On the multiplicative structure of network traffic

Rudolf H. Riedi*
Department of Electrical and Computer Engineering
Rice University 6100 South Main Street Houston, TX 77005, USA

1 Introduction

One crucial property of high-speed network traffic is the presence of *long-range dependence* (LRD), which was demonstrated convincingly in the landmark paper of Willinger et al [1]. This feature is inadequately described by classical traffic models such as homogeneous Markov or Poisson models, yet has important performance implications in particular on queuing delays [1, 2, 3, 4, 5].

Among the numerous LRD traffic models, the Gaussian nature and strong scaling properties of *fractional Brownian motion* (fBm) have made it most amenable to study large time scale asymptotics [5, 4] which are of importance for design and management. Being Gaussian in nature, though, fBm-based models can fail to capture essential traffic properties such as positivity and burstiness in particular on small time scales (see Figure 1) which are relevant to control and performance evaluation.

Leaving the realm of 'classical' linear processes which are additive in kind, [6] proposed multiplicative cascades to model traffic loads. We feel that this simple class of multifractals is convenient to develop some of the main differences between additive and multiplicative processes. It constitutes the bridge between multifractal measures and multifractal processes, has algorithmic advantages and possesses 'conservation of mass' which simplifies estimation procedures. We conclude with an outlook into wider classes multifractal processes.

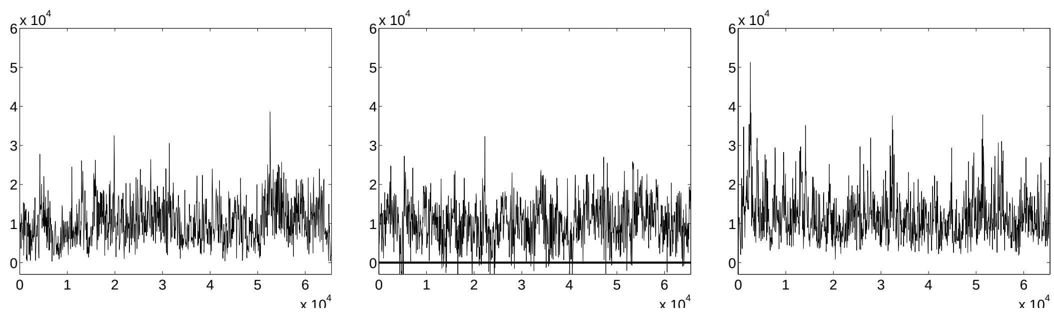

Figure 1: *Real data traffic AUCK [7] and two tree-based models that fit the variance on all dyadic scales. On the left the trace, in the middle the additive WIG model and on the right the multiplicative MWM model (see text) at 64 msec time resolution.*

*This work was supported by NSF, grant no. ANI-9979465 (RENE), by ONR, grant no. N00014-99-10813 and by DARPA Agreement No. F30602-00-0557 (NMS) Email: riedi@rice.edu. URL: www.dsp.rice.edu.

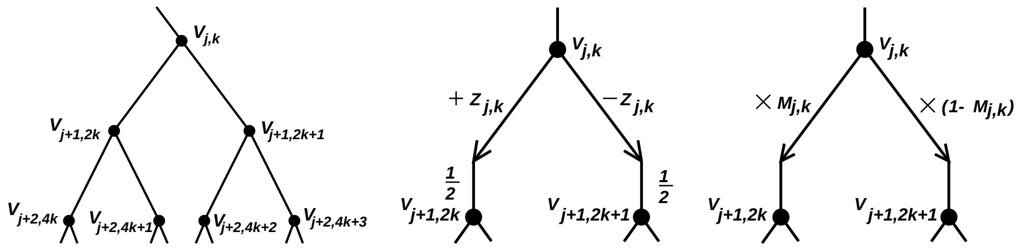

Figure 2: *(Left) The multiscale tree representation can be computed from data flowing up the tree according to (2). For synthesis flow down the tree: (Middle) Additive tree-model WIG (see (3)); (Right) Multiplicative tree-model MWM (see (4)).*

2 Stochastic modeling on a tree

LRD can be viewed as a multi-resolution phenomenon. Indeed, a discrete-time, wide-sense stationary random process $\{X_k,\ k \in \mathbf{Z}\}$ is said to exhibit LRD if its auto-covariance $r_X[k]$ decays as $\simeq k^{2H-2}$ with $1/2 < H < 1$ and is not sumable, or equivalently, if the variance of the aggregates

$$X_k^{(m)} \ = \ X_{km-m+1} + \cdots + X_{km} \tag{1}$$

behaves as m^{2H} [8]. One example of an LRD process is *fractional Gaussian noise* G (fGn), the increment process of fBm for which $r_G[k] = \frac{\sigma^2}{2}(|k+1|^{2H} - 2|k|^{2H} + |k-1|^{2H}) \simeq k^{2H-2}$ and $\mathrm{var}(X^{(m)}) = \sigma^2 m^{2H}$.

For the purpose of analysis as well as synthesis a Haar wavelet tree is a most efficient and simple tool to represent a discrete process on multiple resolutions (see Figure 2): The nodes at the level j of the binary tree are

$$V_{j,k} = X_k^{(2^{n-j})} = V_{j+1,2k} + V_{j+1,2k+1}. \tag{2}$$

Clearly, every parent node is the sum of its two children at the next finer scale. The tree structure allows for fast $O(N)$ pyramidal algorithms to process N data points. More general wavelets could be used to broaden the class of models [9, 10].

2.1 Additive Synthesis: Gaussian tree models

The additive nature of the multiscale tree (see Figure 2 and (2)) suggests a simple additive multiscale model for LRD traffic [9, 10]. Starting at the root of the tree $V_{0,0}$, we model children nodes iteratively using *independent additive random innovations*, $Z_{j,k}$'s, through

$$V_{j+1,2k} = (V_{j,k} + Z_{j,k})/2, \qquad V_{j+1,2k+1} = (V_{j,k} - Z_{j,k})/2. \tag{3}$$

Due to the Central Limit Theorem, we expect the distributions of the nodes $V_{j,k}$ to become Gaussian as $j \to \infty$. The $Z_{j,k}$'s must be identically distributed within scale j to provide a first-order stationary process. We call this model the *wavelet-domain independent Gaussian* (WIG) model. Choosing $V_{0,0} \sim \mathcal{N}(m2^n, \sigma^2 2^{2nH})$ and $Z_{j,k} \sim \mathcal{N}(0, \sigma^2(2^{2-2H} - 1)2^{2(n-j)H})$ provides a Gaussian process with the same tree statistics, in particular the same $\mathrm{var}(X^{(2^{n-j})})$, as fGn.

2.2 Multiplicative synthesis: Multifractal Processes

The WIG model is approximately Gaussian by construction. This is in sharp contrast with the fact that network traffic signals (such as loads and interarrival times) can be highly "spiky" and non-Gaussian (recall Figure 1). Retaining the simplicity of a tree-based WIG model, the *multifractal wavelet model* (MWM) of [6] achieves a better match using independent *multiplicative innovations* $M_{j,k}$ as

$$V_{j+1,2k} = V_{j,k} M_{j,k}, \qquad V_{j+1,2k+1} = V_{j,k}(1 - M_{j,k}). \tag{4}$$

By choosing $M_{j,k} \in (0,1)$ and $V_{0,0}$ to be positive, we ensure positive values at all nodes on the tree. Requiring that the multipliers $M_{j,k}$ be symmetric about 1/2 and identically distributed within scale ensures first-order stationarity of the process at scale j. Beta and even simple discrete distributions for the multipliers have proven useful in applications [6]. Due to the CLT and the multiplicative scheme, the process is approximately lognormal at the finer time scales, and spiky as we are about to explain.

2.3 Multifractal properties

In a nutshell, the multifractal formalism relates the frequency of bursts in one realization to the decay of sample moments. Using Chebyshev's inequality we have for any $q > 0$:

$$P_j\left[V_{j,k} \geq 2^{-ja}\right] = P_j\left[V_{j,k}^q \geq 2^{-jqa}\right] \leq \frac{\mathbb{E}_j[V_{j,k}^q]}{2^{-jqa}}, \tag{5}$$

where P_j is the counting measure of $k \in \{0, \ldots 2^j - 1\}$. Large deviation principles can be used to establish for which q this upper bound (5) is tight in the limit $j \to \infty$ [11]. We may interpret the left hand side of (5) as the chance to observe in one realization arrivals which are very large compared to the time scale. The main term on the right hand side are sample moments. For the MWM, assuming i.i.d. multipliers $M_{j,k}$, we have $\mathbb{E}[V_{j,k}^q] = (\mathbb{E}[M^q])^j$ and find for (5) an exponential decay rate of $\inf_q(qa + \log_2 \mathbb{E}[M^q])$, a convex function of a.

For the WIG with fGn scaling, that is $V_{j,k} \sim \mathcal{N}(0, \sigma^2 2^{-2jH})$, we find $\mathbb{E}[V_{j,k}^q] = c_q \cdot (2^{-qH})^j$ and the exponential rate of (5) $\inf_q(q\alpha - qH)$ degenerates. Indeed, due to a result of Adler $X_k^{(m)}$ is for fGn typically at the most of the order of m^H for *all* t of one realization [11]. For real world traffic traces, however, one finds widely varying burst strength a in any realization, just like the MWM [6].

2.4 Doubly stochastic modeling

In a doubly stochastic approach we model the nodes $V_{j,k}$ of the multi-resolution tree as noisy observables with stochastic background parameters, say their *means* $\theta_{j,k}$, which we are really interested in. We assume that the observations conditioned on knowing the parameters, i.e. $U_{j,k} = V_{j,k}|\theta$ are independent. Two cases are of most interest: Gaussian $U_{j,k} \sim \mathcal{N}(\theta_{j,k}, \sigma_{j,k}^2)$ and Poissonian $U_{j,k} \sim \mathcal{P}(\theta_{j,k})$ observables.

Note first that by additivity of the expected value the parameters $\theta_{j,k}$ form a tree where again the parent is the sum of the two children. In the Gaussian case conditional independence determines the variances: $\sigma_{j,k}^2 = \sigma_{j+1,2k}^2 + \sigma_{j+1,2k+1}^2$. For synthesis downstream on the tree we compute easily that

$$U_{j+1,2k}|U_{j,k} = \begin{cases} \mathcal{N}(\frac{1}{2}U_{j,k} + \frac{\theta_{j+1,2k} - \theta_{j+1,2k+1}}{2}, \frac{1}{2}\sigma_{j,k}^2) & \text{in the Gaussian case,} \\ \text{Binom}(U_{j,k}, \frac{\theta_{j+1,2k}}{\theta_{j,k}}) & \text{in the Poisson case.} \end{cases} \tag{6}$$

This demonstrates impressively that Gaussian resp. Poissonian noise models lead naturally to additive resp. multiplicative tree models for the mean-process. Thus, we are confirmed that multiplicative models are more accurate and appropriate on small scales, where individual arrivals are relevant, whereas on large scales, the high mean from large aggregations makes Poisson and Gaussian distributions very similar and an additive, Gaussian model becomes as appropriate, and potentially easier to treat analytically.

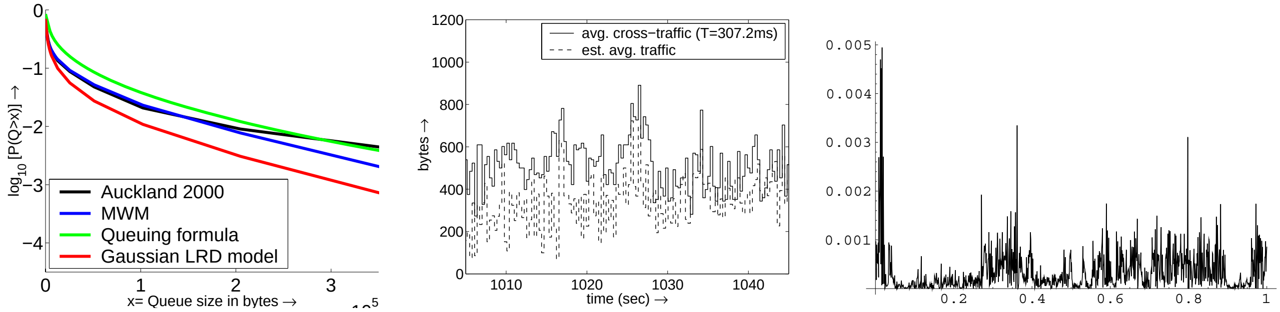

Figure 3: *Left: Queue tail probabilities obtained from an numerical experiment for a real data trace (AUCK), in comparison with fitted, synthetic WIG and MWM traces. Shown are also analytical expressions (MSQ) for the two models. Middle: Model based, adaptive inference of cross traffic over a shared link (see [13]). Right: A sample path of a Multifractal Product of Processes (see [15]).*

2.5 Multiplicative network traffic?

Apart from the LRD match, in which WIG and MWM have the same accuracy, the multiplicative MWM outperforms its additive competitor WIG as a model of network traffic in terms of multifractal properties as seen above, in its queuing behavior both in simulation as well as in an analytical approach (see [12] and Figure 3 left), and allows to adaptively infer volumes of cross traffic that an individual connection experiences on a router (see [13] and Figure 3 middle). As for the cause of the apparent multiplicative structure only clues are available at this time, pointing to the transport protocol TCP as the correct layer to investigate [14].

3 Multifractal Product of Processes

To overcome the inherent non-stationarity of tree-based processes, [15] introduced the *multifractal product of stationary processes* (MPSP):

$$A_n(t) = \int_0^t \prod_{j=0}^{n} \Lambda^{(j)}(s) \; \mathrm{d}s. \tag{7}$$

In the simplest case the multipliers $\Lambda^{(j)}(t/2^j)$ are i.i.d. stationary processes with $\mathbb{E}\Lambda^{(i)}(t) = 1$. The product right hand side of (7) converges only in distributional sense to a meaningful limit, since it will converge pointwise to zero almost everywhere. In more simple words, a multiplier $\Lambda^i(t)$ can be thought of as a local change in the arrival rate where one is interested actually in the integrated 'total traffic load' process $A_n(t)$. The covariance structure of the $\Lambda^{(i)}$ determines whether A_n converges in $\mathcal{L}_2$. Note that this construction guarantees conservation only in the mean, so A_n is *not* a coarse scale aggregation of the limiting process. This makes parameter estimation hard.

It should be noted that the MWM could be formally written in the form (7) except that the multipliers $\Lambda^{(i)}$ would not be stationary since they change value exactly at dyadic time instances. A simple example of an MPSP is then to let the $\Lambda^{(i)}$ be piecewise constant with independent values which change according to a Poisson process. In this case, A_n converges in $\mathcal{L}_2$ if $\mathrm{var}(\Lambda^{(i)}) < 1$ and the multifractal properties is ruled by higher order moments, i.e. the decay rate of (5) is here [15]

$$\inf_q (q\alpha + \log_2 \mathbb{E}[(\Lambda/2)^q]).$$

References

[1] W. Willinger, M. Taqqu, R. Sherman, and D. Wilson, "Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level," *IEEE/ACM Trans. Networking (Extended Version)*, vol. 5, pp. 71–86, Feb. 1997.

[2] A. Erramilli, O. Narayan, and W. Willinger, "Experimental queueing analysis with long-range dependent traffic," *IEEE/ACM Transactions on Networking*, vol. 4, pp. 209–223, April 1996.

[3] V. Paxson and S. Floyd, "Wide area traffic: the failure of Poisson modeling," *IEEE/ACM Trans. Networking*, vol. 3, pp. 226–244, June 1995.

[4] N. Duffield and N. O'Connell, "Large deviations and overflow probabilities for the general single-server queue, with applications," *Math. Proc. Cambr. Phil. Soc.*, vol. 118, pp. 363–374, 1995.

[5] I. Norros, "A storage model with self-similar input," *Queueing Systems*, vol. 16, pp. 387–396, 1994.

[6] R. Riedi, M. S. Crouse, V. Ribeiro, and R. G. Baraniuk, "A multifractal wavelet model with application to TCP network traffic," *IEEE Trans. Info. Theory, Special issue on multiscale statistical signal analysis and its applications*, vol. 45, pp. 992–1018, April 1999.

[7] NLANR, "Auckland-II trace archive." http://moat.nlanr.net/Traces/Kiwitraces/, 2000. trace 20000125-143640, corresponding to 3:11:28 hours worth of mostly TCP traffic.

[8] D. R. Cox, "Long-range dependence : A review," *Statistics : An Appraisal*, pp. 55–74, 1984.

[9] P. Flandrin, "Wavelet analysis and synthesis of fractional Brownian motion," *IEEE Trans. Inform. Theory*, vol. 38, pp. 910–916, Mar. 1992.

[10] A. H. Tewfik and M. Kim, "Correlation structure of the discrete wavelet coefficients of fractional Brownian motion," *IEEE transactions on information theory*, vol. 38, pp. 904–909, March 1992.

[11] R. H. Riedi, "Multifractal processes," *in: "Long range dependence : theory and applications", eds. Doukhan, Oppenheim and Taqqu*, to appear 2001.

[12] V. Ribeiro, R. Riedi, M. S. Crouse, and R. G. Baraniuk, "Multiscale queuing analysis of long-range-dependent network traffic," *IEEE Trans. Networking*, submitted.

[13] V. Ribeiro, R. Riedi, M. Coates, and R. G. Baraniuk, "Multifractal cross-traffic estimation from end-to-end measurements," *Proceedings ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management, Monterey, CA*, Sept. 2000.

[14] R. H. Riedi and W. Willinger, *Self-similar Network Traffic and Performance Evaluation*, ch. Toward an Improved Understanding of Network Traffic Dynamics, pp. 507–530. Wiley, 2000. Chapter 20, K. Park and W. Willinger eds.

[15] P. Mannersalo, I. Norros, and R. Riedi, "Multifractal products of stochastic processes," *COST257*, 1999, 31.