



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Performance Evaluation 61 (2005) 129–141

**PERFORMANCE
EVALUATION**
An International
Journal

www.elsevier.com/locate/peva

On TCP and self-similar traffic

Daniel R. Figueiredo^{a,*}, Benyuan Liu^b, Anja Feldmann^c, Vishal Misra^d,
Don Towsley^a, Walter Willinger^e

^a *University of Massachusetts, Department of Computer Science, Amherst, MA, USA*

^b *University of Massachusetts, Department of Computer Science, Lowell, MA, USA*

^c *Technical University of Munich, Computer Science Department, Garching bei München, Germany*

^d *Columbia University, Department of Computer Science, New York, NY, USA*

^e *AT&T Labs—Research, Florham Park, NJ, USA*

Available online 7 January 2005

Abstract

We re-examine the same TCP trace that was used by Veres et al. [A. Veres, M. Boda, The chaotic nature of TCP congestion control, in: Proceedings of the IEEE INFOCOM, 2000] to claim that TCP creates self-similar traffic. A careful reassessment of their data analysis shows that this claim is not justified and suggests that the TCP trace in question is not consistent with (asymptotic second-order) self-similarity or long-range dependence (LRD). We illustrate the reasons that led to the claim in [A. Veres, M. Boda, The chaotic nature of TCP congestion control, in: Proceedings of the IEEE INFOCOM, 2000] and provide some practical guidelines for assessing a statistical characteristic of trace data such as LRD that is defined in strictly asymptotic terms. Our conclusion is in full agreement with the findings obtained from analyzing a much longer TCP trace (resulting from repeating the same simulation as in [A. Veres, M. Boda, The chaotic nature of TCP congestion control, in: Proceedings of the IEEE INFOCOM, 2000], but running it for a longer period) and with analytical results derived from a detailed Markovian model of TCP. These results show that the traffic generated by a long-lived TCP connection, while exhibiting pronounced correlations over a predictable finite range of time-scales, cannot be (asymptotically second-order) self-similar or exhibit LRD. Our work serves as a reminder of the importance of careful trace analysis and detailed examination (and cross-validation) of alternative explanations when establishing or characterizing the generality of any particular finding about Internet traffic.

© 2004 Elsevier B.V. All rights reserved.

Keywords: TCP; Traffic analysis; Long-range dependence

* Corresponding author.

E-mail addresses: ratton@cs.umass.edu (D.R. Figueiredo); bliu@cs.uml.edu (B. Liu); anja@in.tum.de (A. Feldmann); misra@cs.columbia.edu (V. Misra); towsley@cs.umass.edu (D. Towsley); walter@research.att.com (W. Willinger).

0166-5316/\$ – see front matter © 2004 Elsevier B.V. All rights reserved.

doi:10.1016/j.peva.2004.11.004

1. Introduction

In this paper we revisit the claim made in a recent article by Veres and Boda [1], namely that “. . . TCP congestion control creates self-similar traffic . . .” [1, Section VII]. This claim is based on a statistical analysis of traffic generated by an individual TCP flow (so-called TCP microflow) filtered out from an aggregate trace that was collected during the course of a simple NS-2 simulation experiment described in [1]. The analysis of this TCP microflow reported in [1] consists of applying a number of existing methods for testing for self-similarity in a time series and concludes that the traffic generated by such flows exhibits long-range dependence (LRD) or, equivalently, is (asymptotically second-order) self-similar with $H > 0.5$.

While [1] has motivated a number of follow-up studies that have attempted to either explain or corroborate the idea that TCP generates self-similarity, some of these papers have essentially taken for granted the results reported in [1]. For example, [2,3] claim that TCP is capable of inducing traffic that is consistent with self-similarity, while [4] uses a model to show why TCP congestion control mechanisms can lead to self-similar traffic. However, more consistent with the ideas that will be presented here, other papers have supported the hypothesis that TCP, by itself, may not be capable of generating self-similar traffic. For example, [5,6] argue that congestion control mechanisms of TCP can produce strong correlations structure in the traffic pattern, but that such behavior is not consistent with self-similarity.

Three main ingredients form the basis of any statistical analysis: (1) the quality of the measurements used to support the finding; (2) the integrity of the data analysis that allows one to conclude the presence of a particular characteristic; and (3) the scope of possible alternative explanations that were considered in explaining the observed behavior. By revisiting the claim that TCP can generate self-similar traffic, we focus in this paper on the latter two requirements. Concerning (1), since the TCP trace data used in [1] was generated by the widely used NS-2 simulator [7] involving a simple simulation scenario, we will not address the issue of trace quality in this work.¹ We note, though, that the simulation scenario used in [1] is subject to anomalies for much longer simulations as reported in [8]. The original TCP trace data analyzed in [1] was also made publicly available in early February 2001 [9].²

With respect to requirement (2), the integrity of the data analysis, using the same data set considered in [1], we first repeat in Section 2 the statistical analysis presented in [1] and show that there are problems. In particular, we indicate that the evidence in favor of the reported LRD property or asymptotic second-order self-similarity (with $H > 0.5$) is at best inconclusive. Being more constructive, we then illustrate a little-used but generally applicable technique for assessing whether or not a given data set, irrespective of its size, can be viewed as being consistent with statistical features that are defined in terms of asymptotic limits (e.g., LRD). When applying this technique to the TCP trace in question, we find that while LRD cannot be ruled out, there is clear evidence that the data set may not be consistent with LRD (i.e., the 95% confidence intervals of the estimated Hurst parameters consistently contains the value $H = 0.5$). To confirm this suspicion, we repeated the same simulation experiment considered in [1], but executed it for a longer time period to produce a much longer TCP trace. When using the same technique to analyze this larger data set, it is easily seen that the TCP trace is in fact not consistent with LRD or asymptotic second-order self-similarity with $H > 0.5$. Thus, based on the data used in [1], the claim that “TCP generates self-similarity” is premature and when relying on additional data generated from the same simulation

¹ The NS-2 simulation script used in [1] is available at <http://comet.ctr.columbia.edu/~veres/Chaos.html>.

² The data set is available at <http://www.comet.columbia.edu/~veres/DATASETS/Datasets.html>.

used in [1], we find that the measured TCP traces are not consistent with LRD; i.e., by itself, TCP does not generate (asymptotically) self-similar traffic. In this sense, our reassessment of the analysis in [1] argues that the main effort in any statistical analysis should not consist of applying existing techniques but of establishing whether or not the results of these techniques can be trusted for the data at hand, and we provide a concrete and practical means for actively encouraging such efforts.

With a different perspective, but also motivated by [1], the work of [6] is in the spirit of requirement (3) and argues that long-lived flows TCP cannot generate traffic that is consistent with long-range dependence behavior. The authors in [6] use modeling and simulation to show that the internal mechanisms of TCP, namely congestion avoidance and congestion control, are responsible for creating a sustained correlation structure over a finite range of time-scales, and that this range can be analytically predicted and is generally small. Despite its Markovian nature, the model closely matches the correlation structure of simulation traces. Here we will parameterize the model described in [6] to match the behavior of an individual flow in the simulation scenario used in [1]. Relying on the analytical techniques described in [6], we obtain the theoretical correlation structure for the traffic generated by such a TCP session across the whole range of time-scales. We show that the predicted power spectral of the traffic matches the behavior of an individual TCP flow obtained from the simulation scenario of [1]. As we will discuss, both cases exhibit a pronounced correlation structure over a similar range of small time scales, beyond which the correlations becomes negligible.

The rest of the paper is organized as follows. In Section 2, we present a careful statistical analysis of the publicly available TCP trace and of a much longer TCP trace that we generated using the same simulation setup described in [1] and check whether or not the traces are consistent with LRD or asymptotic second-order self-similar. Section 3 presents the analysis of the TCP model, supporting the findings of Section 2. Conclusions are presented in Section 4.

2. Revisiting the data analysis in [1]

In this section we first revisit the analysis performed in [1] and illustrate why some of the reported findings are problematic. We then present and apply a practical approach related to earlier ideas discussed by Mandelbrot [10] and Tukey [11] for investigating whether or not a given time series (irrespective of its length) is consistent with mathematical concepts such as LRD.

2.1. Redoing the analysis in [1]

In the following, we consider the trace `microflow_4h.log` used in [1], representing an individual TCP flow extracted from an aggregate stream of long-lived TCP connections, and consisting of 149,999 observations, where each observation gives the number of bytes sent by the source per 0.1 s. Because of the strange initialization effect in the time series, which was also noted by the authors,³ we ignore the first 10 s of the trace and analyze the modified time series D of length 149,899 consisting of the observations 101, 102, . . . , 149, 999 of the original trace.

We first apply the following well-known methods for checking for LRD to the time series D :

³ See the data description given at <http://www.comet.columbia.edu/~veres/DATASETS/Datasets.html>.

- Absolute moments method (Fig. 1(a)).
- Periodogram method (Fig. 1(b)).
- R/S method (Fig. 1(c)).
- Wavelet-based method (Fig. 1(d)).

For a brief description of these methods and of the corresponding Splus functions code, see <http://math.bu.edu/people/murad/methods/>. For a more detailed overview of the wavelet-based method, and respective Matlab code, see <http://www.emulab.ee.mu.oz.au/~darryl/>. A recently developed and publicly available tool that integrates these methods is described in [12].

A quick glance at [1, Fig. 13] suggests that the time series `microflow_4h.log` is consistent with LRD: each plot suggests the presence of a linear regime, with resulting H -estimates that are consistently around $H \approx 0.8$. However, when comparing our plots shown in Fig. 1 to those depicted in [1, Fig. 13], we can easily identify a number of inconsistencies in the arguments that lead to such a conclusion. In particular, in terms of the available data, since the time series D consists of about 150,000 observations, aggregation-based analysis methods such as the ones listed above should show the results for aggregation

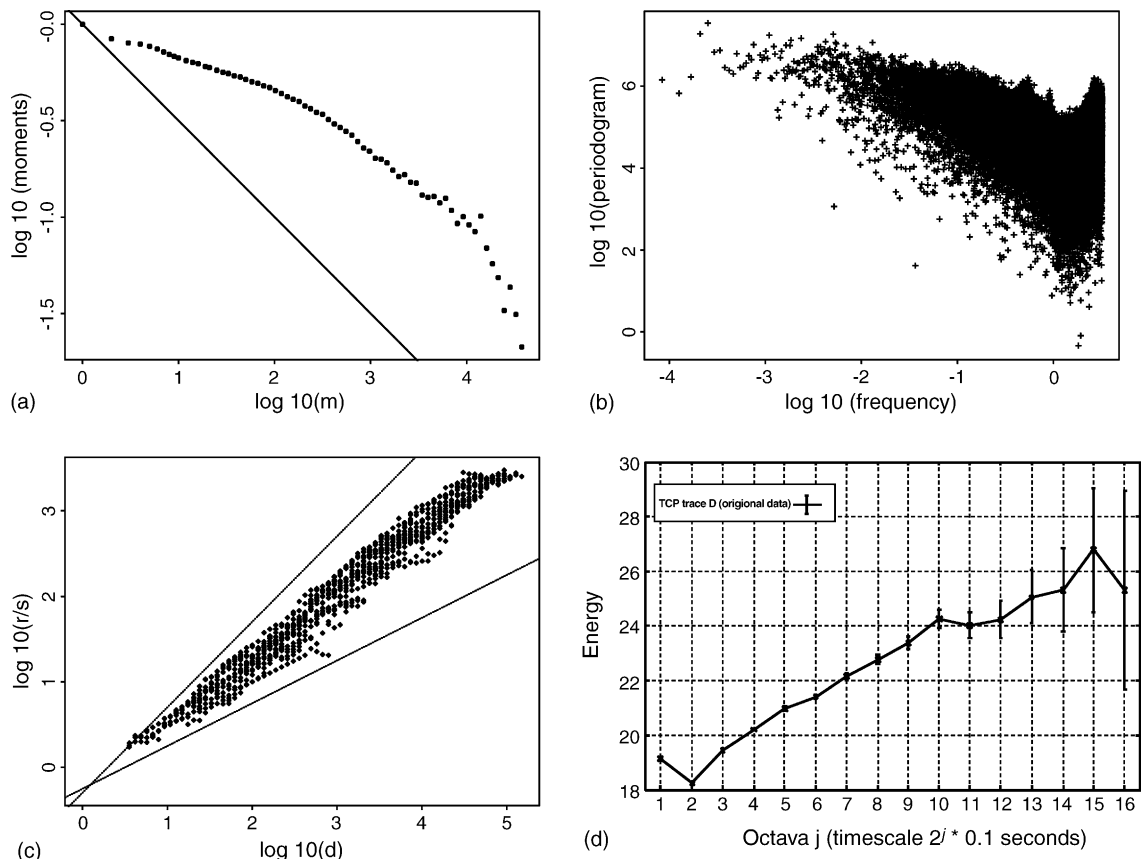


Fig. 1. Applying well-known methods to the time series D : (a) absolute moments method; (b) periodogram method; (c) R/S method; and (d) wavelet-based method.

levels $2^1, 2^2, \dots, 2^{14}, 2^{15}$, up to about 2^{16} (when using base 2), or equivalently (when using base 10), $10^1, 10^2, 10^3, 10^4$, up to about $10^{4.5}$ or so. While plot (c) in [1, Fig. 13] fully exploits the possible aggregation levels, almost half of the large aggregation levels are not shown in plot (a) of that same figure, and some of the large but still informative scales are also absent in plot (d).

The implications for ignoring these large aggregation levels can be serious. Consider for example [1, Fig. 13(a)] which shows the results of applying the *absolute moments method* to the time series `microflow_4h.log`, but only showing the aggregation levels from 2^1 to about $2^9 = 512 \approx 10^{2.7}$. Based on this particular plot, concluding that the underlying time series exhibits LRD with $H \approx 0.8$ (see [1, Section IV]) seems justified. However, as already noted and depicted in Fig. 1(a), when showing the result of this analysis for all aggregation levels up to about $2^{16} \approx 10^{4.5}$, a very different picture emerges. In fact, in view of Fig. 1(a), its counterpart in [1, Fig. 13(a)] focuses exclusively on the transient regime (i.e., from 1 to $10^{2.7}$ or so), even though this regime is largely meaningless when considering an inference technique that is asymptotic in nature and hence has to focus on and exploit the available large aggregation levels (i.e., from about $10^{2.5}$ to $10^{4.2}$ or so). Indeed, by showing the whole range of available aggregation levels, Fig. 1(a) provides evidence that while the underlying time series contains an “interesting” and pronounced short-range correlation structure, the long-range correlations appear to be consistent with those of short-range-dependent processes (e.g. white noise or Markovian models), i.e., with $H \approx 0.5$ (slope of the solid line). Using the latest addition to the tool-kit of inference techniques for LRD behavior, i.e., the *wavelet-based method* (or its close cousin, the *power spectral analysis method*—see Section III) and comparing the results shown in [1, Fig. 13(d)] and Fig. 1(d) supports this observation. Representing another inherently asymptotic inference technique, the relevant information about the data being (in)consistent with LRD behavior cannot be gleaned from the transient region of the wavelet-based energy function (left part of the plots) but has to exploit to the fullest possible degree—within the limits dictated by the data—the large scale behavior of the energy function (right part of the plots). In the case at hand, the emergence of an essentially flat energy function over the larger time-scales (indicating that the underlying data is not consistent with LRD behavior) is already noticeable in [1, Fig. 13(d)] but becomes more pronounced when fully exploiting all of the available larger scales as is done in Fig. 1(d). Somewhat more subtle signs of the same problem (i.e., emphasis on the transient region when the focus should be on the asymptotic regime) can be seen when applying the popular *periodogram method* or the commonly used *R/S method* ([1, Fig. 13(b) and (c)] or Fig. 1(b) and (c), respectively).

2.2. Moving beyond the analysis in [1]

The above-mentioned problem should come as no surprise and is a direct consequence of having to rely on a finite-sized data set when trying to infer a statistical characteristic such as LRD, whose mathematical definition is strictly asymptotic in nature and involves taking a limit of large time scales or aggregation levels. Realizing that statistical inference techniques such as those shown in Fig. 1 (or in [1, Fig. 13]) can only be used to conclude whether or not a given time series (irrespective of its length) is consistent with LRD, showing that it is inconsistent with LRD eliminates LRD from being an appropriate statistical description of the data at hand. On the other hand, demonstrating that the given time series is consistent with LRD does not necessarily rule out alternative statistical descriptions that may have little to do with LRD (e.g., powerful SRD approximations) but—given the data at hand—cannot be distinguished from LRD. In statistical terms, this situation corresponds to cases where the 95% confidence intervals

associated with a reliable Hurst parameter estimate contains the value $H = 0.5$. In view of the discussion in Section 2.2, the question is whether a more systematic exploitation of the available observations is feasible and, if so, whether it can provide more information to determine if the time series D is either consistent or inconsistent with LRD. Although we focus here on the problem of inferring LRD or, equivalently, (asymptotic second-order) self-similarity, the approach discussed below applies more generally to other statistical concepts. For a more detailed discussion of this approach and some illustrative examples, see for example [13,14].

To answer the above question, we focus on the wavelet-based inference method, but the framework outlined below applies to other inference techniques as well. Note that Fig. 1(d) results from taking a strictly “static” view of the available data; that is, we take the time series D as such (all 149,899 observations), apply the wavelet-based method, and plot the resulting energy function in Fig. 1(d). As mentioned earlier, the interpretation of this wavelet plot is somewhat ambiguous, especially when discounting the large scales. However, this ambiguity can be significantly reduced by taking a more “dynamic” view of the data at hand in the sense of Tukey’s notion of “borrowing strength from large data sets” [11] and Mandelbrot’s “sequential (moment) estimate plots” [10]. To illustrate these little-known but powerful concepts, we start with an initial subset D_0 of D of length n_0 and consider successively larger series D_1, D_2, \dots, D_N of length $n_0 < n_1 < \dots < n_N$ with n_N equal or close to the total number of observations of D . For each time series D_i we plot the corresponding wavelet energy function and obtain the corresponding Hurst parameter estimate H_i , together with its 95% confidence interval CI_i . In particular, setting $n_i = 2^{10+i}$, $0 \leq i \leq 6$, Fig. 2(a) shows the resulting family of wavelet energy functions for octaves $j = 1, 2, \dots, 10 + i - 1$. When focusing for each energy function on roughly the top half of the available scales (e.g., for $i = 0$, consider scales or octaves 6–9), the resulting estimates H_i and their corresponding 95% confidence intervals CI_i are depicted in Fig. 2(b).

The main motivation for this dynamic analysis of the time series D is that it supports the careful exploration of the consistency of an assumed model (e.g., LRD model like fractional Gaussian noise) or associated model parameters (e.g., Hurst parameter H) as the number of observations increases. In




Fig. 2. “Dynamic” wavelet-based analysis of time series D : energy function plots for time series D_0, D_1, \dots, D_6 , where D_i is of length $n_i = 2^{10+i}$, $0 \leq i \leq 6$ (a); estimated H_i -values with corresponding 95% confidence intervals CI_i for time series D_i , $0 \leq i \leq 6$ (b).

particular, assuming that one and the same underlying (stationary) process generated the time series D in the first place, increasing the number of observations should have only the following two main effects. First, the H_i -values should stabilize, and in addition, their accuracy (measured in terms of the widths of the CI_i 's) should improve in such a way that as i increases, CI_i should more or less contain CI_{i+1} . In fact, with this dynamic view of D , being inconsistent with SRD (i.e., $H = 0.5$) corresponds to obtaining a roughly nested sequence of CI_i 's that does not include the value 0.5 for all CI_j 's, where $j > i$ for some fixed i . On the other hand, being consistent with LRD does not mean that SRD can be ruled out—as long as the value $H = 0.5$ is ultimately included in the nested sequence of CI_i 's, the data set at hand cannot be used to discriminate between LRD and SRD. Fig. 2(b) is a clear indication that the TCP trace D is not inconsistent with SRD. Note that the relatively large confidence intervals in Fig. 2(b) reflect the fact that our focus is on the large time scales/octaves, where the observed variability in the wavelet energy plots in Fig. 2(a) produces a wide range of viable H -estimates. For details about how the CI_i 's depend on the underlying scales, see [15].

To probe further, we combine this “dynamic wavelet-based method” with the simple idea of running the simulation longer and collecting more data, thereby allowing the statistical analysis to account for larger time-scales. Note that while asking for larger data sets is, in general, a luxury when dealing with measurement experiments in real networks, it is a perfectly valid and often highly effective method for clarifying statistical uncertainties when analyzing trace data from a simulation. To illustrate, running the same simulation scenario as in [1], we modified the ns-2 source code to collect the observations internally, without generating a huge raw trace file. With this modification, we were able to obtain a 672 h-long trace in a couple of hours of simulated time. The “static” wavelet analysis of the time series E representing an individual TCP flow consisting of 24,209,431 observations and extracted from this long trace is shown in Fig. 3 and provides clear evidence that the energy function plot flattens out after scale 10 or so; i.e., the TCP trace is not consistent with LRD. The results of applying the “dynamic” version of the wavelet-based method are even more illuminating and are depicted in Fig. 4. The left plot shows the wavelet energy functions corresponding to selected subsets E_i of E , where E_i contains the first $n_i = 2^{10+i}$ observations of E , $i = 6, 8, 10, 12$ and 14 . To generate the plot to the right, we consider for each ($6 \leq i \leq 14$) time series E_i and the corresponding wavelet energy function and obtain the Hurst parameter estimate H_i , together with its 95% confidence intervals CI_i , by focusing again on roughly the top half of the available scales (e.g., for $i = 6$, we choose scales or octaves 10–15). Clearly, the plots confirm that the energy function of the trace E does begin to flatten out after times-scale 2^{11} or so, just as [1, Fig. 13(d)] hinted at, or as Fig. 1(d) suggested, or as Fig. 2 indicated. While the trace exhibits pronounced correlations over small-to-medium time-scales, it is clearly void of any correlations over time scales beyond 2^{11} or so (about 200 s). This observation is succinctly captured in Fig. 4(b)—after about $i = 6$ or so, the H_i estimates stabilize around the value $H = 0.5$ and the accuracy of that value improves consistently as we consider larger and larger sets E_i . In short, applying the “dynamic” wavelet method in the context of the TCP trace E shows unambiguously that within the setting considered in [1], TCP does not generate self-similar traffic.

There are numerous other problems associated with using the afore-mentioned predominantly graphical techniques for inferring whether a given trace is consistent with LRD behavior or not. Some of these problems have their origins in the presence of strong periodic components or other non-stationary characteristics in the data at hand, and avoiding them typically require tedious and time-consuming “digging” in the data. Fortunately, the simulation scenario considered is simple enough such that the traces considered in the analysis are stationary.

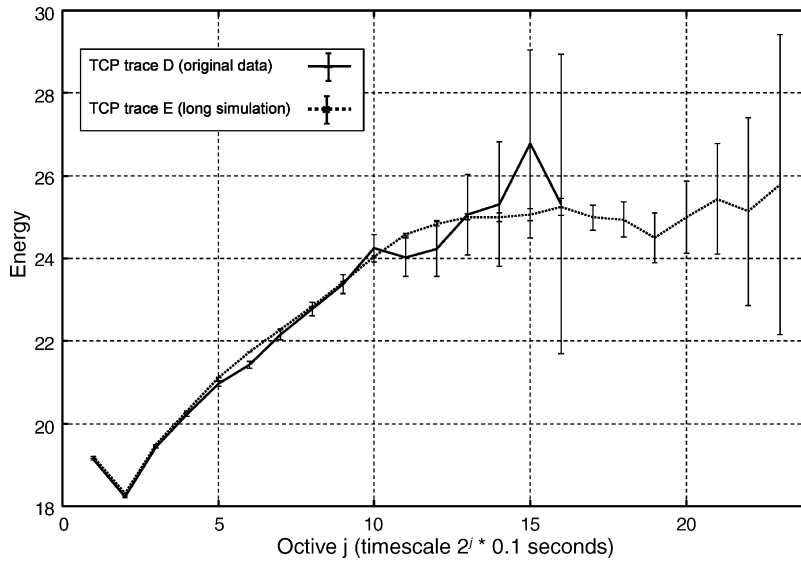


Fig. 3. “Static” wavelet-based analysis of TCP traces *D* and *E*.

3. TCP model and traffic analysis

In order to support our previous findings, which were obtained through analysis of simulation traces, we now focus on an analytical model for TCP and study the statistical characteristics of the traffic generated. Understanding traffic patterns of accurate analytical models for TCP will serve as cross-validation of results obtained through simulation.




Fig. 4. “Dynamic” wavelet-based analysis of TCP trace *E*: energy function plots for time series E_6, E_8, \dots, E_{14} , where E_i is of length $n_i = 2^{10+i}$ (a); estimated H_i -values with corresponding 95% confidence intervals CI_i for time series $E_i, 6 \leq i \leq 14$ (b).

Analytical models have been quite successful in capturing the dynamics of TCP. Among these, several Markovian models have been developed to predict the steady state behavior of TCP [16–19]. Some of these models have been shown to be accurate predictors of TCP behavior over diverse scenarios. This is the case for the model proposed in [16] which captures the congestion avoidance and congestion control behavior of TCP-Reno using a discrete-time Markov model. It was validated using simulation and measurement results and appears adequate to capture the essence of TCP traffic [16,20,21]. In particular, it was shown that this model accurately predicts the TCP throughput and packet sending rate under several packet loss probabilities and round-trip times.

The analysis that follows is based on the Markovian TCP model introduced in [6], which is a slightly modified version of the model proposed in [16]. This detailed model captures the behavior of a single long-lived TCP flow and will be used to understand the statistical characteristics of the flow's data rate. We assume the reader is familiar with TCP-Reno and provide only a brief explanation of the model.

TCP behavior is modeled in terms of “rounds”, where a round represents the transmission of the entire current congestion window. The main assumption is that the round-trip time of the connection is larger than the time required to transmit all packets in the congestion window. This assumption, which is often the case for wide-area networks, supports the notion of rounds. The model also assumes that the packet loss process is independent from round to round. However, if a packet is lost in a round, which occurs with a given probability p , then all subsequent packets in that same window are also lost. This loss process is motivated by the universality of drop-tail queues within the Internet and the observation that the round-trip times are typically larger than the time required to send all packets in any given window, which prevents TCP from reacting within a congestion window. The state of the Markov process is given by the tuple of $(W_i, C_i, L_i, E_i, R_i)$, $i = 0, 1, \dots$, where i is the number of the current round. W_i represents the window size for round i ; C_i helps model the delayed ACK behavior of the TCP receiver, $C_i = 0$ indicates the first of the two rounds and $C_i = 1$ indicates the second; L_i is the number of packets lost in the previous round; E_i denotes whether the connection is in a timeout state and the value of the back-off exponent in round i ; R_i indicates if a packet being sent in the timeout phase (TO) is either a retransmission ($R_i = 1$) or a new packet ($R_i = 0$). The transition matrix of the process is obtained from the behavior of TCP-Reno by examining all possible transitions from a given state. The holding time in each state of the model depends on the outgoing transitions. Time is measured in units of RTT. Let R be the value of the RTT. The model assumes that the value of retransmission timer R_{TO} is a fixed integer multiple of R . Note that p and R_{TO} are the sole parameters of the model.

Associated with each state of the Markov chain is a packet sending rate, which is dependent on the current window size and on the state holding time. Let the sequence S_i , $i = 1, 2, \dots$ be the packet sending rate associated with state $(W_{i-1}, C_{i-1}, L_{i-1}, E_{i-1}, R_{i-1})$. The interested reader is referred to the original work in [16] for more technical details and justifications for the model or to [6] for a description of the modified model (state space and transition probability matrix) and the packet sending rate.

3.1. Model analysis

In order to evaluate the traffic characteristics of the model, we will apply the analytical techniques described and used in [6]. The analysis consists of numerically evaluating the theoretical autocorrelation function for the traffic sending rate associated with the Markov chain and numerically computing the respective power spectral density (PSD). The latter is then mapped into the time domain to allow direct

comparison with results obtained from the wavelet-based method applied to simulation traces. A succinct description of this analysis process follows.

We will assume that the initial state distribution of the Markov model is the steady state probability distribution. Then, the autocorrelation function of the traffic sending rate is given by $\rho(\tau) = E[S_1 S_{\tau+1}]$; $\tau = 0, 1, 2, \dots$. Note that $E[S_1]$ is the expected steady state traffic sending rate of the model. We use the Tangram-II modeling tool [22] to construct the Markov chain and compute the autocorrelation function $\rho(\tau)$. The power spectral density of a discrete time (stationary) stochastic process is defined as the discrete Fourier transform of its autocorrelation function. In our case, the PSD function is given by $\psi(f) = \rho(0) + 2 \sum_{n=1}^{\infty} \cos(2\pi fn) \rho(n)$, since $\rho(\tau)$ is real. The last step is mapping $\psi(f)$ into the time domain, which is obtained through the transformation $\psi^*(t) = \psi((1/t)f_0)$, where f_0 is determined by the sampling rate which in our case is equal to one round-trip time (the time unit of the model).

The motivation for computing the PSD and then mapping it into the time domain is the close relationship between the PSD method and wavelet-based method. The wavelet method consists of estimating the variance of the wavelet coefficients (which result from the discrete wavelet transform) of the underlying time series. These estimates measure the “energy” in the time series at particular time-scales and correspond to estimates of the power spectral density of the time series at frequencies corresponding to those time-scales. Thus, the wavelet-based method (more precisely, the continuous wavelet-transform method, a close relative of discrete wavelet-transform method used here) yields essentially an estimate of the PSD of the process that is assumed to have generated the time series being analyzed.

The discrete wavelet-transform method is commonly used for analyzing the scaling behavior of a time series over different time-scales [15,23]. This method is computationally efficient and is robust against certain types of non-stationarity features that may be present in measured traces (e.g., trends, diurnal cycles) which can cause problems for most of the other known analysis techniques. However, even this more robust method can generate biased results under some conditions (e.g., pronounced periodic components over some time-scales [24]). This potential pitfall shows that cross-validation and further exploration of the data are important to confirm such statistical findings.

3.2. Numerical results

In order to obtain the parameters for the TCP model, namely p and R_{TO} , we simulate the network scenario used in [1] and gather statistics for an individual TCP flow. We use the long simulation trace described earlier to collect these two parameters for the model. The measured average packet loss probability of the long trace was around 0.160, which agrees with the packet loss ratio reported in [1] for their shorter trace. To obtain the measured average R_{TO} value in units of RTT, we collected all TCP estimates for the RTT (round-trip time) and R_{TO} in the course of the simulation. The measured average was 0.163 and 0.515 s, respectively, yielding a ratio of 3 when rounded to the nearest integer. Thus, to compare the analysis of a simulation trace with the TCP model, we use the parameters $p = 0.160$ and $R_{TO} = 3R$. We note that the simulation parameters used in [1] are not consistent with the “round” assumption of the TCP model, since the propagation delay is relatively small compared to transmission delay.

Fig. 5 shows the correlation structure of the traffic sending rate obtained from the TCP model by plotting $\psi^*(t)$ as described earlier. The curve indicates the presence of a sustained correlation structure over a finite range of time-scales, which disappears for larger time-scales, in particular for values above $2^{11}R$. We can observe two distinct regimes in the correlation structure, one regime over a range of lower time-scales (up to 2^4R), and another over a range of larger time-scales (beyond 2^4R). As pointed out in




Fig. 5. Power spectral analysis of the TCP model with parameters equal to the simulated model.

[6], this reflects the interaction of the congestion avoidance and timeout mechanisms, which govern the structure of the traffic correlation in different time-scales.

In [6], the authors present an analytical procedure to predict the largest time-scale over which the model exhibits a sustained correlation. This analysis is based on writing the power spectral density function as a sum over the eigenvalues of the transition probability matrix of the Markov chain. The dominant eigenvalue is then used to relate the boundary of the time-scale range to the pole of the PSD. Applying Eq. (4) of [6], we obtain the value $2^{10.03} R$ (170.5 s) for the highest time-scale associated with the model. This value is marked with a cross in the plot of Fig. 5. From the figure, we observe that the predicted time-scale is relatively accurate and falls on the knee of the plotted curve.

Recall Fig. 3 which is comparable to Fig. 5 and shows the “static” wavelet analysis of the long TCP trace data obtained through simulation. As discussed, we can establish a direct comparison between the results obtained from the two methods (with the caveat that the time-scales are in units of 0.1 and 0.163 s, respectively). We observe that the two curves present a very similar range of time-scales under which TCP traffic exhibits sustained correlation (up to around 200 s). Although the shape of the two curves do not exactly match for lower time-scales, the model clearly captures the characteristics of the correlation structure of the traffic sending rate at larger time-scales, and in particular the knee in the curve at around 200 s. Since Markovian models cannot generate LRD traffic, the close correspondence between model and simulation results is evidence that the latter is not consistent with LRD, and can be more accurately described and explained in terms of a strictly short-range-dependent model of TCP.

Finally, the results depicted in Fig. 5 are consistent with the results of Section 2 and indicate that within the experimental setting described in [1] or the mathematical modeling framework considered in this section, TCP cannot generate LRD traffic. At best, it can be claimed that under certain conditions, TCP can generate a sustained correlation structure over some initial but finite range of time-scales. Moreover, the largest time-scale in this range can be predicted analytically with reasonable accuracy.

4. Conclusion

This paper considers the question whether or not a single long-lived TCP flow can generate traffic that is consistent with (asymptotically) self-similar behavior, or that exhibits long-range dependence. Over the past couple of years, the research community has been intrigued by the possibility that TCP—on its own (i.e., without any variability of TCP connection sizes)—could give rise to self-similar traffic, such as claimed in [1]. To this end, we revisited and expanded upon the analysis of a publicly available traffic trace that was used in [1] to support the original claims of TCP and self-similarity. We also performed an analysis of a previously developed detailed Markovian model of TCP with parameters that match the scenario encountered in the simulation. Both of our methodologies conclusively demonstrate that, while possessing a sustained correlation structure over a finite range of time-scales, traffic generated by a long-lived TCP flow is *not* consistent with (asymptotically) self-similar behavior or LRD. In fact, the analysis of the Markovian model can with reasonable accuracy predict the time-scale below which TCP gives rise to sustained correlations and above which a long-lived TCP flow exhibits necessarily a trivial (i.e., white noise-type) correlation structure.

Besides answering the question about TCP and self-similarity, we have also addressed the larger issue of adequately interpreting the results of a statistical analysis of trace data. In particular, we argue here that the mere appearance of a linear region in a double-log plot (as most popular statistical tests for self-similarity are) should not lead to an premature conclusion of self-similarity or long-range dependence. LRD is an asymptotic property, and while any inference method for LRD can only be used to conclude whether or not the time series at hand is consistent with LRD (irrespective of the size of the time series), every effort should be made to ensure that the analysis exploits the available data sets to their fullest possible extent. In support of this effort, we describe in this paper a “dynamic” version of the wavelet-based method that applies to other inference techniques as well and provides an efficient mean for checking in a systematic manner whether or not a given time series is consistent with LRD or some other asymptotically defined concept. We also demonstrate that coupling such a systematic analysis of trace data with an appropriate level of abstract mathematical modeling is beneficial and provides for a natural form of cross-validation of proposed explanations for any particular finding.

Acknowledgements

Research of DRF, BL, DT have been supported in part by NSF under awards EIA-0080119, ANI-0070067 and ANI-0085848 and in part by DARPA under contract DOD F30602-00-0554. Research of VM has been supported in part by NSF under award ANI-0238299 and in part by DOE under award DE-FG02-02ER25528. Research of DRF funded in part by a scholarship from CAPES (Brazil).

References

- [1] A. Veres, M. Boda, The chaotic nature of TCP congestion control, in: Proceedings of the IEEE INFOCOM, 2000.
- [2] W. Feng, P. Tinnakornsriruphap, The failure of TCP in high-performance computational grids, in: Proceedings of the ACM/IEEE Supercomputing 2000: Conference on High Performance Networking and Computing, 2000.
- [3] N. Wisitpongphan, J. Peha, Effect of TCP on self-similarity of network traffic, in: Proceedings of the 12th International Conference on Computer Communications and Networks, 2003.

- [4] B. Sikdar, K.S. Vastola, The effect of TCP on the self-similarity of network traffic, in: Proceedings of the Conference on Information Science and Systems, The John Hopkins University, Baltimore, MD, 2001.
- [5] L. Guo, M. Crovella, I. Matta, How does TCP generate pseudo-self-similarity? in: Proceedings of the International Workshop on Modeling, Analysis and Simulation of Computer and Telecommunications Systems (MASCOTS), 2001; addendum to published paper in Technical Report BU-CS-2001-026.
- [6] D.R. Figueiredo, B. Liu, V. Misra, D. Towsley, On the autocorrelation structure of TCP traffic, *Comput. Networks* 40 (3) (2002) 339–361.
- [7] LBL, Xerox PARC, UCB, USC/ISI, ACIRI, Network simulator—NS (version 2). URL: <http://www.isi.edu/nsnam/ns>.
- [8] D.R. Figueiredo, B. Liu, D. Towsley, On the oscillatory behavior of TCP sessions in ns-2 simulator, in preparation.
- [9] end2end-interest mailing list archives. <http://www.postel.org/mailman/listinfo/end2end-interest>.
- [10] B.B. Mandelbrot, *Fractals and Scaling in Finance*, Springer, Berlin, 1997.
- [11] J.W. Tukey, *Exploratory Data Analysis*, Addison-Wesley, Reading, MA, 1977.
- [12] T. Karagiannis, M. Faloutsos, M. Molle, A user-friendly self-similarity analysis tool, *ACM SIGCOMM Comput. Commun. Rev.* 33 (3) (2003) 81–93.
- [13] W. Willinger, D. Alderson, L. Li, A pragmatic approach to dealing with high-variability in network measurements, in: Proceedings of the ACM/SIGCOMM Internet Measurement Conference (IMC'04), Taormina, Italy, 2004.
- [14] W. Willinger, D. Alderson, J.C. Doyle, L. Li, More “normal” than normal: scaling distributions and complex systems, in: R.G. Ingalls, M.D. Rossetti, J.S. Smith, B.A. Peters (Eds.), *Proceedings of the 2004 Winter Simulation Conference (WSC'04)*, Washington, DC, 2004.
- [15] P. Abry, D. Veitch, Wavelet analysis of long-range-dependent traffic, *IEEE Trans. Inform. Theory* 44 (1) (1998) 2–15.
- [16] J. Padhye, V. Firoiu, D. Towsley, A stochastic model of TCP Reno congestion avoidance and control, Tech. Rep. 99-02, Department of Computer Science, University of Massachusetts, Amherst, MA, February 1999.
- [17] E. Altman, K. Avrachenkov, C. Barakat, A stochastic model of TCP/IP with stationary random losses, in: Proceedings of the ACM SIGCOMM, 2000, pp. 231–242.
- [18] M.A. Marsan, E. de Souza e Silva, R. Lo Cigno, M. Meo, A Markovian model for TCP over ATM, *Telecommun. Syst. J.* 12 (4) (1999) 341–368.
- [19] C. Casetti, M. Meo, A new approach to model the stationary behavior of TCP connections, in: Proceedings of the IEEE INFOCOM, 2000, pp. 367–375.
- [20] M. Handley, S. Floyd, TCP-friendly simulations. <http://www.aciri.org/mjh/results.ps.gz>.
- [21] J. Bolliger, T. Gross, U. Hengartner, Bandwidth modelling for network-aware applications, in: Proceedings of the IEEE INFOCOM, 1999.
- [22] R. Leão, E. de Souza e Silva, S. de Lucena, A set of tools for traffic modeling, analysis and experimentation, in: Proceedings of the 11th International Conference on Modelling Tools and Techniques for Computer and Communication System Performance Evaluation (TOOLS'2000), LNCS 1786, Springer, Schaumburg, USA, 2000, pp. 40–55.
- [23] D. Veitch, P. Abry, A wavelet based joint estimator of the parameters of long-range dependence, *IEEE Trans. Inform. Theory* 45 (3) (1999) 878–897.
- [24] P. Huang, A. Feldmann, W. Willinger, A non-intrusive, wavelet-based approach to detecting network performance problems, in: Proceedings of the First ACM SIGCOMM Internet Measurements Workshop, San Francisco, CA, 2001, pp. 213–227.