Multiscale Connection-level Analysis of Network Traffic

Shriram Sarvotham, Rudolf Riedi and Richard Baraniuk

Department of Electrical and Computer Engineering, Rice University 6100 Main st., Houston, TX, 77005

Abstract

Network traffic exhibits drastically different statistics, ranging from nearly Gaussian marginals and long range dependence at very large time scales to highly non-Gaussian marginals and multifractal scaling on small scales. This behavior can be explained by decomposing traffic into two components according to the connection bandwidth: the small bandwidth component absorbs most traffic and is Gaussian, while large bandwidth component constitutes virtually all of the small scale bursts. Based on this understanding, we propose a novel traffic model that parsimoniously accounts for user behavior, network topology, and the heterogeneous distribution of network bandwidths.

1. Introduction

Numerous studies have found that aggregate traffic exhibits *fractal* or *self-similar* scaling behavior, that is, the traffic "looks statistically similar" on all time scales [1]. This discovery led to new fractal aggregate traffic models. *Frac-tional Gaussian noise* (fGn), the most widely applied fractal model, is a Gaussian process with strong scaling behavior. A strong argument for fGn in networks is that often aggregate traffic can be viewed as a superposition of a large number of independent individual ON/OFF sources that transmit at the same rate but with heavy-tailed ON durations [2, 3]. In the limit of infinitely many sources, the ON/OFF model converges to fGn. The *long range dependence* LRD in the resulting process stems from the heavy-tailed nature of the ON durations.

Unfortunately, fGn is unrealistic for bursty non-Gaussian traffic. For instance, when the standard deviation of the traffic exceeds its mean, a considerable portion of an fGn traffic synthesis is negative. These failings have motivated more complicated models for aggregate traffic such as multifractals and infinitely divisible cascades [4, 5]. However, while more statistically accurate, these models lack network relevance in their parameterizations. In particular, they do not account for *why* bursts occur in network traffic.

The aggregated total load in the ON/OFF model converges to fBm as the number of sources go to infinity first, followed by letting the time interval for aggregation of load go to infinity. However, if we reverse the order and consider the limit where the time interval tends to infinity first and then the number of sources, we get a Lévy stable motion [6], a process with independent and stationary increments but with infinite variance. See [7] for an overview. We argue through a careful analysis of several real world traces that network traffic is a mixture of both limits.

In this paper, we present the alpha/beta model in which the fGn part of traffic constitutes the *beta* component, and the Lévy part of traffic constitutes the *alpha* component. We show that the alpha and beta components are linked to the connection level information and hence have networking relevance. Alpha component comes from connections that have large files to send and have high available bandwidth. Beta component comes from connections that have small available bandwidth and/or small files to transfer. We explain the alpha/beta model through the heterogeneity of connection bandwidths and round-trip times (RTT). We also suggest a fast scheme using wavelets to separate the alpha and beta components. Finally, we present a multiscale non-Gaussianity analysis on the alpha and beta components.

2. Separation of Traffic

2.1. Alpha and beta components of traffic

Connection-level information enables us to conduct a refined analysis of traffic bursts. In Gaussian aggregate traffic models (such as the classical ON/OFF model [3]), traffic bursts arise from a large number of connections transmitting bytes or packets simultaneously. That is, bursts stem from a kind of "constructive interference" of many connections. With connection-level information, we can test this hypothesis. If it were true, then we should observe in real traffic traces a large number of active connections during bursts. However, Figure. 1(a) and (b) demonstrate that this is not the case. Bursts in bytes-per-time generally do not coincide with large values of connections-per-time.

Quite to the contrary, a careful analysis of many real

Email: {shri, riedi, richb}@rice.edu; Web: spin.rice.edu. This work was supported by NSF, DARPA, and DOE.



Fig. 1. (a) Bytes-per-time and (b) number of connections-per-time of an aggregate traffic trace collected in Auckland, NZ [8].

traces [9] reveals that generally very few high-rate connections dominate during a burst. In fact in most cases only one connection dominates. This motivates us to propose a connection-level separation of network traffic. We aggregate the dominant connections and call it the alpha component of traffic. The remaining connections aggregate to give the beta component of traffic. ¹ Our procedure thus decomposes an aggregate traffic trace into

total traffic = alpha traffic + beta traffic.
$$(1)$$

See Figure 2 for real data example.

We have applied the alpha/beta traffic decomposition to many real-world traffic traces, from Auck [8] to LBL [10] and found tremendous consistency in our results [9]. The statistical properties of the components can be summarized as follows.

Beta traffic: At time-scales coarser than the RTT, the beta component is very nearly Gaussian and strongly LRD (i.e., approximately fGn), provided a sufficiently large number of connections are present. Moreover, the beta component carries the same fractal scaling (LRD) exponent as the aggregate traffic [9].

Alpha traffic: The alpha component constitutes a small fraction of the total workload but is entirely responsible for the bursty behavior. Alpha traffic is highly non-Gaussian.

It is notable that this decomposition in networking terms (based on connection-level information) achieves a separation in statistical terms.

2.2. Fast alpha/beta separation using wavelets

The computationally intense connection-level separation of alpha and beta traffic does not lend itself to massive data processing or on-line monitoring [9]. Fortunately, approximate separation of alpha and beta traffic can be done using a wavelet-based thresholding scheme that does not require explicit connection information. This scheme is based on the fact that we can treat the beta component as "noise" and the alpha component as the "signal", and use well-known denoising techniques to separate the two. We use wavelet based denoising techniques, with coefficient thresholding. For colored denoising (since beta traffic is colored noise, fGn), we use different thresholds for wavelet coefficients at different scales. Kaplan and Kuo [11] have shown that for Haar wavelet, the variance progression of the wavelet transform of fGn with Hurst parameter H satisfies a power-law decay:

$$\operatorname{var}(W_{i,k}) = \sigma^2 \, 2^{(2H-1)(j-1)} \, (2 - 2^{2H-1}). \tag{2}$$

In a colored denoising scheme, the threshold at each scale is made proportional to the expected standard deviation of the wavelet coefficients at that scale. Thus, knowing the Hurst parameter, we can fix the threshold at each scale using equation (2). Johnstone *et al* [12] have shown that this thresholding scheme is optimal for colored denoising. For more details, see [9].

3. Origins of Alpha and Beta Traffic

We now argue that the end-to-end bandwidth is a crucial factor that determines whether a connection dominates or not. We define *alpha groups* as those end-to-end groups which contain at least one alpha connection. Let us focus for a moment on the top row of Figure 4, which shows three *alpha* groups. For each connection in a group, we determine the *peak rate*, i.e., its maximum sending rate for the time period T under study, by sliding a moving window along the connection duration. We observed that typically *all* connections of an alpha group with a *sufficiently large* transfer load

¹By analogy to the dominating *alpha* and submissive *beta* members of the animal kingdom.



Fig. 2. Decomposition of the traffic trace into the sum of a bursty alpha component (Lévy) and a Gaussian beta component (fGn).



Fig. 3. (a) Isolated bursts in the real trace. (b) Bursts isolated using the wavelet separation scheme.

are alpha connections. Thus, we verified that alpha connections naturally group according to their end-to-end host pairs. This leads us to conclude that bursts emerge indeed from end-to-end properties.

For the beta-groups (see bottom row of Figure 4) the connections with large loads are obviously not getting as much bandwidth as they could consume. This provides further confirmation to our claim that heterogeneity in bottleneck bandwidths is the cause of burstiness in traffic.

In summary, we conclude that all burst causing connections are due to *large file transfers with a large bottleneck* bandwidth in its end-to-end path.

Connection bandwidth is determined by the connection pipe size and RTT through

$$bandwidth = pipe size / RTT.$$
(3)

To see whether the heterogeneity in connection bandwidths is primarily due to heterogeneity in pipe size or RTT, we look at how the RTT correlates with the peak rate. For the trace under consideration, the correlation coefficient between 1/RTT and peak rate is 0.75, and the correlation coefficient between pipe size and peak rate is 0.10. Consistent results show that RTT plays a major role.

Consider a scenario where a relatively small number of hosts locate close to the measured link while the majority of hosts are far away. As a consequence, the distribution of RTT is bimodal: a small number of connections with small RTTs and the remaining with large RTTs. See Figure 5 for a real trace example. Let us further assume that the TCP congestion windows of all connections are equal.

Applying the superposition ON/OFF model or the infinite source Poisson model [7] to the scenario above, we obtain the following results. The alpha traffic is composed of the few connections with small RTTs and, therefore, high sending rates. The TCP clock ticks faster for these sources, therefore we approximate this component as the ON/OFF limit at infinite time, i.e., by i.i.d. Lévy stable noise. So the alpha traffic brings in burstiness but not LRD.

The beta traffic is made up from the bulk of remaining connections with large RTTs and low rates and is well approximated by the ON/OFF limit at infinite source number, i.e., by fGn. It inflicts LRD on the overall traffic but is not



Fig. 4. Plot of peak rate and total transfer for all connections which share the same pair of source and destination hosts. Note the high peak rates in the top row (alpha connections). Notice the starvation of connections in the groups of the bottom row and the overall low peak rate (beta connections).

as bursty as the alpha traffic.

We conclude that both connection RTT and pipe size are behind the heterogeneity in connection bandwidths, and RTT plays a much bigger role. Therefore, network traffic burstiness is due to a small number of high-bandwidth connections (alpha connections) which transmit large files and have small RTTs.

4. Multiscale analysis of Network Traffic

In this section, we present a mutiscale analysis on the network traffic components. We use the well-known Gaussianity measure, kurtosis, to measure the deviation from Gaussianity of the signal at different time-scales. Kurtosis of a random variable X is defined as the ratio of its fourth cental moment to the square of its variance

$$\operatorname{kurtosis}(X) = \frac{E(X-\mu)^4}{\sigma^4}.$$
 (4)

The kurtosis for Gaussian random variable is 3. Random variables that have slower decay of tails have a kurtosis greater than 3, and those that have faster decay of tails have kurtosis less than 3.

Fig. 5. RTT histogram for the Auck-2 trace; connections in the first 300,000 packets.

Figure 6 indicates that neither an additive model nor a multiplicative cascade model provides a satisfactory fit of Gaussianity across time scales compared with the real trace. We plot the kurtosis curves for alpha, beta, as well as the overall traffic. Due to inherent Gaussianity, the kur-

Fig. 6. Kurtosis-time-scale curves for (a) real trace; (b) Multiplicative cascade; (c) Additive cascade model.

tosis curves for the additive model are constant at 3 (within simulation error tolerance) at all time scales. While for the multiplicative model, the kurtosis curves are notably higher the those of the real trace, showing a lower decay of tail distribution compared with the real traffic data.

The analysis motivates additive-multiplicative mixture models for the components of network traffic. As part of our future work, we are investigating different mixture models to capture the multiscale kurtosis behavior of network traffic.

5. Conclusions

We have proposed a framework for analyzing and modeling network traffic that takes into account the crucial connectionlevel information that aggregate analysis ignores. The topological variability of the network enters through the distribution of RTTs and bottlenecks link speeds. In a real world situation these distributions will depend on the particular location where the measurements are taken. Client behavior will determine both the LRD component as well as how often large files are transferred over large bottlenecks. The alpha-beta traffic model has implications not only for modeling, but also for simulation, synthesis, estimation, prediction, performance evaluation and understanding of traffic dynamics. This approach also opens a clear alley towards studying the influence of RTT distributions on network traffic dynamics which is currently under way.

As future work, we intend to study the impact of alpha and beta components in queueing behavior. We have shown that the alpha component contributes to large build-up of queues. This suggests scope for AQM [13] techniques that treat alpha and beta packets differently. Another area of future work is to link the alpha/beta model to the network topology.

6. References

 P. Abry, R. Baraniuk, P. Flandrin, R. Riedi, and D. Veitch, "Multiscale nature of network traffic," *IEEE Signal Process-* ing Mag., vol. 19, no. 3, pp. 28-46, May 2002.

- [2] M. Crovella and A. Bestavros, "Self-similarity in World Wide Web traffic. Evidence and possible causes," *IEEE/ACM Transactions on Networking*, vol. 5, pp. 835– 846, December 1997.
- [3] W. Willinger, M. Taqqu, R. Sherman, and D. Wilson, "Selfsimilarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level," *IEEE/ACM Trans. Networking (Extended Version)*, vol. 5, no. 1, pp. 71–86, Feb. 1997.
- [4] R. H. Riedi, M. S. Crouse, V. Ribiero, and R. G. Baraniuk, "A multifractal wavelet model with application to TCP network traffic," *IEEE Trans. Inform. Theory*, vol. 45, no. 3, pp. 992– 1018, April 1999.
- [5] A. Feldmann, A. C. Gilbert, and W. Willinger, "Data networks as cascades: Investigating the multifractal nature of Internet WAN traffic," *Proc. ACM/Sigcomm* 98, vol. 28, pp. 42-55, 1998.
- [6] M. S. Taqqu, W. Willinger, and R. Sherman, "Proof of a fundamantal result in self-similar traffic modeling," *Computer Communications Review*, vol. 27, no. 2, pp. 5–23, 1997.
- [7] W. Willinger, V. Paxson, R. Riedi, and M. Taqqu, Long range dependence : theory and applications, chapter Long range dependence and Data Network Traffic, Wiley, 2002.
- [8] NLANR Network Traffic Traces, "http://pma.nlanr.net/traces,".
- [9] S. Sarvotham, R. Riedi, and R. Baraniuk, "Connection-level analysis and modeling of network traffic," Tech. Rep., ECE Dept., Rice Univ., July 2001.
- [10] LBL Traces, "http://ita.ee.lbl.gov/html/traces.html," .
- [11] L. M. Kaplan and C.-C. J. Kuo, "Fractal estimation from noisy data via discrete fractional Gaussian noise (DFGN) and the Haar basis," *IEEE Trans. Signal Proc.*, vol. 41, no. 12, pp. 3554–3562, Dec. 1993.
- [12] I. M. Johnstone and B. W. Silverman, "Wavelet threshold estimators for data with correlated noise," J. Royal Stat. Soc. B, no. 59, pp. 319–351, 1997.
- [13] S. Floyd and V. Jacobsen, "Random early detection gateways for congestion avoidance," *IEEE/ACM Trans. Networking* (*Extended Version*), vol. 1, no. 4, pp. 397–413, Aug. 1993.