

DETRENDED FLUCTUATION ANALYSIS OF THE TCP-RED ALGORITHM*

XI CHEN, SIU-CHUNG WONG and CHI K. TSE[†]
Department of Electronic and Information Engineering
Hong Kong Polytechnic University, Hong Kong
[†]*encktse@polyu.edu.hk*

LJILJANA TRAJKOVIĆ
School of Engineering Science
Simon Fraser University, Vancouver, British Columbia, Canada

Received December 20, 2008; Revised May 5, 2009

It has been observed that Internet gateways employing Transport Control Protocol (TCP) and the Random Early Detection (RED) control algorithm may exhibit instability and oscillatory behavior. Most control methods proposed in the past have been based on analytical models that rely on statistical measurements of network parameters. In this paper, we apply the detrended fluctuation analysis (DFA) method to analyze stability of the TCP-RED system. The DFA is used to analyze time-series data and generate *power-law scaling exponents*, which indicate the long-range correlations of the time series. We quantify the stability of the TCP-RED system by examining the variation of the DFA power-law scaling exponent when the system parameters are varied. We also study the long-range power-law correlations of TCP window periods.

Keywords: DFA, TCP/IP protocol, RED gateway, dynamics, stability.

1 Introduction

Internet applications, such as world wide web, file transfer, and email, are transmitted using the Transmission Control Protocol (TCP) [Jacobson, 1988; Jacobson, 1990; Brakmo & Peterson, 1995]. TCP enables multiple information sources to compete for a fair share of physical bandwidth allocated between two end-points of a connection. When the number of sources increases, control of the traffic

*An earlier version of this paper was presented at the *2008 IEEE International Symposium on Circuits and Systems*.

flows becomes necessary. Active Queue Management (AQM) has been used to control the traffic flow in bottleneck gateways. The Random Early Detection (RED) [Floyd & Jacobson, 1993] gateway was introduced to provide early feedback to the TCP senders by marking (or dropping) packets before buffers become full thus preventing large number of packets from being lost. However, it has been reported that TCP gateways employing RED may exhibit instability and oscillations that degrade the transmission performance [Misra *et al.*, 2000; Firoiu & Borden, 2000; Floyd *et al.*, 2001; Ranjan *et al.*, 2004; Gao *et al.*, 2005; Nga *et al.*, 2008].

The TCP-RED algorithms have been implemented in the *ns-2* network simulator [DARPA, 2007]. The *ns-2* simulator enables analysis and verification of communication networks in a controlled environment. A number of analytical models have been developed to characterize the TCP-RED system and analyze its oscillatory behavior [Chen *et al.*, 2007; Chen *et al.*, 2008; Tinnakornsrisuphap & Makowski, 2003; Kuusela *et al.*, 2001a; Kuusela *et al.*, 2001b]. The stability conditions derived using these analytical models often depend on statistical measurements of various parameters and their upper and/or lower bounds. These measured parameters may be non-stationary.

The Hurst parameter H has traditionally been used to describe the degree of self-similarity of a time series. Several methods may be used to estimate the Hurst parameter, such as the absolute value (AV) method, R/S plot, variance method, Whittle estimator, and periodogram. However, traditional methods may be inaccurate in estimating scaling properties such as self-similarity and long-range dependence (LRD) [Molnár *et al.*, 1997; Teverovsky & Taqqu, 1997; Roughan & Veitch, 1999; Figueiredo *et al.*, 2005]. The inaccuracy of these methods is often due to the non-stationarity of the traffic data and the use of finite sets of trace data. The Hurst parameter values $0 < H < 0.5$, $H = 0.5$, and $0.5 < H < 1$ correspond to *power-law anti-correlation*, *memoryless process*, and *persistent long-range power-law correlation*, respectively. The characteristic of power-law anti-correlation in a time series is that it assumes some large values after it has assumed enough small values, and vice versa. The white noise is an example of memoryless process.

The detrended fluctuation analysis (DFA) has recently become a widely used technique for the determination of fractal scaling properties and the detection of long-range correlations in noisy, non-stationary time series. The DFA method has been employed to shown a relation between the DFA exponent β and the Hurst parameter H [Movahed *et al.*, 2006]:

$$\begin{cases} H = \beta & \text{if } 0 < \beta < 1 \\ H = \beta - 1 & \text{if } \beta \geq 1. \end{cases} \quad (1)$$

DFA is a method for detecting long-range power-law correlations in seemingly non-stationary, noisy or randomized signals [Peng *et al.*, 1994; Peng *et al.*, 1995; Hu *et al.*, 2001; Chen *et al.*, 2002; Chen *et al.*, 2005]. It has been widely used in computer science, biodynamics, bioinformatics, economics, and meteorology [Peng *et al.*, 1994; Buldyrev *et al.*, 1995; Peng *et al.*, 1995; Heneghan & McDarby, 2000; Siwy *et al.*, 2002; Janosi & Muller, 2005; Santhanam *et al.*, 2006; Masugi, 2006; Carpena *et al.*, 2007]. The outcome is usually given in terms of a parameter called *power-law scaling exponent*. Given a time

series (waveform), a chosen statistical variable is observed for various time scales. When this statistical property is plotted over various time scales, a straight line on a log-log scale implies a long-range power-law correlation for the observed variable. Its slope is equal to the power-law scaling exponent. This line may display a deflection (*crossover*) at a certain time scale where the slope abruptly changes. The interpretations of power-law scaling exponents and the crossovers are system dependent. In this paper, we apply the DFA method to two time series of a TCP-RED system: the period of individual TCP windows and the queue length in the bottleneck RED gateway of the TCP-RED system that may undergo a change of stability under variation of system parameters. We analyze the change in scaling exponents and the locations of the crossovers. The main contribution of this work is that the scaling exponent found by applying DFA may be used as a stability indicator for TCP-RED gateways. We also find that individual TCP flows may have LRD properties, which are closely related to the instability of the TCP-RED system. Furthermore, we propose a model to describe the similarity between the flows and show that there is observable dissimilarity between the stable and unstable cases.

The paper is organized as follows. In Section 2, we describe the TCP-RED system and illustrate the DFA method. In Section 3, we analyze the long-range power-law correlation property of the queue length using DFA and provide a physical interpretation of the obtained power-law scaling exponent values. Long-range power-law correlations for individual TCP connections are analyzed in Section 4. We conclude with Section 5.

2 Overview of Algorithms: TCP, RED, and DFA

2.1 TCP Congestion Control Algorithm

We assume a window based flow control between the two end-points of a TCP connection. The TCP sender organizes a stream of data into a sequence of windows of packets. Each window contains w packets. When the TCP receives a packet from the sender, it returns an acknowledgment packet. Upon receiving the first acknowledgment packet from the last window of data sent, the TCP sender begins sending packets within the next window. The duration between sending a packet and receiving the acknowledgment is called the *round-trip-time* (RTT). Therefore, a TCP sender may send w packets for each RTT. If the window size w is adaptively selected to reach its maximum share of throughput according to the network conditions, the TCP communication eventually reaches its steady-state as a self-regulated RTT clocked system. Several TCP algorithms have been employed for adjusting the window size according to network conditions (Tahoe [Jacobson, 1988], Reno [Jacobson, 1990], Vegas [Brakmo & Peterson, 1995]). Both TCP Tahoe and TCP Reno employ an Additive Increase Multiplicative Decrease (AIMD) algorithm for congestion avoidance. TCP Vegas uses a more proactive approach for controlling window size. In this study, we consider TCP Reno, being the most widely implemented version of TCP.

The TCP congestion control mechanism employs four algorithms: slow start, congestion avoidance, fast retransmit, and fast recovery. TCP uses slow start to detect network bandwidth immediately

following connection establishment or timeout. The algorithm starts with a window size w of one or two packets that doubles for each RTT when acknowledgments are received. When w exceeds a threshold, the algorithm enters the congestion avoidance phase. In a typical congestion avoidance phase, the window size w increases linearly by $1/w$ for each returned acknowledgment (additive increase). Thus, after a window of w packets is transmitted, w will be increased by 1 within one RTT. Window size w is reduced to half its current value (multiplicative decrease) if there is an indication of congestion (the receipt of three duplicated acknowledgments of a packet or the presence of an Explicit Congestion Notification (ECN) header information received from a gateway and forwarded in the acknowledgment sent by the receiver). The window then follows the additive increase algorithm. When three duplicated acknowledgments of a packet are received, the TCP sender enters the fast retransmit phase where lost packets are retransmitted without waiting for timeout. The fast recovery algorithm then retransmits all lost packets until a non-duplicate acknowledgment is received. The congestion avoidance phase follows, where the TCP sender attempts to optimize its fair share of the link bandwidth connecting the two TCP end-points. If at least one link on this communication path becomes the bottleneck, an RED gateway regulates the traffic flow in this bottleneck link at its optimum capacity.

2.2 The RED Algorithm

The RED mechanism calculates a weighted moving average x_k of the current queue size q_k as:

$$x_k = (1 - \alpha)x_{k-1} + \alpha q_k, \quad (2)$$

where $\alpha \in (0, 1)$ is the filter resolution. A probability p_k is calculated based on the current value of x_k :

$$p_b = \begin{cases} 0 & 0 \leq x_k < X_{\min} \\ \frac{x_k - X_{\min}}{X_{\max} - X_{\min}} p_{\max} & X_{\min} \leq x_k \leq X_{\max} \\ p_{\max} + \frac{1 - p_{\max}}{X_{\max}} (x_k - X_{\max}) & X_{\max} < x_k \leq 2X_{\max} \\ 1 & 2X_{\max} \leq x_k \leq B \end{cases} \quad (3)$$

$$p_k = \frac{p_b}{1 - c_m p_b}, \quad (4)$$

where X_{\min} and X_{\max} are minimum and maximum thresholds of x_k , respectively; B is the buffer size; p_{\max} is the maximum threshold of p_k ; and c_m is the number of packets that arrived after the last marking. Each incoming packet is marked (or dropped) with probability p_k . Therefore, the *target-queue-length* q_0 can be maintained by appropriately choosing system parameters, such as X_{\min} , X_{\max} , and p_{\max} .

2.3 The TCP-RED System

We consider a system of N TCP flows passing through a bottleneck RED gateway shown in Fig. 1 [Firoiu & Borden, 2000; Ranjan *et al.*, 2004]. We assume that the system operates in its desired stable condition,

satisfying a designed objective [Floyd & Jacobson, 1993]. Without loss of generality, the system is assumed ECN capable. The N long-lived TCP connections are controlled by the ECN markings of the RED gateway, which aims to maintain an even distribution of the bandwidth in the bottleneck link among the N TCP flows.

The average queue size can be maintained at the target queue level $(X_{\min} + X_{\max})/2$. The system can be analyzed using small-signal linearization and perturbation for the calculation of instability boundaries [Chen *et al.*, 2007].

The nonlinear dynamics of the aggregated flow at the RED gateway was analyzed using the *ns-2* network simulator and the parameters shown in Table 1. As the RED filter resolution α varies, the system crosses the boundary of stability. Stable and unstable RED queue length waveforms of the TCP-RED system for two values of α are shown in Fig. 2 [Chen *et al.*, 2007]. The waveforms observed over various time scales exhibit *self-similarity*. In this paper, we analyze its effect on the queue length stability.

2.4 Self-Similarity and DFA Method

Detrended fluctuation analysis (DFA) method [Peng *et al.*, 1995; Hu *et al.*, 2001] is a popular approach to analyze self-similarity and long-range dependence (LRD) properties of signals.

Consider a signal $s(i)$, where $i = 1, \dots, L$, and L is the length of the signal. An integrated time series $y(i)$ of $s(i)$ is obtained as:

$$y(i) = \sum_{j=1}^i (s(j) - \bar{s}), \quad (5)$$

where \bar{s} is the mean of $s(j)$ given as:

$$\bar{s} = \frac{1}{L} \sum_{j=1}^L s(j). \quad (6)$$

The integrated time series $y(i)$ is divided into m bins of equal length l , where $ml = L$. A least-squares fitted line $y_{k_l}(i)$ of $y(i)$, also called the local trend of $y(i)$, is calculated for the k th bin of length l , where $k = 1, \dots, m$. The series $y(i)$ is detrended by subtracting the local trend $y_{k_l}(i)$, for $i = (k-1)l+1, \dots, kl$, as:

$$Y_l(i) = y(i) - y_{k_l}(i). \quad (7)$$

For the k th bin of length l , the root mean square (rms) fluctuation for the integrated and detrended time series is calculated as:

$$f_k(l) = \sqrt{\frac{1}{l} \sum_{j=(k-1)l+1}^{kl} [Y_l(j)]^2}. \quad (8)$$

Thus, the averaged rms fluctuation for the entire time series is given as:

$$F(l) = \sqrt{\frac{1}{m} \sum_{k=1}^m f_k^2(l)}. \quad (9)$$

This computation is repeated for all time scales l to obtain the behavior of $F(l)$. Scale invariant signals with power-law correlations possess a power-law relationship between the rms fluctuation function $F(l)$ and the scale l , given by::

$$F(l) \sim l^\beta. \quad (10)$$

If $\log F(l)$ increases linearly with $\log l$, the slope of the line $\log F(l)$ versus $\log l$ gives the *power-law scaling exponent* β . $F(l)$ is also called the scaling function. We analyze the scale exponent β for the TCP-RED system.

3 Long-Range Power Law Correlations in TCP Flows

3.1 Power-Law Scaling Exponent and Stability

Fig. 3 illustrates the application of the DFA method to the waveforms of the queue length in the TCP-RED system with parameters shown in Table 1. The instantaneous queue length in the RED gateway is recorded every 1 ms over the 5,000 s interval. The power-law scaling exponent (slope of the line) changes when the RED filter resolution α is varied. For each DFA curve, we observe two deflections (crossovers) and three linear regions. We thus obtain three power-law scaling exponents for each DFA curve over the entire range of bin size l , labeled β_1 , β_2 , and β_3 for regions 1, 2, and 3, respectively. A plot of exponent β_1 versus the RED filter resolution α is shown in Fig. 4. Also shown are the stability boundaries for the TCP-RED system. They are obtained by waveform inspection [Chen *et al.*, 2007]. The value of β_1 , which varies between 1.5 and 2.0, provides a clear indication of the stability of the TCP-RED system. Values of β_1 below ~ 1.85 correspond to stable region and above ~ 1.85 correspond to unstable region of operation.

3.2 Interpretation from a Waveform Viewpoint

In the analysis of time series, the DFA exponent value allows a consistency characterization of seemingly nonstationary time series [Hu *et al.*, 2001]. Two time series are *DFA similar* if they have nearly equal DFA exponent values. Hence, if a stationary time series that is *DFA similar* to the RED queue length time series could be identified, the two would have consistent characteristics.

The unstable RED queue length waveform shown in Fig. 2(b) resembles a sinusoidal wave. The DFA curve for a sinusoidal waveform with power-law scaling exponent 2 is shown in Fig. 5. The crossover depends on the frequency of the sinusoidal wave. Also shown in Fig. 5 is the DFA curve for a sawtooth waveform with a power-law scaling exponent 1.5. Such a sawtooth waveform is the native waveform of the stable TCP-RED congestion algorithm. DFA exponent values between 1.5 and 2.0 may be generated by combining “sinusoidal” and “sawtooth” waveforms. For example, a DFA exponent value 1.75 may be obtained by combining “sinusoid” and “sawtooth” with an amplitude ratio of 2:1, as shown in Fig. 5.

4 Long-Range Power-Law Correlations in TCP Window Periods

In Section 3, the FDA method was applied to a time series of the RED queue length aggregated from all TCP sources. Hence, it would be of interest to analyze individual TCP sources. In this Section, we consider the time series that enables probing into the TCP sources. We consider the time series $T_j(i)$, where $j = 1, \dots, 170$ is the flow number and $T_j(i)$ is the time interval between two multiplicative decreases of the j th TCP flow. $T_j(i)$ is, therefore, the TCP window period. Figure 6 shows the DFA scaling exponent of $T_1(i)$ by varying α , the filtering parameter of RED. The results follow a linear relationship, which indicates the feasibility of using the DFA method and the existence of power-law relationship. We fit the traces using a straight line whose slope is equal to the DFA exponent. Note that the scaling behavior for small and large scales are not identical [Hu *et al.*, 2001]. The raw data length ranges roughly from 10 to 10^9 . For the DFA method, we use for all connections a shortened data length, between $10^{2.3}$ and $10^{6.9}$. Two reasons for using reduced data lengths are that the DFA method is less accurate at the two ends of the scaling range and that the range between $10^{2.3}$ and $10^{6.9}$ of the scaled DFA results shows a linear relation on the *log-log* plot, indicating the suitability of applying the DFA method.

To avoid confusion with the DFA exponent β used in Section 3, let β_{T_j} denote the DFA exponent of $T_j(i)$, the time series of the j th TCP window period. Even though β_{T_j} s are not identical, they follow a distribution as shown in Fig. 7 for all α 's.

The value of DFA scaling exponent of the TCP window period for the TCP-RED system is given by $\beta_T = \frac{1}{N} \sum_{j=1}^N (\beta_{T_j})$. As shown in Fig. 8, $\beta_T < 0.5$ represents power-law anti-correlation. The degree of the power-law anti-correlation is related to the stability of the TCP-RED system. Larger power-law anti-correlation (smaller value of β_T) imply more stable system.

Anti-correlation may be observed in the time series of Fig. 9. We also observe the common characteristic of the anti-correlation of a time series: several periods of larger window size are always followed by several periods of smaller window size. Figure 9 can also be explained in the view of feedback control systems. A connection having an increase in the window size, which represents an increase in flow rate, will always trigger a decrease in the window size of another connection to maintain a constant flow rate. However, when the system is unstable, the TCP sending rate becomes less organized and appears memoryless, as indicated by $\beta_T \approx 0.5$.

5 Conclusion

Based on the data collected from the *ns-2* simulations, long-range power-law correlations of both the queue length waveforms at the RED gateway and the TCP window period series from the TCP sources have been analyzed using the DFA method. The results show that the power-law scaling exponent varies with the stability of the RED gateway. The main contribution of this work is the finding that

the scaling exponent varies with the relative stability of the RED gateway. The scaling exponent is independent of the stationarity of the queue length or TCP source periods and, hence, it can be used at the gateway or the TCP source nodes as an indicator for the stability of the TCP-RED system.

References

- Brakmo, L. S. & Peterson, L. L. [1995] "TCP Vegas: end-to-end congestion avoidance on a global Internet," *IEEE J. Selected Areas in Comm.*, **13**(8): 1465–80.
- Buldyrev, S. V., Goldberger, A. L., Havlin, S., Mantegna, R. N., Matsu, M. E., Peng, C. K., Simons, M. & Stanley, H. E. [1995] "Long-range correlation properties of coding and noncoding DNA sequences: GenBank analysis," *Phys. Rev. E* **51**: 5084.
- Carpena, P., Bernaola-Galvan, P., Coronado, A. V., Hackenberg, M. & Oliver, J. L. [2007] "Identifying characteristic scales in the human genome," *Phys. Rev. E* **75**: 032903.
- Chen, Z., Ivanov P. C., Hu, K., & Stanley, H. E. [2002] "Effect of nonstationarities on detrended fluctuation analysis," *Phys. Rev. E* **65**: 041107.
- Chen, Z., Hu, K., Carpena. P., Bernaola-Galvan, P., Stanley, H. E. & Ivanov P. C. [2005] "Effect of nonlinear filters on detrended fluctuation analysis," *Phys. Rev. E* **71**: 011104.
- Chen, X., Wong, S. C., Tse, C. K. & Trajković, Lj. [2007] "Stability analysis of RED gateway with multiple TCP Reno connections," *Proc. Int. Symp. Circuits and Systems*: 1429-1432.
- Chen, X., Wong, S. C., Tse, C. K. & Trajković, Lj. [2008] "Stability study of the TCP-RED system using detrended fluctuation analysis," *Proc. IEEE Int. Symp. Circuits and Systems*: 324–327.
- Defense Advanced Research Projects Agency. [2007] "The network simulator–ns-2," [Online] <http://www.isi.edu/nsnam/ns/>.
- Figueiredo, D. R., Liu, B., Feldmann, A., Misra, V., Towsley, D. & Willinger, W. [2005] "On TCP and self-similar traffic," *Performance Evaluation*, **61**(2-3): 129–141.
- Firoiu, V. & Borden, M. [2000] "A study of active queue management for congestion control," *Proc. IEEE Infocom*, **3**: 1435–1444.
- Floyd, S., Gummadi, R. & Shenker, S. [2001] "Adaptive RED: an algorithm for increasing the robustness of RED's active queue management," [Online] <http://www.icir.org/floyd/red.html>.
- Floyd, S. & Jacobson, V. [1993] "Random early detection gateways for congestion avoidance," *IEEE/ACM Trans. Netw.* **1**(4): 397–413.
- Gao, J. B., Rao, N. S. V., Hu, J. & Ai, J. [2005] "Quasi-periodic route to chaos in the dynamics of Internet transport protocols," *Phys. Rev. Lett.* **94**(19): 198702.
- Heneghan, C. & McDarby, G. [2000] "Establishing the relation between detrended fluctuation analysis and power spectral density analysis for stochastic processes," *Phys. Rev. E* **62**: 6103.

- Hu, K., Ivanov, P. C., Chen, Z., Carpena, P. & Stanley, H. E. [2001] “Effect of trends on detrended fluctuation analysis,” *Phys. Rev. E* **64**(1): 011114.
- Jacobson, V. [1988] “Congestion avoidance and control,” *Proc. Symp. on Communications Architectures and Protocol*,: 314–329.
- Jacobson, V. [1990] “Modified TCP congestion avoidance algorithm,” [Online] <ftp://ftp.ee.lbl.gov/email/vanj.90apr30.txt>.
- Janosi, I. M. & Muller, R. [2005] “Empirical mode decomposition and correlation properties of long daily ozone records,” *Phys. Rev. E* **71**: 056126.
- Kuusela, P., Lassila, P., Virtamo, J. & P. Key, [2001a] “Modeling RED with idealized TCP sources,” *Proc. 9th IFIP Conf. Performance Modeling and Evaluation of ATM & IP Networks*: 155–166.
- Kuusela, P., Lassila, P. & Virtamo, J. [2001b] “Stability of TCP-RED congestion control,” *Proc. ITC-17*: 655–666.
- Masugi, M. [2006] “Recurrence plot-based approach to the analysis of IP-network traffic in terms of assessing nonstationary transitions over time,” *IEEE Trans. Circ. Syst. I, Reg. Papers* **53**(10): 2318–2326.
- Misra, V., Gong, W. B. & Towsley, D. [2000] “Fluid-based analysis of a network of AQM routers supporting TCP flows with an application to RED,” *ACM SIGCOMM Computer Comm. Review*, **30**(4): 151–160.
- Molnár, S., Vidács, A., & Nilson, A. A. [1997] “Bottlenecks on the way towards fractal characterization of network traffic: estimation and interpretation of the Hurst parameter,” *Proc. Int. Conf. on Performance & Management of Complex Communication Networks*.
- Movahed, M. S., Jafari, G. R., Ghasemi, F., Rahvar, S. & Tabar, M. R. R. [2006] “Multifractal detrended fluctuation analysis of sunspot time series,” *J. of Stat. Mech.: Theory and Experiment*, **602**: 3–12.
- Nga, J. H. C., Iu, H. H. C., Ling, B. W. K. & Lam, H. K. [2008] “Analysis and control of bifurcation and chaos in averaged queue length in TCP/RED model,” *Int. J. of Bifur. and Chaos* **18**(8): 2449–2459.
- Paxson, V. & Floyd, S. [1995] “Wide-area traffic: the failure of Poisson modeling”, *IEEE/ACM Trans. on Networking* **3**(3): 226–244.
- Peng, C. K., Buldyrev, S. V., Havlin, S., Simons, M., Stanley, H. E. & Goldberger, A. L. [1994] “Mosaic organization of DNA nucleotides,” *Phys. Rev. E* **49**(2): 1685–1689.
- Peng, C. K., Havlin, S., Stanley, H. E. & Goldberger, A. L. [1995] “Quantification of scaling exponents and crossover phenomena in nonstationary heartbeat time series,” *Chaos* **5**(1): 82–90.
- Ranjan, P., Abed, E. H. & La, R. J. [2004] “Nonlinear instabilities in TCP-RED,” *IEEE/ACM Trans. Netw.* **12**(6): 1079–1092.

- Roughan, M. & Veitch, D. [1999] “Measuring long-range dependence under changing traffic conditions,” *Proc. IEEE Infocom*,: 1513–1521.
- Santhanam, M. S., Bandyopadhyay, J. N. & Angom, D. [2006] “Quantum spectrum as a time series: Fluctuation measures,” *Phys. Rev. E* **73**: 015201.
- Siwy, Z., Ausloos, M. & Ivanova, K. [2002] “Correlation studies of open and closed state fluctuations in an ion channel: analysis of ion current through a large-conductance locust potassium channel,” *Phys. Rev. E* **65**: 031907.
- Teverovsky, V. & Taqqu, M. S. [1997] ”Testing for long-range dependence in the presence of shifting means or a slowly declining trend using a variance type estimator,” *J. of Time Series Analysis*, **18**: 279–304.
- Tinnakornsrisuphap, P. & Makowski, A. M. [2003] “Limit behavior of ECN/RED gateways under a large number of TCP flows,” *Proc. IEEE INFOCOM*: 873–883.

List of Figures

1	A system of N TCP flows, from S_i to D_i , where $i = 1, 2, \dots, N$, passing through a common bottleneck link between G_1 and G_2	12
2	Simulated RED queue length waveforms using <i>ns-2</i> simulator for filter resolutions $\alpha = 0.0001$ and $\alpha = 0.0008$ illustrate stable waveforms (a), (c), and (e) and unstable waveforms (b), (d), and (f), respectively. Figures (c) and (e) are enlarged views of (a) while (d) and (f) are enlarged views of (b). There are 170 TCP connections. Each connection shares a fixed bandwidth of 1.5 Mbps in the bottleneck link.	13
3	DFA scaling functions obtained from RED instantaneous queue length for various choices of α	14
4	DFA scaling exponents in region 1 of Fig. 3 for varying α	15
5	DFA method applied to stationary signals.	16
6	DFA scaling exponent of $T_1(i)$, time series of flow number 1 of the 170 TCP connections, for $\alpha = 0.1, 0.002, 0.0015, 0.001, 0.0008, 0.0006, 0.0005, 0.0004, 0.0001$, and 0.00001 . . .	17
7	DFA results and their distribution for $T_j(i)$, time series of the 170 TCP connections, where j is the connection number, for $\alpha = 0.1$	18
8	DFA scaling exponents of $T_j(i)$, time series for the 170 TCP connections, for varying α . . .	19
9	An illustration of two TCP flows competing for bandwidth in a stable TCP-RED system. When one flow has a high frequency, the other has a low frequency, and vice versa. . . .	20




Fig. 1: A system of N TCP flows, from S_i to D_i , where $i = 1, 2, \dots, N$, passing through a common bottleneck link between G_1 and G_2 .

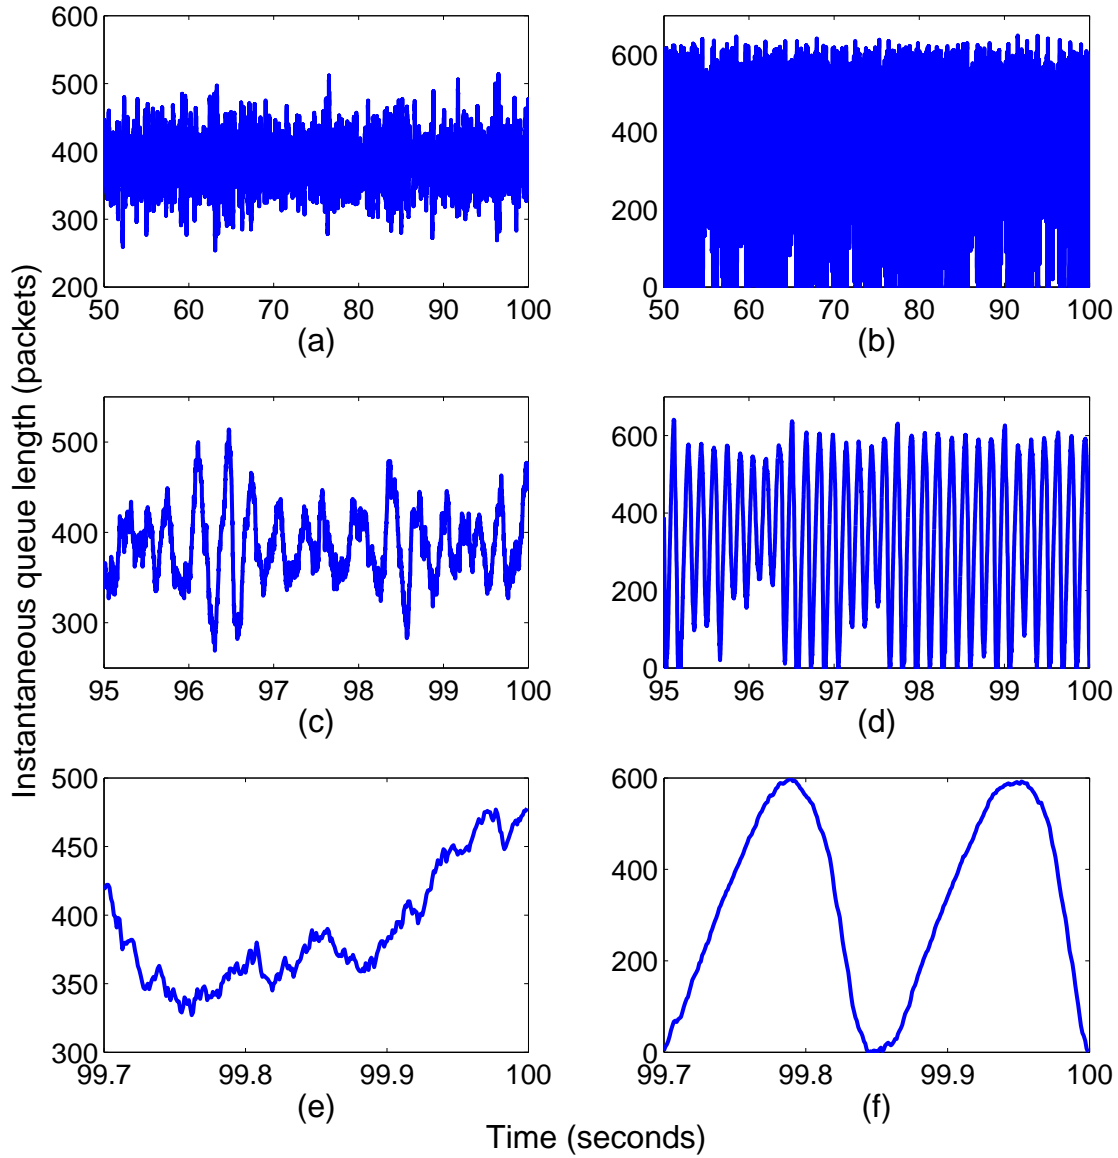


Fig. 2: Simulated RED queue length waveforms using *ns-2* simulator for filter resolutions $\alpha = 0.0001$ and $\alpha = 0.0008$ illustrate stable waveforms (a), (c), and (e) and unstable waveforms (b), (d), and (f), respectively. Figures (c) and (e) are enlarged views of (a) while (d) and (f) are enlarged views of (b). There are 170 TCP connections. Each connection shares a fixed bandwidth of 1.5 Mbps in the bottleneck link.




Fig. 3: DFA scaling functions obtained from RED instantaneous queue length for various choices of α .




Fig. 4: DFA scaling exponents in region 1 of Fig. 3 for varying α .

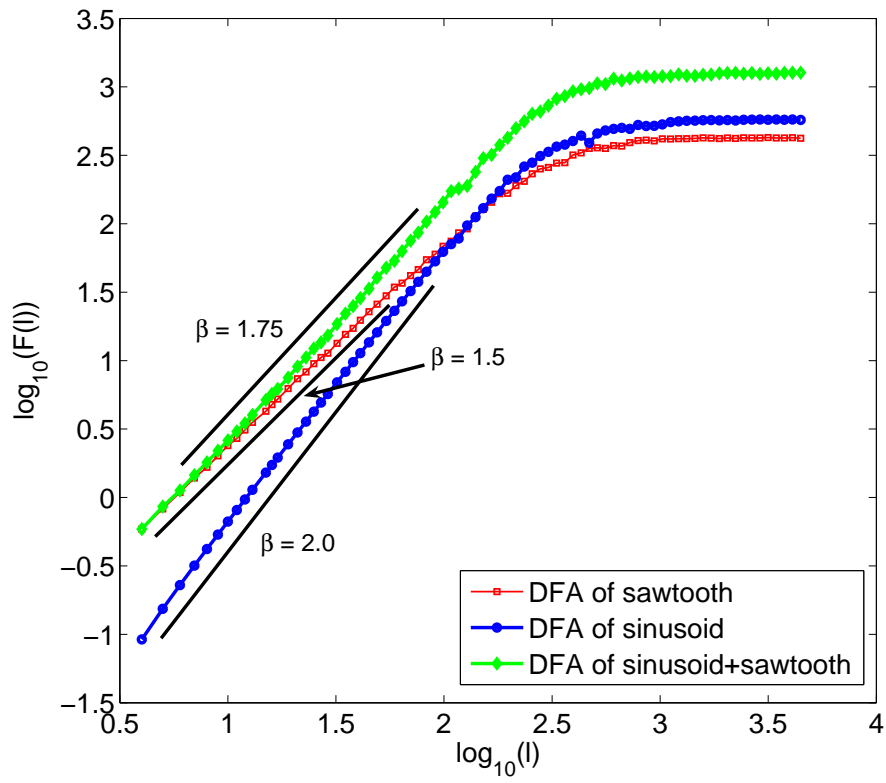


Fig. 5: DFA method applied to stationary signals.




Fig. 6: DFA scaling exponent of $T_1(i)$, time series of flow number 1 of the 170 TCP connections, for $\alpha = 0.1, 0.002, 0.0015, 0.001, 0.0008, 0.0006, 0.0005, 0.0004, 0.0001$, and 0.00001 .

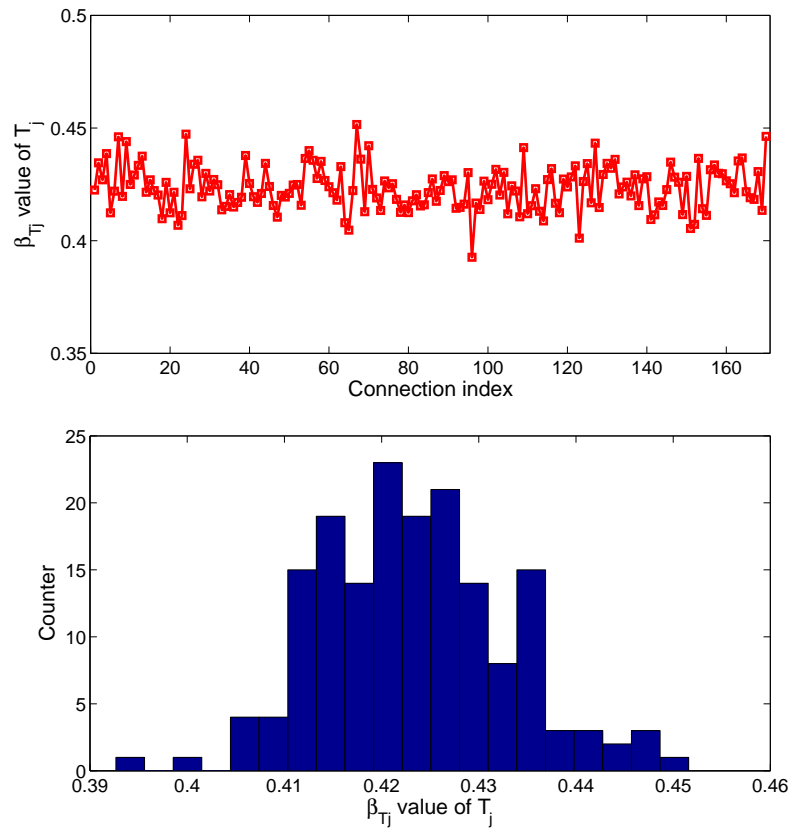


Fig. 7: DFA results and their distribution for $T_j(i)$, time series of the 170 TCP connections, where j is the connection number, for $\alpha = 0.1$.

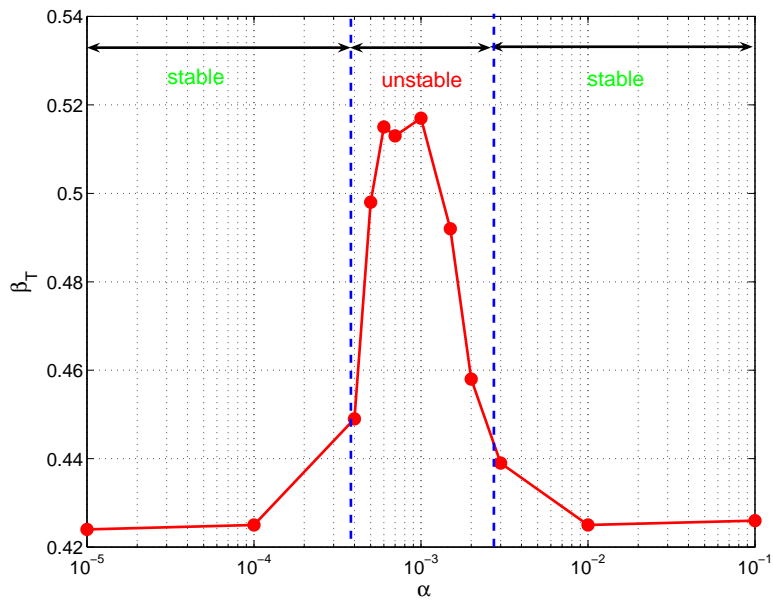


Fig. 8: DFA scaling exponents of $T_j(i)$, time series for the 170 TCP connections, for varying α .




Fig. 9: An illustration of two TCP flows competing for bandwidth in a stable TCP-RED system. When one flow has a high frequency, the other has a low frequency, and vice versa.

List of Tables

1	<i>ns-2</i> Simulation Parameters.	22
---	--	----

Table 1: *ns-2* Simulation Parameters.

<i>ns-2</i> parameter	value
N	170
C/N	1.5 Mbps
R_o	12 ms
r_o	19.46 ms
X_{\min}	256 packets
q_0	384 packets
X_{\max}	512 packets
Bottleneck delay	6 ms
RED redwait	false
TCP ECN	1