.3, pp. 1566–1579.

[8] Bhattacharya, R. N., Gupta, V. K., and Waymire, E. (1983). The Hurst Effect under Trends. *Journal of Applied Probability*, Vol. 20, pp. 649–662.

[9] Boyarsky, A. and Gora, P. (1997). *Laws of Chaos : Invariant Measures and Dynamical Systems in One Dimension.* Birkhauser, Boston.

[10] Bunimovich, I. A. and Sinai Ya. G. (1981). Statistical Properties of a Lorentz Gas with Periodic Configuration of Scatterers. *Communications in Mathematical Physics*, Vol. 78, pp. 479–497.

[11] Chang, C. S. (1994). Stability, Queue Length, and Delay of Deterministic and Stochastic Networks.*IEEE Transactions on Automatic Control*, Vol. 39, No. 5, pp. 913–931.

[12] Crawford, J. D. and Cary J. R. (1983). Decay of Correlation's in a Chaotic Measure-Preserving Transformation. *Physica 6D*, pp. 223–232.

[13] Dembo, A. and Zeitouni, O. (1993). *Large Deviation Techniques and Applications.* Jones and Bartlett Publishers, Boston.

[14] Duffield, N.G., Lewis, J.T., O'Connell, N., Russell, R., and Toomey, F. (1994). Statistical issues raised by the Bellcore data. *IEE Eleventh UK Teletraffic Symposium. Performance Engineering in Telecommunications Networks,* Cambridge, UK, Vol. 1 pp. 1–8.

[15] Ellis, R. S. (1988). Large Deviation for the Empirical Measure of a Markov Chain with an Application to the Multivariate Empirical Measure. *Annals of Probability*, Vol. 16, pp. 1496–1508.

[16] Erramilli, A. and Singh, R. P. (1995). An application of deterministic chaotic maps to model packet traffic. *Queuing Systems : Theory and Applications*, Vol. 20, pp. 171–206.

[17] Erramilli, A., Narayan, O., and Willinger, W. (1996). Experimental queuing analysis with long-range dependent packet traffic.*IEEE/ACM Transactions on Networking*, April 1996, vol.4, no.2, pp. 209–23.

[18] Feldmann, A., Gilbert, A. C., and Willinger, W. (1998). Data Networks as Cascades : Investigating the Multifractal Nature of Internet WAN Traffic. *Proceedings of the ACM SIGCOMM 1998*, ACM Press, pp. 42–55.

[19] Heyman, D.P. and Lakshman, T.V. (1996). What are the implications of long-range dependence for VBR-video traffic engineering?. *IEEE/ACM Transactions on Networking*, June 1996, vol.4, no.3, pp. 301–317.

[20] Klemes, V. (1974). The Hurst phenomenon : a puzzle ? *Water Resources Research*, Vol. 10, pp. 675–688.

[21] Konstantopoulos, T. and Lin, S-J. (1996). Fractional Brownian Motion and Levy Motions as Limits of Stochastic Traffic Models. *Proceedings of the 34th Annual Allerton Conference on Communications, Control, and Computing*, Urbana, Illinois, pp. 913–922.

[22] Lasota, A. and Mackey, M. C. (1994). *Chaos, Fractals, and Noise : Stochastic Aspects of Dynamics*. Springer Verlag, New York.

[23] Leland, W. E., Taqqu, M. S., Willinger, W. and Wilson, D. V. (1994). On the Self-Similar Nature of Ethernet Traffic (Extended Version). *IEEE/ACM Transactions on Networking*, Vol. 2, No. 1, pp. 1–15.

[24] Likhanov, N., Tsybakov, B. and Georganas, N. D. (1995). Analysis of an ATM Buffer with Self-Similar ("Fractal") Input Traffic. *Proceedings of the 14th Annual IEEE Infocom*, IEEE Press, pp. 985–992.

[25] Loynes, R. M. (1962). The Stability of Queues with Non-Independent Interarrival and Service Times. *Proceedings of the Cambridge Philosophical Society*,Vol. 58, pp. 497–520.

[26] Niwa, T. (1978). Time Correlation Functions of a One-Dimensional Infinite System. *Journal of Statistical Physics*, Vol. 18, No. 3, pp. 309–317.

[27] Norros, I. (1994). A storage model with self-similar input. *Queuing Systems : Theory and applications*, Vol. 16, pp. 387–396.

[28] Parulekar, M. and Makowski, A. M. (1996). Tail Probabilities for a Multiplexer with Self-similar Traffic. *Proceedings of the IEEE INFOCOM 1996*, IEEE Press, pp. 1452–1459.

[29] Paxson, V. and Floyd, S. (1995). Wide Area Traffic: The Failure of Poisson Modeling. *IEEE/ACM Transactions on Networking*, Vol. 3, No. 3, pp. 226–244.

[30] Petersen, K. (1983). *Ergodic Theory* Cambridge studies in advanced mathematics, Vol. 2, Cambridge University Press, New York.

[31] Reidi, R. H. and Crouse, M. S. (1999). A Multifractal Wavelet Model with Application to Network Traffic. *IEEE Transactions on Information Theory*, Vol. 45, No. 3, pp. 992–1018.

[32] Sinai, Ya. G. (1994). *Topics in Ergodic Theory.* Princeton Mathematical Series, Vol. 44, Princeton University Press, Princeton, New Jersey.

[33] Taqqu, M.S., Willinger, W., and Sherman, R. (1997). Proof of a fundamental result in self-similar traffic modeling. *Computer Communication Review*, April 1997, vol.27, no.2, pp. 5–23.

[34] Teverovsky, V. and Taqqu, M. (1995). Testing for Long-Range-Dependence in the Presence of Shifting Means or a Slowly Declining Trend, using a Variance-type Estimator. *Journal of Time Series Analysis*, Vol. 18, No. 3, pp. 270–304.

[35] Vamvakos, S. and Anantharam, V. (1998). On the Departure Process of a Leaky Bucket System with Long-Range Dependent Input Traffic. *Queuing Systems : Theory and Applications*, Vol. 28, Nos. 1-3, pp. 191–214.

[36] Veitch, D. and Abry. P (1999). A Wavelet-Based Joint Estimator of the Parameters of Long-Range Dependence. *IEEE Transactions on Information Theory*, Vol. 45, No. 3, pp. 878–897.

[37] Venkateshwara Rao, B., Krishnan, K.R., and Heyman, D.P. (1996). Performance of finite-buffer queues under traffic with long-range dependence. *Proceeding of the IEEE GLOBECOM 1996*, IEEE Press, pp. 607–611.

[38] Veres, A. and Boda, M. (2000). The Chaotic Nature of TCP Congestion Control. *Proceedings of the IEEE INFOCOM 2000*, IEEE Press, pp. 1715–1723.

[39] Veres, A., Kenesi, Zs., Molnar, S., and Vattay, G. (2000). On the Propagation of Long-Range Dependence in the Internet. *Proceedings of the IEEE SIGCOMM 2000*, ACM Press, pp. 243–254.

[40] Vivaldi, F., Casati, G., and Guarneri, I. (1983). Origin of Long-Time Tails in Strongly Chaotic Systems. *Physical Review Letters*, Vol. 51, No. 9, pp. 727–730.

[41] Willinger, W., Taqqu, M.S., and Erramilli, A. (1996). A Bibliographical Guide to Self-similar Traffic and Performance Modeling for Modern High-speed Networks. In {\em Stochastic Networks : *Theory and Applications*, edited by F. Kelly, S. Zachary, and I. Ziedins, Oxford University Press, Oxford, pp. 339–366.

[42] Willinger, W., Taqqu, M.S., Sherman, R., and Wilson, D.V. (1997). Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level. *IEEE/ACM Transactions on Networking*, Feb. 1997, vol.5, no.1, pp. 71–86.

Venkat Anantharam

Received October 27, 2004

EECS Department, University of California,
Berkeley, CA 94720
E–mail: *ananth@vyasa.eecs.berkeley.edu*