

University of California

Postprints

Year 2004

Paper 1267

Long-range dependence - Ten years of Internet traffic modeling

T Karagiannis

M Molle

M Faloutsos


T Karagiannis, M Molle, and M Faloutsos, "Long-range dependence - Ten years of Internet traffic modeling" (2004). IEEE Internet Computing. 8 (5), pp. 57-64. Postprint available free at: <http://repositories.cdlib.org/postprints/1267>

Posted at the eScholarship Repository, University of California.
<http://repositories.cdlib.org/postprints/1267>

Long-range dependence - Ten years of Internet traffic modeling

Abstract

Self-similarity and scaling phenomena have dominated Internet traffic analysis for the past decade. With the identification of long-range dependence (LRD) in network traffic, the research community has undergone a mental shift from Poisson and memory-less processes to LRD and bursty processes. Despite its widespread use, though, LRD analysis is hindered by our difficulty in actually identifying dependence and estimating its parameters unambiguously. The authors outline LRD findings in network traffic and explore the current lack of accuracy and robustness in LRD estimation. In addition, the authors present recent evidence that packet arrivals appear to be in agreement with the Poisson assumption in the Internet core.



Long-Range Dependence

Ten Years of Internet Traffic Modeling

Self-similarity and scaling phenomena have dominated Internet traffic analysis for the past decade. With the identification of long-range dependence (LRD) in network traffic, the research community has undergone a mental shift from Poisson and memory-less processes to LRD and bursty processes. Despite its widespread use, though, LRD analysis is hindered by our difficulty in actually identifying dependence and estimating its parameters unambiguously. The authors outline LRD findings in network traffic and explore the current lack of accuracy and robustness in LRD estimation. In addition, the authors present recent evidence that packet arrivals appear to be in agreement with the Poisson assumption in the Internet core.

**Thomas Karagiannis,
Mart Molle,
and Michalis Faloutsos**
University of California, Riverside

Traffic modeling and analysis is a fundamental building block of Internet engineering and design. We can't replicate the Internet and study it as a whole, so we rely on thorough analysis of network measurements and their transformation into models to help explain the Internet's functionality and improve its performance.

About 10 years ago, the introduction of long-range dependence (LRD) and self-similarity revolutionized our understanding of network traffic. (LRD means that the behavior of a time-dependent process shows statistically significant correlations across large time scales; self-

similarity describes the phenomenon in which the behavior of a process is preserved irrespective of scaling in space or time.) Prior to that, researchers considered Poisson processes (that is, the packet arrival process is memory-less and interarrival times follow the exponential distribution) to be an adequate representation for network traffic in real systems.¹ LRD flew in the face of conventional wisdom by stating that network traffic exhibits long-term memory (its behavior across widely separated times is correlated). This assertion challenged the validity of the Poisson assumption and shifted the community's focus from

assuming memory-less and smooth behavior to long memory and bursty behavior.

In this article, we provide an overview of what the community has learned from 10 years of LRD research; we also identify the caveats and limitations of our ability to detect LRD. In particular, we want to raise awareness on two issues: that identifying and estimating LRD is far from straightforward, and that the large-scale aggregation of the Internet's core might have shifted packet-level behavior toward being a Poisson process. Ultimately, measuring and modeling the Internet requires us to constantly reinvent models and methods.

Self-Similarity in Internet Traffic

Ample evidence collected over the past decade suggests the existence of LRD, self-similarity and heavy-tailed distributions (meaning large values can exist with non-negligible probability) in various aspects of network behavior.

Before we look at the major advances in LRD research, we must first describe LRD and self-similarity in the context of time-series analysis.

Stochastic Time Series

Let $X(t)$ be a stochastic process. In some cases, X can take the form of a discrete time series $\{X_t\}$, $t = 0, 1, \dots, N$, either through periodic sampling or by averaging its value across a series of fixed-length intervals. We say that $X(t)$ is stationary if its joint distribution across a collection of times t_1, \dots, t_N is invariant to time shifting. Thus, we can characterize the dependence between the process's values at different times by evaluating the process's *autocorrelation function* (ACF), which is $\rho(k)$. The ACF measures similarity between a series X_t and a shifted version of itself X_{t+k} :

$$\rho(k) = \frac{E[(X_t - \mu)(X_{t+k} - \mu)]}{\sigma^2}, \quad (1)$$

where μ and σ are the mean and standard deviations, respectively, for X .

Also of interest is a time series' aggregated process $X_k^{(m)}$:

$$X_k^{(m)} = \frac{1}{m} \sum_{i=km}^{(k+1)m-1} X_i, \quad k = 0, 1, 2, \dots, \left\lfloor \frac{N}{m} \right\rfloor - 1. \quad (2)$$

Intuitively, $\{X_k^{(m)}\}$ describes the average value of the time series across "windows" of m consecutive values from the original time series. If $\{X_k^{(m)}\}$ were independent and identically distributed, then $\text{Var}(X^{(m)}) = \sigma^2/m$. However, if the sequence exhibits

long memory, then the aggregated process's variance converges to zero at a much slower rate than $1/m$.²

Self-Similarity and LRD

A stationary process X is long-range dependent if its autocorrelations decay to zero so slowly that their sum doesn't converge – that is, $\sum_{k=1}^{\infty} |\rho(k)| = \infty$. Intuitively, memory is built-in to the process because the dependence among an LRD process's widely separated values is significant, even across large time shifts.

A stochastic process X is self-similar if

$$X(at) = a^H X(t), \quad a > 0,$$

where the equality refers to equality in distributions, a is a scaling factor, and the self-similarity parameter H is called the *Hurst exponent*. Intuitively, self-similarity describes the phenomenon in which certain process properties are preserved irrespective of scaling in space or time.

Second-order self-similarity describes the property that a time series' correlation structure (ACF) is preserved irrespective of time aggregation. Simply put, a second-order self-similar time series' ACF is the same for either coarse or fine time scales. A stationary process X_t is second-order self-similar³ if

$$\rho(k) = 1/2 [(k+1)^{2H} - 2k^{2H} + (k-1)^{2H}], \quad 0.5 < H < 1 \quad (3)$$

and asymptotically exactly self-similar if

$$\lim_{k \rightarrow \infty} \rho(k) = 1/2 [(k+1)^{2H} - 2k^{2H} + (k-1)^{2H}], \quad 0.5 < H < 1.$$

Second-order self-similar processes are characterized by a hyperbolically decaying ACF and used extensively to model LRD processes. Conversely, quickly decaying correlations characterize short-range dependence. From these definitions, we can infer that LRD characterizes a time series if $0.5 < H < 1$. As $H \rightarrow 1$, the dependence is stronger.

For network-measurement processes, X refers to the number of packets and bytes at consecutive time intervals, meaning that X describes the volume of bytes/packets observed in a link every time interval t .

Self-Similarity in Internet Traffic

Leland and colleagues' pioneering work provided the first empirical evidence of self-similar charac-

teristics in LAN traffic.⁴ They performed a rigorous statistical analysis of Ethernet traffic measurements and established its self-similar nature. Specifically, they observed that Internet traffic variability was invariant to the observed time scale – that is, traffic didn't become smooth with aggregation as fast as the Poisson traffic model indicated. Subsequently, Paxson and Floyd described the failure of using Poisson modeling in wide-area Internet traffic.⁵ They demonstrated that packet interarrival times for Telnet and FTP traffic were described by heavy-tailed distributions and characterized by burstiness, which indicated that the Poisson process underestimated both burstiness and variability. In addition, they proved that large-scale correlations characterized wide-area traffic traces, concluding, “We should abandon Poisson-based modeling of wide-area traffic for all but user session arrivals.”

These two landmark studies nudged researchers away from traditional Poisson modeling and independence assumptions, which were discarded as unrealistic and overly simplistic. The nature of the congestion produced from self-similar network traffic models had a considerable impact on queuing performance,⁶ due in large part to variability across various time scales. Further studies proved that Poisson-based models significantly underestimated performance measures, showing that self-similarity resulted in performance degradation by drastically increasing queuing delay and packet loss.⁷

Self-similarity's origins in Internet traffic are mainly attributed to heavy-tailed distributions of file sizes.^{8,9} Several studies correlated the Hurst exponent with heavy-tailed distributions, indicating that extremely large transfer requests could occur with non-negligible probability.

Apart from LRD, Internet traffic presents complex scaling and multifractal characteristics. Many simulations and empirical studies illustrate how scaling behavior and the intensity of the observed dependence is related to the scale of observation. Specifically, loose versus strong dependence exists in smaller versus larger time scales, respectively. The change point is usually associated with either the round-trip time (RTT) or intrusive “fast” flows with small interarrival times.^{10,11}

Despite the overwhelming evidence of LRD's presence in Internet traffic, a few findings indicate that Poisson models and independence could still be applicable as the number of sources increases in fast backbone links that carry vast numbers of distinct flows, leading to large vol-

umes of traffic multiplexing.¹² In addition, other studies¹³ point out that several end-to-end network properties seem to agree with the independence assumptions in the presence of nonstationarity (that is, statistical properties vary with time).

LRD Estimation and Its Limitations

The predominant way to quantify LRD is through the Hurst exponent, which is a scalar, but calculating this exponent isn't straightforward. First, it can't be calculated definitively, only estimated. Second, although we can use several different methods to estimate the Hurst exponent, they often produce conflicting results, and it's not clear which provides the most accurate estimation.

We can classify Hurst exponent estimators into two general categories: those operating in the time domain and those operating in the frequency or wavelet domain. Due to space constraints, we can't give a complete description of all available estimators, but an overview appears elsewhere.¹⁴

Time-domain estimators investigate the power-law relationship between a specific statistical property in a time series and the time-aggregation block size m : LRD exists if the specific property versus m is a straight line when plotted in log-log scale. This line's slope is an estimate of the Hurst exponent, so time-domain estimators imply two presuppositions for LRD to exist: statistically significant evidence that the relevant points do indeed represent a straight line, and the line's slope is such that $0.5 < H < 1$ (the Hurst exponent H depends on this slope). These estimators use several methodologies: R/S (rescaled range statistic), absolute value, variance, and variance of residuals.

Naturally, frequency- and wavelet-domain estimators operate in the frequency or wavelet domain. Similarly to the time-domain methodologies, they examine if a time series' spectrum or energy follows power-law behavior. These estimators include the Periodogram, the Whittle, and the wavelet Abry-Veitch (AV) estimators.¹⁵

We can test these estimation methodologies' capabilities by first examining their accuracy on synthesized LRD series and then testing their ability to discriminate LRD behavior when applied to non-LRD data sets. In agreement with similar findings in earlier studies,^{14,16} our findings demonstrate that no consistent estimator is robust in every case: estimators can hide LRD or report it erroneously. Furthermore, each estimator has

different strengths and limitations. We used the software package SELFIS (publicly distributed at our Web site, www.cs.ucr.edu/~tkarag) to perform the experiments described next.

Estimator Accuracy on Synthesized LRD Series

The most extensively used self-similar processes for simulating LRD are fractional Gaussian noise (fGn) and fractional Auto Regressive Integrated Moving Average (ARIMA) processes. fGn is an increment of fractional Brownian motion (fBm) (a random walk process with dependent increments); fGn is a Gaussian process and its ACF is given by Equation 3. The fractional ARIMA(p,d,q) model is a fractional integration of the *autoregressive moving average*, or ARMA(p,q), model. Fractional ARIMA processes describe LRD series when $0 < d < 0.5$, in which $H = d + 0.5$.

We tested each estimator against two different types of synthesized long-memory series: fractional ARIMA and fGn.¹⁷ For each Hurst value between 0.5 and 1 (using a step of 0.1), we generated 100 fGn and 100 fractional ARIMA synthesized data sets of 64 Kbytes. Figure 1 reports the average estimated Hurst value for these data sets for each estimator as well as the 95 percent confidence intervals of the mean (that is, the range of values that has a high probability of containing the mean). However, these intervals are so close to the average that they're barely discernible. Although many estimators and generators exist, we used and evaluated the most common and widely used ones.

Figure 1 shows significant variation in the estimated Hurst exponent value between the various methodologies, especially as the Hurst exponent tends to 1, where the intensity of long-range dependence is larger. Frequency-domain estimators seem to be more accurate. In the case of the fGn synthesized series, Whittle and Periodogram estimators fall exactly on top of the optimal estimation line. The Whittle estimator has the a priori advantage of being applied to a time series whose underlying structure matches the assumptions under which the estimator was derived. The wavelet AV estimator always overestimates the Hurst exponent's value (usually by 0.05). Overall, time-domain estimators fail to report the correct Hurst exponent value, underestimating it by more than 20 percent. (In Figure 1, lines clustered under the optimal estimation line represent these estimators.) When we used

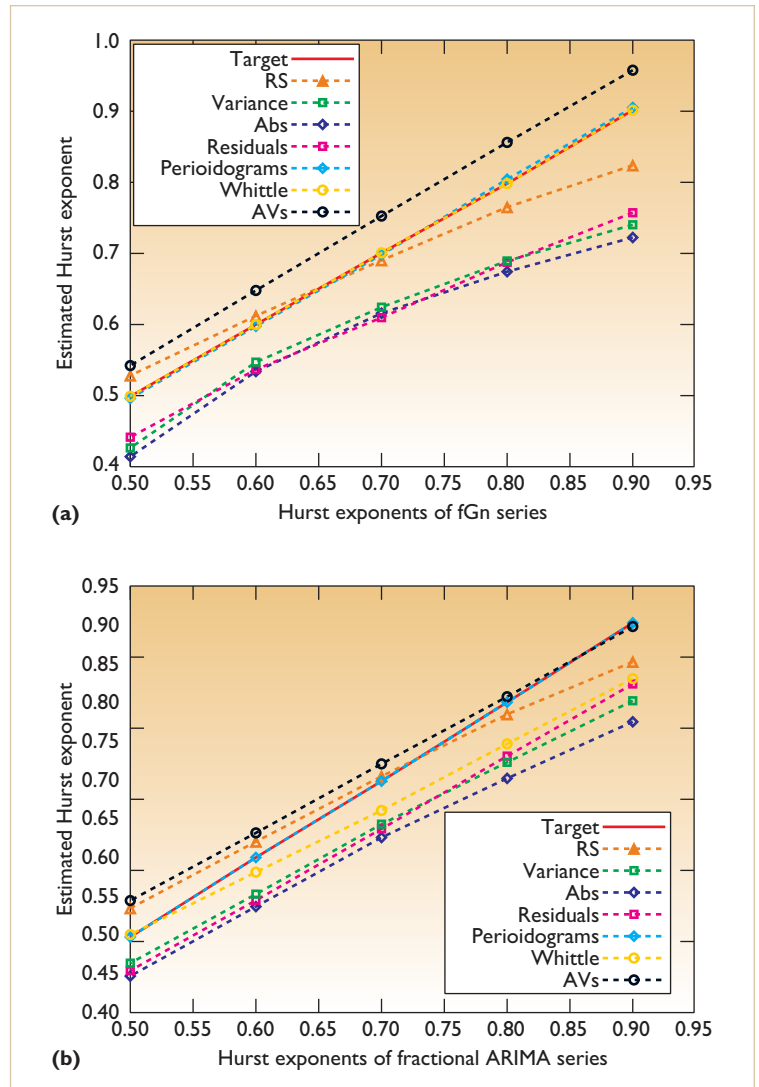


Figure 1. Estimating Hurst exponent values. We tested the performance of estimators on (a) fractional Gaussian noise (fGn) estimator and (b) fractional ARIMA (Auto Regressive Integrated Moving Average) synthesized time series. The target line is the optimal estimation. In both cases, time-domain estimators, represented by the lines clustered below the target line, failed to capture the synthesized Hurst exponent value, especially as H tended to 1. Frequency-based estimators appear to be more accurate, following the target line closer.

fractional ARIMA to synthesize the time series, the estimations were generally closer to the optimal estimation line. However, none of the estimators consistently followed the optimal line across all Hurst values.

Discrimination of LRD Behavior in Deterministic Series

To study the estimations' sensitivity, we examined

the effects of various phenomena common to time-series analysis, such as periodicity, noise, and trend (where the mean of the process is steadily increasing or decreasing). Our analysis revealed that the presence of such processes significantly affects estimators. Furthermore, most methodologies fail to distinguish between LRD and such phenomena, and falsely report LRD in deterministic non-LRD time series. We examined four cases and learned that, essentially, no estimator is consistently robust in every case. Each one evaluates different statistics to estimate the Hurst exponent, which requires the examination of many estimators to get an overall picture of the time series' properties. Applying signal-processing techniques and methodologies could help us overcome some of these limitations, but networking practitioners aren't necessarily familiar with such practices.

Cosine plus white Gaussian noise. In our first test, we applied the estimators to periodic data sets and then synthesized the series with white Gaussian noise and a cosine function: $A\cos(\alpha x)$. Periodicity can mislead the Whittle, Periodogram, R/S, and AV methods into falsely reporting LRD. The Hurst exponent estimation depends mainly on A , so the estimations approach 1 as A increases. Thus, as the amplitude increases, estimations become less reliable. If the amplitude is large and the period is small, Whittle always estimates the Hurst exponent to be 0.99. (Whittle estimates of 0.99 represent the failure of robust estimation.)

fGn series plus white Gaussian noise. We next examined the effect of noise on LRD data. We found that all estimators underestimate the Hurst exponent in the presence of noise, but with the exception of Whittle and the wavelet estimator, the difference is negligible. Depending on the signal-to-noise ratio and the fGn series' Hurst exponent value, however, these two estimators could significantly underestimate the Hurst exponent – by more than 20 percent in some cases.

fGn series plus a cosine function. In studying the effect of periodicity on LRD data, we found that all estimations were affected by its presence. Depending on the cosine function's amplitude, time-domain estimators tend to underestimate the Hurst exponent. On the other hand, frequency-based methodologies overestimate the Hurst exponent. As we increase the cosine function's amplitude, estimates tend toward 1.

Trend. The definition of LRD assumes stationary time series. To study the impact of nonstationarity on the estimators, we therefore synthesized various series with different decaying or increasing trends. We also examined combinations of previous categories (white Gaussian noise and cosine functions) with trend. In every case, the Whittle estimate was consistently 0.99; the Periodogram method's estimates for the Hurst exponent were greater than 1, whereas self-similarity is only defined for $H < 1$. No other methodology produced statistically significant estimations.

Examining the Poisson Assumption in the Backbone

We studied the Poisson assumption's validity on several OC48 (2.5 Gbps) backbone traces taken from CAIDA (Cooperative Association for Internet Data Analysis) monitors located at two different SNET OC48 links belonging to two US tier-1 Internet service providers (ISPs).

To capture the traces, we used Linux-based monitors with Dag 4.11 network cards and packet-capture software originally developed at the University of Waikato and currently produced by Endace. We analyzed various backbone traces: August 2002 (backbone 1, eight hours), January 2003 (backbone 1, one hour), April 2003 (backbone 1, eight hours), May 2003 (backbone 1, 48 hours; backbone 2, two hours), and January 2004 (backbone 2, one hour).

Our analysis demonstrates that backbone packet arrivals appear to agree with the Poisson assumption,^{12,18} but our traces also appear to agree with self-similarity and past LRD findings. A more elaborate discussion of our findings as well as a traffic characterization that reconciles these contradictory results appears elsewhere;¹⁸ there, we argue how Internet traffic demonstrates a nonstationary, time-dependent Poisson process and, when viewed across very long time scales, exhibits the observed LRD.

To test the Poisson traffic model's validity, we must examine two key properties: whether packet interarrival times follow the exponential distribution, and whether packet sizes and interarrival times appear mutually independent. Congestion in today's Internet usually appears on access links rather than in the backbone where ISPs overprovision their networks: traffic characteristics can vary in such links, which means our findings might not apply.

Distribution of Packet Interarrival Times

An interarrival-time distribution consists of two portions, one that contains back-to-back packets and another with packets guaranteed to be separated by idle time. For heavily utilized links, interarrival times are a function of packet sizes because many packets are sent back to back. For overprovisioned links, the distribution tends to contain most probability in the “idle” portion (where packets are separated by idle time).

We can closely approximate packet interarrival times for our traces by using an exponential distribution. Figure 2 shows the packet interarrival distributions for two of the backbone traces. The Complementary Cumulative Distribution Function (CCDF) of packet interarrival times is a straight line when the y-axis is plotted in log scale, which corresponds to exponential distribution.

To highlight the differences between current backbone traces and past Ethernet-link traces, Figure 2 also shows the CCDF of interarrival times for the famous BC-pAug89 trace, which was first used to prove LRD in network traffic in the pioneering work of Leland and colleagues.⁴ It was recorded at 11:25 (EDT) on 29 August 1989 from an Ethernet at the Bellcore Morristown Research and Engineering facility.

Figure 2 shows a minor discrepancy between our traces and the exponential distribution for small values (that is, less than $6\mu s$; $5\mu s$ is the time required for the transmission of 1,500-byte packets in an OC48 link) of the interarrival times. This discrepancy is caused by *train-like* interarrivals (back-to-back packets not separated by any intermediate idle time) during busy periods at the upstream router. However, the Poisson traffic model assumption does not require that interarrival times follow a perfect exponential distribution. In fact, these deviations and short-range artifacts can be incorporated into the Poisson model as “packet trains.”¹⁹

Independence

We separately examined and showed that packet sizes and interarrival times appear to be uncorrelated in our traces. We validated the independence by using various tests, such as the ACF and cross-correlation function (XCF), visual examination of conditional probabilities and scatter plots, the Box-Ljung statistic, and Pearson’s chi-square test for independence. Other researchers have used similar tools in the literature to test the independence hypothesis.^{5,13}

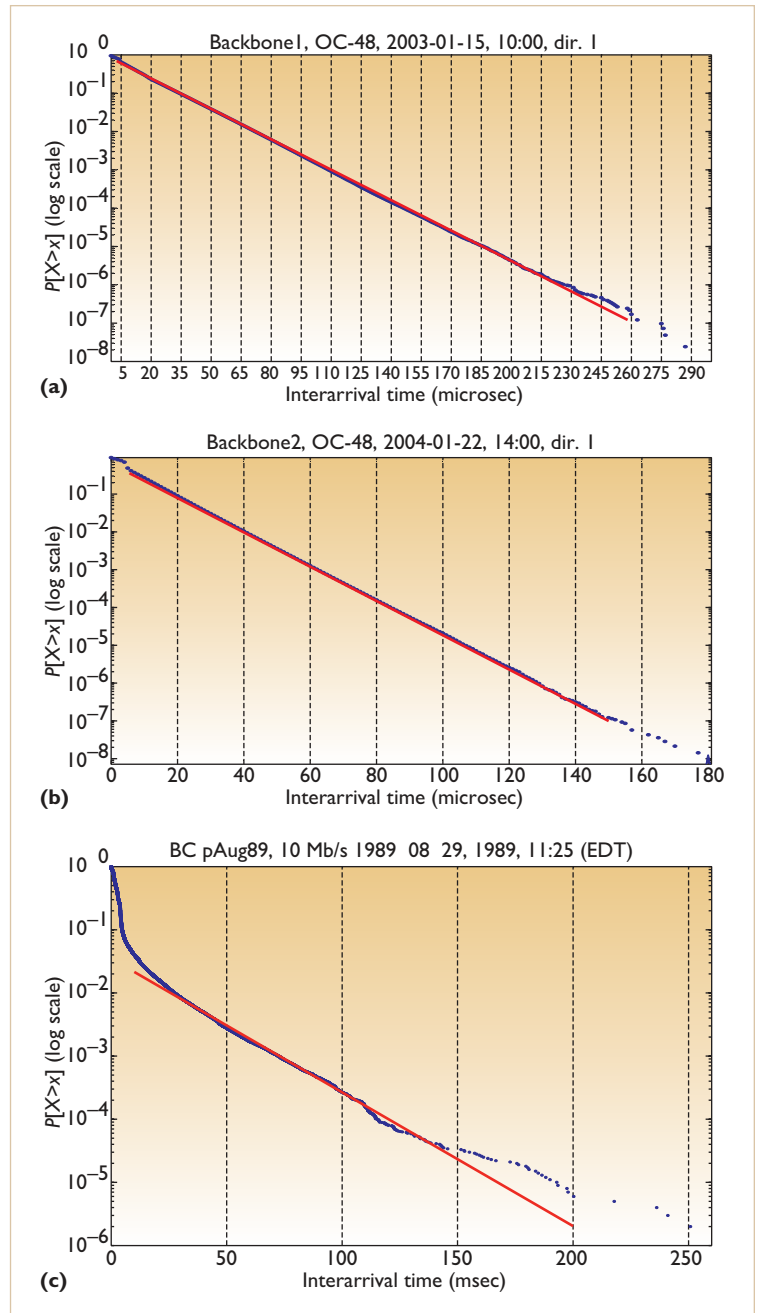


Figure 2. Complementary Cumulative Distribution Function of packet interarrival times over two backbone networks. For OC48 link traces on (a) January 2003, backbone 1, and (b) January 2004, backbone 2, as well as for (c) the BC-pAug89 1989 Bellcore trace, the y-axis is plotted in logarithmic scale. We can approximate the distributions of OC48 traces with an exponential distribution (straight line in log-linear scale), but the BC-pAug89 data set clearly deviates from the exponential distribution.

Using the ACF, we examined two different time series. The *sizes* series consisted of the actual packet sizes as individual packets arrive, and the *inter-*

arrival series consisted of timestamp differences between consecutive packets. Apart from limited correlation at small time lags (less than 20), sizes and interarrivals weren't correlated. The trivial correlation at small time lags close to zero was due to back-to-back packets, as described earlier. The XCF between sizes and interarrivals points to independence beyond minimal correlation at small lags.

Independence was also suggested by the Box-Ljung statistic Q_k , defined as

$$Q_k = n(n+2) \sum_{i=1}^k \frac{\rho_i^2}{n-i},$$

where ρ_i is the autocorrelation coefficient for lags $1 \leq i \leq k$ and n is the series' length. To test the null hypothesis (that is, independence), we compared the Q_k statistic with the chi-square distribution, which had k degrees of freedom. We applied the test for varying numbers of consecutive packet arrivals for both the interarrival times and packet sizes. The Box-Ljung statistic shows that we can consider that both variables are not correlated with 95 percent confidence for up to a certain number of consecutive packet arrivals. The point at which dependence appears differs with the trace and time within the trace – for example, independence holds for 20,000 consecutive packet interarrivals on average according to the test for the January 2003, backbone 1 trace. For the packet-sizes series, the average is approximately 16,000 consecutive packet arrivals.

We validated these findings by applying Pearson's chi-square test for independence. In all cases, we accepted the null hypothesis for similar numbers of consecutive interarrivals (as with the Box-Ljung statistic), provided that we apply the test to the "idle" portion of the distribution (that is, using interarrival times larger than $6 \mu\text{s}$ to remove back-to-back packet effects). Independence held for hundreds of thousands of consecutive interarrivals for the May 2003, backbone 2 trace.

LRD

Despite the Poisson characteristics of packet arrivals, our traces and analysis agreed with previous findings, showing that LRD characterizes backbone traffic. However, the intensity of correlation depends on the scale of observation. Specifically, in all traces analyzed, we saw a dichotomy in the scaling in agreement with previous studies;^{10,11} The intensity of LRD depends on the scale. The change point is within the millisecond scale,

albeit slightly different for each case, but the pattern is the same: at scales below the change point, the Hurst exponent is just above 0.6. At larger scales, it varies between 0.7 and 0.85 depending on the trace and the estimator used. We studied the series of byte and packet counts with smallest aggregation level at $10 \mu\text{s}$.

Conclusions

The findings we've presented here might further challenge established beliefs. They reflect an extremely dynamic and constantly evolving network expanding in size and complexity. Further analysis of other backbone links as well as links near the network's periphery seems compelling. We could very well discover that individual links exhibit varying behavior, especially at small time scales. Why should traffic be an exception to the Internet's diversity?

The problem of characterizing Internet traffic is not one that can be solved easily, once and for all. As the Internet increases in size and the technologies connected to it change, we must constantly monitor and reevaluate our assumptions to ensure that our conceptual models correctly represent reality. □

Acknowledgments

This work was supported by grants from the US National Science Foundation from the Advanced Network Infrastructure and Research (ANIR-9985195) and the Information and Data Management (IDM-0208950) programs.

References

1. L. Kleinrock, *Queueing Systems, Volume II: Computer Applications*, John Wiley & Sons, 1976.
2. J. Beran, *Statistics for Long-Memory Processes*, Chapman and Hall, 1994.
3. K. Park and W. Willinger, "Self-Similar Network Traffic: An Overview," *Self-Similar Network Traffic and Performance Evaluation*, Wiley-Interscience, 2000, 1–39.
4. W.E. Leland et al., "On the Self-Similar Nature of Ethernet Traffic" *IEEE/ACM Trans. Networking*, vol. 2, no. 1, 1994, pp. 1–15.
5. V. Paxson and S. Floyd, "Wide Area Traffic: The Failure of Poisson Modeling," *IEEE/ACM Trans. Networking*, vol. 3, no. 3, 1995, pp. 226–244.
6. A. Erramilli, O. Narayan, and W. Willinger, "Experimental Queueing Analysis with Long-Range Dependent Packet Traffic," *IEEE/ACM Trans. Networking*, vol. 4, no. 2, 1996, pp. 209–223.
7. K. Park, G. Kim, and M.E. Crovella, "On the Relationship Between File Sizes Transport Protocols, and Self-Similar

- Network Traffic," *Int'l Conf. Network Protocols*, IEEE CS Press, 1996, pp. 171–180.
8. M.E. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," *IEEE/ACM Trans. Networking*, vol. 5, no. 6, 1997, pp. 835–846.
 9. W. Willinger et al., "Self-Similarity through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level," *IEEE/ACM Trans. Networking*, vol. 5, no. 1, 1997, pp. 71–86.
 10. A. Feldmann et al., "The Changing Nature of Network Traffic: Scaling Phenomena," *ACM Computer Comm. Rev.*, vol. 28, Apr. 1998, pp. 5–29.
 11. Z.L. Zhang et al., "Small-Time Scaling Behaviors of Internet Backbone Traffic: An Empirical Study," *Proc. IEEE Infocom*, IEEE CS Press, 2003, pp. 1826–1836.
 12. J. Cao et al., "On the Nonstationarity of Internet Traffic," *Sigmetrics/Performance*, ACM Press, 2001, pp. 102–112.
 13. Y. Zhang et al., "On the Constancy of Internet Path Properties," *Proc. ACM Sigcomm Internet Measurement Workshop*, ACM Press, 2001, pp. 197–212.
 14. M.S. Taqqu and V. Teverovsky, "On Estimating the Intensity of Long-Range Dependence in Finite and Infinite Variance Time Series," *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*, Birkhauser, 1998, pp. 177–217.
 15. P. Abry and D. Veitch, "Wavelet Analysis of Long-Range Dependence Traffic," *IEEE Trans. Information Theory*, vol. 44, no. 1, 1998, pp. 2–15.
 16. T. Karagiannis, M. Faloutsos, and R.H. Riedi, "Long-Range Dependence: Now You See It, Now You Don't!" *Proc. IEEE Global Telecommunications Conf. Global Internet Symp.*, 2002.
 17. V. Paxson, "Fast Approximation of Self Similar Network Traffic," *Computer Comm. Rev.*, vol. 27, no. 5, 1997, pp. 5–18.
 18. T. Karagiannis, M. Molle, and M. Faloutsos, "A Nonstationary Poisson View of Internet Traffic," *Proc. IEEE Infocom*, IEEE CS Press, 2004.
 19. R. Jain and S. Routhier, "Packet Trains: Measurements and a New Model for Computer Network Traffic," *IEEE J. Select. Areas Comm.*, vol. 4, no. 6, 1986, pp. 986–995.

Thomas Karagiannis is a PhD candidate in the Department of Computer Science and Engineering at the University of California, Riverside. His technical interests include Internet traffic measurements, analysis of Internet traffic dynamics, Internet protocols, and peer-to-peer networks. Karagiannis received a BSc in the Department of Applied Informatics at University of Macedonia, Thessaloniki, Greece. He is a member of IEEE. Contact him at tkarag@cs.ucr.edu.

Mart L. Molle is a professor in the Department of Computer Science and Engineering at the University of California, River-

side. His research interests include the performance evaluation of protocols for computer networks and of distributed systems. Molle received a BSc (Hons.) in mathematics/computing science from Queen's University at Kingston, Canada, and an MS and PhD in computer science from the University of California, Los Angeles. He is a member of the IEEE. Contact him at mart@cs.ucr.edu.

Michalis Faloutsos is a faculty member in the Computer Science Department at the University of California, Riverside. His interests include Internet protocols and measurements, multicasting, and ad hoc networks. Faloutsos received a BS in electrical and computer engineering from the National Technical University of Athens and an MSc and PhD in computer science from the University of Toronto. Contact him at michalis@cs.ucr.edu.