TCP AIMD Dynamics Over Internet Connections

Jianbo Gao and Nageswara S. V. Rao

***Abstract*— The precise nature of TCP dynamics over Internet connections has not been well understood, since the existing results are solely analytical or based on simulations. We employ the time-dependent exponent curves and logarithmic displacement curves to study TCP AIMD congestion window-size traces over Internet connections. We show that these dynamics have two dominant parts, a stochastic component in response to network traffic and a deterministic chaotic component due to the nonlinearity of protocol. These dynamics can be largely characterized as anomalous diffusions with a very large exponent.**

***Index Terms*— TCP AIMD, transport dynamics, chaotic dynamics, Internet traces.**

I. Introduction

FOR next generation Internet applications such as remote instrument control and computational steering, it is important to understand the transport dynamics to sustain the needed control loops over wide-area connections. For example, high jitter levels can lead to destabilized control, which in turn could damage remote instruments or result in runaway computations on supercomputers. Several applications currently utilize TCP (Transmission Control Protocol), a widely deployed transport protocol, to implement such control loops. These TCP dynamics are a result of the highly non-linear nature of its AIMD (Additive Increase and Multiplicative Decrease) congestion control mechanism reacting to stochastic dynamics of Internet traffic, more precisely the packet losses. Thus it is natural to expect TCP dynamics to be complicated, as indicated by preliminary studies [2], [21], [17], [16]. The first study on TCP dynamics using chaos theory is by Veres and Boda [21]. Using the ns-2 simulator, they have studied the congestion window dynamics of competing TCP streams. While these works demonstrate that transport dynamics can be chaotic, their scope is limited since they rely solely on simulation [21] or analysis [17], or use network configurations that are not typical of current Internet such as active queues [16]. To a large extent the exact nature of TCP dynamics over actual Internet connections is still an open question. There are two major difficulties: (a) technologically, high quality measurements of network transport variables are hard to collect over live Internet connections, and (b) analytically, it is very difficult to handle the co-existence of deterministic and stochastic components of transport dynamics. As a result, methods from standard chaos and linear stochastic theories are unable to offer much new perspective on TCP dynamics.

Manuscript received March 12, 2004. The associate editor coordinating the review of this letter and approving it for publication was Prof. Changcheng Huang.

Jianbo Gao is with the Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611 (e-mail: gao@ece.ufl.edu).

Nageswara S. V. Rao is with the Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831 (e-mail: raons@ornl.gov).

Digital Object Identifier 10.1109/LCOMM.2005.01018.

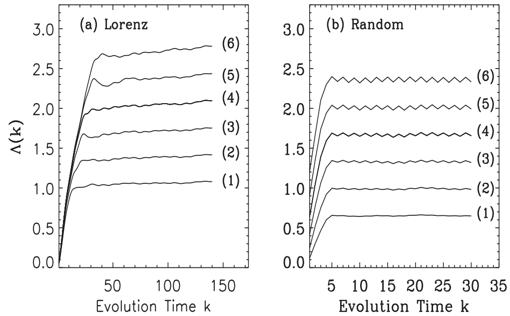

Fig. 1. $\Lambda(k)$ curves for (a) the chaotic Lorenz attractor with $m = 4$, $L = 3$ and sampling time 0.03, and (b) a random process with $m = 6$, $L = 1$. Curves numbered $1, \cdots, 6$ correspond to shells with sizes $(2^{-(i+1)/2}, 2^{-i/2})$, $i = 5, 6, \cdots, 10$.

We utilize the recent web100/net100 [13] instruments to collect time series of TCP congestion window-sizes for Internet connections. We analyze these series using the time-dependent exponent curves [4] and the logarithmic displacement curves [5]. Our analysis shows the dominance of both chaotic and stochastic components in TCP AIMD dynamics. Indeed, these dynamics can be largely characterized as anomalous diffusions with a very large diffusional exponent that arise in other disciplines [9]. This note assumes basic terminology from chaos theory which can be found in books [1].

II. Chaos and Noise

For the congestion window-size data $w(i),\ i = 1, \cdots, n$, we construct vectors of the form: $V_i = [w(i), w(i+L), \ldots, w(i+(m-1)L)]$, where m is the embedding dimension and L the delay time. The embedding theorem [20] states that when the embedding dimension m is larger than twice the box counting dimension of the attractor, then the dynamics of the original system can be studied from a single scalar time series. Incidentally, the embedding dimension used by Veres and Boda [21] is 2, which is low and might call for a re-examination of some of their conclusions.

The time-dependent exponent curves $\Lambda(k)$ are defined by [4]:

$$\Lambda(k) = \left\langle \ln \left(\frac{\|V_{i+k} - V_{j+k}\|}{\|V_i - V_j\|} \right) \right\rangle \qquad (1)$$

where the angle brackets denote ensemble averages of all possible pairs of (V_i, V_j), and k is called the evolution time. The computation is carried out for a sequence of shells, $r \leq \|V_i - V_j\| \leq r + \Delta r$, where r and Δr are prescribed small distances. For true low-dimensional chaotic systems, the curves $\Lambda(k)$ for different shells form a common envelope. The slope of the envelope estimates the largest positive Lyapunov exponent. An example is given in Fig. 1(a) for the well-known chaotic Lorenz system, which shows a common envelope at the lower left corner. The existence of that common envelope

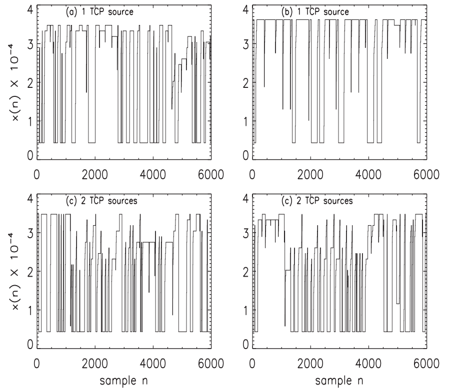

Fig. 2. Time series for the congestion window-size (in Bytes) measured on the Internet. $w(n)$ has been scaled by 10^4.

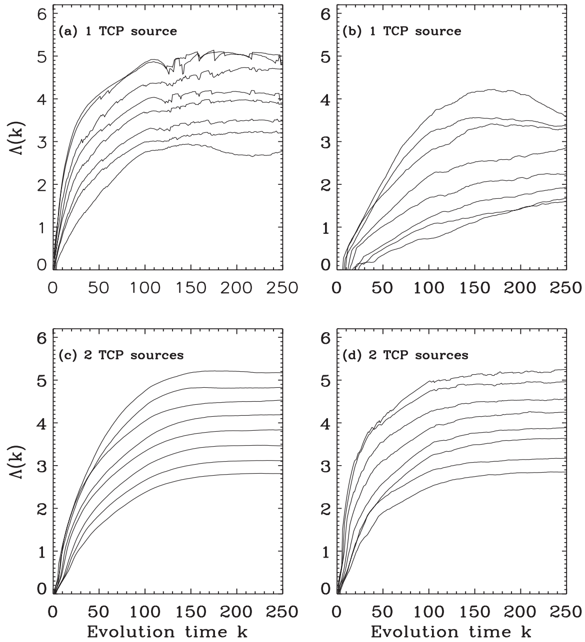

Fig. 3. $\Lambda(k)$ curves for the congestion window-size corresponding to Fig. 2.

guarantees a robust positive estimate for Lyapunov exponent no matter which shell is used in computation. For non-chaotic systems, however, the common envelope is absent, as shown for example in Fig. 1(b) for the $\Lambda(k)$ curves of a realization of uniformly distributed random variable. At this point it is appropriate to comment that it is often assumed that a positive numerically estimated Lyapunov exponent and non-integral value for the fractal dimension are a sufficient indicator of chaos in a time series. Since $\Lambda(k)/(k\delta t)$, where δt is the sampling time, is more or less equivalent to the largest positive Lyapunov exponent (which can be obtained using conventional ways), we thus see that under this common assumption, random processes or noise can be easily interpreted as chaos. Hence, the common envelope represents a more stringent test for deterministic chaos.

We point out some interesting features of the $\Lambda(k)$ curves: (i) For uniform random data, only for k up to the embedding window-size $(m-1)L$ will $\Lambda(k)$ increase. This can be easily seen from Fig. 1(b). Thus, whenever $\Lambda(k)$ increases for a much larger range of k, it is an indication of non-trivial deterministic structure in the data. (ii) For periodic signals, $\Lambda(k)$ is essentially zero for any k. (iii) For quasi-periodic signals, $\Lambda(k)$ is periodic with an amplitude typically smaller than 0.1, hence, for practical purposes, $\Lambda(k)$ can be considered very close to 0.

III. TCP Traces and Analysis

We measured the time series of congestion window-size for ftp sessions of Gigabyte files over the Internet connection between Oak Ridge National Laboratory and Louisiana State University, which is approximately one thousand network miles long. The sampling time is approximately 1 millisecond and the round trip time is approximately 30 milliseconds. To capture the deterministic component of the dynamics, larger sampling time will not be appropriate. To find out whether one or two TCP sources might make a difference, we measured the congestion window-size with only one or two TCP sources (i.e., single or dual ftp sessions). Here, we present results for four datasets, two with two TCP sources and two with a single TCP source, and we plot this data in Figs. 2(a-d). Power spectral analysis of these data does not show any dominant peaks, and hence the dynamics are not simply oscillatory. The $\Lambda(k)$ curves, corresponding to Fig. 2, are plotted in Fig. 3. In the computations, 3×10^4 points are used, and $m = 10,\ L = 1$. The eight curves, from the bottom to top, correspond to shells of sizes $(2^{-(i+1)/2}, 2^{-i/2})$, $i = 8, 9, \cdots, 15$. We note the following interesting features: (i) The dynamics are complicated and cannot be described as either periodic or quasi-periodic motions, since $\Lambda(k)$ is much larger than 0. (ii) The dynamics cannot be characterized as pure deterministic chaos, since in no case can we observe a well-defined linear envelope. (iii) The data is not simply noisy, since otherwise we should have observed that $\Lambda(k)$ is almost flat when $k > (m-1)L$. The importance of features (ii) and (iii) is that the TCP dynamics contain *both* deterministic and stochastic components. (iv) There are considerable difference between the data with only one TCP source and with two competing TCP sources. In the latter case, the $\Lambda(k)$ curves sharply rise when k just exceeds the embedding window-size, $(m-1)L$. On the other hand, $\Lambda(k)$ for Fig. 3(b) with only one TCP source increases much slower when k just exceeds $(m-1)L$. Also important is that the $\Lambda(k)$ curves in Figs. 3(c,d) are much smoother than those in Figs. 3(a,b). Hence, we can say that the deterministic component of the dynamics is more visible when there are more than one competing TCP sources. This is understandable, since when there are multiple TCP sources, TCP would then play a more dominant role.

Since our data was measured over the Internet connections, which always have background traffic, we believe this data is profoundly more complicated than those obtained from ns-2 simulations [21]. Nevertheless, we anticipate that with two TCP sources, the deterministic component of the network dynamics might be stronger. This was the case, as seen above.

Since the increasing part of the $\Lambda(k)$ curves are not very linear, we next examine if in those regions of k, the displacement $||V_{i+k} - V_{j+k}||$ curves actually increase with k in a power-law manner,

$$||V_{i+k} - V_{j+k}|| \sim k^{\alpha},$$

where α is called the diffusional exponent. Note that $||V_{i+k} - V_{j+k}||$ is the numerator of Eq. (1). For Brownian motions, $\alpha = 1/2$ [7]. This leads to a classification of diffusional processes [9]: (i) $\alpha = 1/2$: normal diffusion. (ii) $\alpha > 1/2$: anomalous diffusion. This type of diffusion plays a key role in the study of noise-induced chaos. In fact, in that context, it may be termed "pre-noise-induced chaos [8]." (iii) $\alpha < 1/2$: sub-diffusion. When the diffusional process takes place near a limit cycle, this behavior is often due to a strong convergent flow to the limit cycle [6]. We plot the evolution of $||V_{i+k} - V_{j+k}||$ with k in a log-log scale as in Figs. 4(a-d). Very interestingly, we have now indeed observed a collection of better defined linear lines. It is most interesting to note that the diffusional exponent is larger than 1/2 for Figs. 4(a,c,d). Hence, we can say that the diffusion is often anomalous. In fact, α is about 1.2 for both Figs. 4(c,d) and about 1.0 for Fig. 4(a). Fig. 4(b) is an exception: for not too small shells, α is close to 1/2. However, when the scale or the shell size is very small, α is about 1.0. This crossover of anomalous diffusion to normal diffusion from small to fairly large scales suggests that the deterministic component of the dynamics at small scales is gradually masked by the stochastic component at large scales. Depending on how noisy the background traffic is, sometimes this crossover may not happen, as in the case of Fig. 4(a). In fact, such crossover is often inhibited when there are competing TCP sources, as in Figs. 4(c,d).

The above qualitative behavior is representative of TCP AIMD dynamics over the Internet. They remained quite robust when the measurements were repeated tens of times for this connection and also for other Internet connections of lengths ranging from few hundred to few thousand miles.

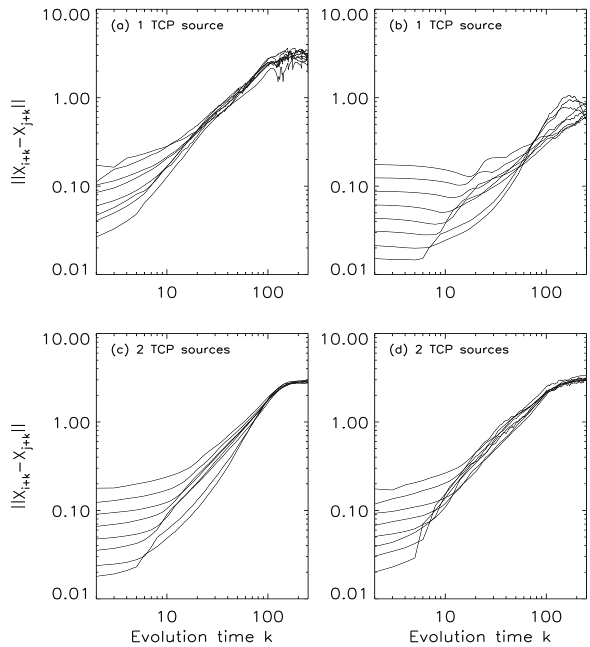

Fig. 4. A log-log plot of the displacement curves for Fig. 3.

IV. IMPLICATIONS

Our analysis provided a new insight into TCP AIMD dynamics over Internet connections that they have a complex mixture of chaotic and stochastic components. There have been several works on TCP dynamics to improve its congestion control [3], [12], [10], [11], [19]. Our approach could be very useful in evaluating these methods, since the conventional methods are not effective. The chaotic component of transport dynamics can be avoided by non-AIMD methods, such as TCP Vegas. But the existence of randomness due to network traffic conditions could be problematic, must be explicitly accounted for to achieve stable dynamics; Rao *et al* [18] utilized the Robbins-Monro method to achieve stable throughput over Internet connections.

ACKNOWLEDGMENTS

This research is sponsored by Defense Advanced Research Projects Agency under MIPR No. K153, National Science Foundation under Grants No. ANI-0229969 and No. ANI-0335185, and by Engineering Research Program and High-Performance Networking Program of Office of Science, U. S. Department of Energy under Contract No. DE-AC0500OR22725 with UT-Battelle, LLC.

REFERENCES

[1] K. T. Alligood, T. D. Sauer, and J. A. Yorke, *Chaos: An Introduction to Dynamical Systems*. New York: Springer-Verlag, 1997.
[2] A. Erramilli and L. J. Forys, "Traffic synchronization effects in teletraffic systems," *Proc. ITC-13*, Copenhagen, 1991.
[3] V. Firoiu and M. Borden, "A study of active queue management for congestion control," *Proc. INFOCOM*, 2000.
[4] J. B. Gao and Z. M. Zheng, "Local exponential divergence plot and optimal embedding of a chaotic time series," *Phys. Lett. A*, vol. 181, pp. 153-158, 1993.
[5] J. B. Gao, *Physica D*, vol. 106, p. 49, 1997.
[6] J. B. Gao, S. K. Hwang, and J. M. Liu, "Effects of intrinsic spontaneous-emission noise on the nonlinear dynamics of an optically injected semiconductor laser," *Phys. Rev. A*, vol. 59, pp. 1582-1585, 1999.
[7] J. B. Gao and W. W. Tung, "Pathological tremors as diffusional processes," *Biological Cybernetics*, vol. 86, pp. 263-270, 2002.
[8] J. B. Gao, S. K. Hwang, and J. M. Liu, "When can noise induce chaos?" *Phys. Rev. Lett.*, vol. 82, pp. 1132-1135, 1999.
[9] J. B. Gao *et al.*, "Noise-induced chaos," *Int. J. Mod. Phys. B*, vol. 13, pp. 3283-3305, 1999.
[10] C. Hollot *et al.*, "A control theoretic analysis of RED," *Proc. INFOCOM*, 2001.
[11] P. Kuusela, P. Lassilaand, and J. Virtamo, "Stability of TCP RED Congestion Control," http://www.tct.hut.fi/tutkimus/com2/publ/, 2000.
[12] V. Misra, W. Gong, and D. Towsley, "A study of active queue for congestion control," *Proc. SIGCOMM*, 2000.
[13] http://www.net100.org
[14] K. Park and W. Willinger, *Self-Similar Network Traffic and Performance Evaluation*. New York: Wiley, 2000.
[15] N. H. Packard *et al.*, "Geometry from a time series," *Phys. Rev. Lett.*, vol. 45, p. 712, 1980.
[16] P. Ranjan, E. H. Abed, and R. La, "Nonlinear instabilities of TCP-RED," *IEEE/ACM Trans. Networking*, 2004, in press..
[17] N. S. V. Rao and L. O. Chua, "On dynamics of network transport protocols," *Proc. of Workshop on Signal Processing, Communications, Chaos and Systems*, 2002.
[18] N. S. V. Rao, Q. Wu, and S. S. Iyengar, "On throughput stabilization of network transport," *IEEE Commun. Lett.*, vol. 8, pp. 66-68, 2004.
[19] R. Srikant, *The Mathematics of Internet Congestion Control*. Birkhauser: Boston, 2004.
[20] F. Takens, "Detecting strange attractors in turbulence," in *Dynamical Systems and Turbulence*, Lecture Notes in Mathematics, ed. by D. A. Rand and L. S. Young, vol. 898. Berlin: Springer-Verlag, p. 366, 1981.
[21] A. Veres and M. Boda, "The chaotic nature of TCP congestion control," *Proc. INFOCOM*, pp. 1715-23, 2000.