

Connection-level Analysis and Modeling of Network Traffic

Shriram Sarvotham, Rudolf Riedi, Richard Baraniuk

Abstract— Most network traffic analysis and modeling studies lump all connections together into a single flow. Such aggregate traffic typically exhibits long-range-dependent (LRD) correlations and non-Gaussian marginal distributions. Importantly, in a typical aggregate traffic model, traffic bursts arise from many connections being active simultaneously. In this paper, we develop a new framework for analyzing and modeling network traffic that moves beyond aggregation by incorporating connection-level information. A careful study of many traffic traces acquired in different networking situations reveals (in opposition to the aggregate modeling ideal) that traffic bursts typically arise from just a few high-volume connections that dominate all others. We term such dominating connections *alpha traffic*. Alpha traffic is caused by large file transmissions over high bandwidth links and is extremely bursty (non-Gaussian). Stripping the alpha traffic from an aggregate trace leaves a *beta traffic* residual that is Gaussian, LRD, and shares the same fractal scaling exponent as the aggregate traffic. Beta traffic is caused by both small and large file transmissions over low bandwidth links. In our alpha/beta traffic model, the heterogeneity of the network resources give rise to burstiness and heavy-tailed connection durations give rise to LRD. Queuing experiments suggest that the alpha component dictates the tail queue behavior for large queue sizes, whereas the beta component controls the tail queue behavior for small queue sizes.

Keywords— network traffic modeling, animal kingdom

I. INTRODUCTION

NETWORK traffic analysis and modeling play a major rôle in characterizing network performance. Models that accurately capture the salient characteristics of traffic are useful for analysis and simulation, and they further our understanding of network dynamics and so aid design and control.

Most traffic analysis and modeling studies to date have attempted to understand *aggregate traffic*, in which all simultaneously active connections are lumped together into a single flow. Typical aggregate time series include the number of packets or bytes per time unit over some interval. 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]. Self similarity endows traffic with *long-range-dependence* (LRD) [1]. Numerous studies have also shown that traffic can be extremely bursty, resulting in a non-Gaussian marginal distribution [2]. These findings are in sharp contrast to classical traffic models such as Markov or homogeneous Poisson. LRD and non-Gaussianity can lead to much higher packet losses than predicted by classical Markov/Poisson queueing analyses [1], [3].

The discovery of self-similar behavior in traffic led immediately to new fractal aggregate traffic models (see [4], [5], for example). *Fractional Gaussian noise* (fGn), the most widely applied fractal model, is a Gaussian process with strong scaling behavior. Due to its Gaussianity, it lends itself to rigorous analytical studies of queueing behavior. Also, approximate fGn can be synthesized rapidly by a variety of different techniques, including wavelets.

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 [6], [7]. In the limit of infinitely many sources, the ON/OFF model converges to fGn.

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 [2], [8]. 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. This is precisely the question we study in this paper.

We will exploit the *connection-level information* contained in most publicly available traffic traces to target bursts directly. This information is typically ignored in a classical aggregate (LRD, fractal, multifractal) analyses.¹ The result is a new framework for analyzing and modeling bursty network traffic.

Department of Electrical Engineering, Rice University, Houston, TX 77005. Email: {shri, riedi, richb}@rice.edu; Web: dsp.rice.edu. This work was supported by NSF, DARPA/AFOSR, DOE, and ONR.

¹See [6], [7] for connection-level studies to explain the underlying causes of self-similar behavior.




Fig. 1. (a) Bytes-per-time and (b) number of connections-per-time of an aggregate traffic trace [9]. (c) Number of bytes per connection (sorted in decreasing order) during a typical burst. Clearly one connection dominates all others.

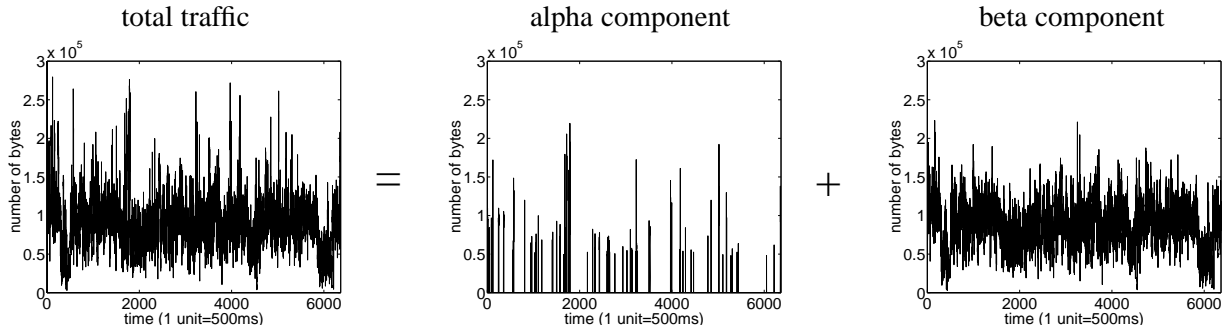


Fig. 2. Decomposition of the traffic trace into the sum of a bursty alpha component and an fGn beta component.

II. CONNECTION-LEVEL TRAFFIC ANALYSIS

We define a connection as the unique four-tuple comprising a source IP address, destination IP address, source port number, and destination port number.

Connection-level information enables us to conduct a refined analysis of traffic bursts. In aggregate traffic models (including the ON/OFF model), 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, Figures. 1(a) and (b) demonstrate that this is not the case. Bursts in bytes-per-time generally do not coincide with large values connections-per-time.

Quite to the contrary, a careful analysis of many real traces [10] reveals that generally *very few high-rate connections dominate during a burst*. In fact in most cases only *one* connection dominates. This surprising finding has far-reaching implications for traffic analysis and modeling.

To explore further, we propose a new analysis technique that exploits connection-level information to separate a measured traffic trace into two distinct components

at a time-scale T of interest.²

1. In each T -second time bin, identify the connection(s) that transmits the largest number of bytes.
2. If the strength of the identified connection(s) is greater than a threshold, then label it as a *dominant connection*. The (large) threshold is chosen based on the mean of the aggregate traffic at time-scale T plus a few standard deviations.

We call the traffic corresponding to the dominant connections the *alpha* component. The residual traffic is called the *beta* component.³ Our procedure thus decomposes an aggregate traffic trace into

$$\text{total traffic} = \text{alpha traffic} + \text{beta traffic}. \quad (1)$$

See Figure 2 for real data example.

We have applied the alpha/beta traffic decomposition to many real-world traffic traces, from Auck [9] to LBL [11] and found tremendous consistency in our results [10]. The statistical properties of the components can be summarized as follows.

Beta traffic: At time-scales coarser than the round-trip time, the beta component is very nearly Gaussian and

²While we set $T = 500$ ms for the analyses in this paper, our results held for a wide range of T .

³By analogy to the dominating *alpha males* and submissive *beta males* observed in the animal kingdom.

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 (see Figure 4(a)).

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.

See [10] for a number of computationally simpler schemes for decomposing traffic into alpha and beta components, including a scheme based on wavelet thresholding that does not require explicit connection information to extract the alpha traffic.

III. ORIGINS OF ALPHA AND BETA TRAFFIC

We have seen that bursts are not caused by a “conspiracy” of many moderate flows, but rather by solitary alpha connections. Here we study the origins of this behavior.

A. Potential causes of bursts

To begin, we list four possible reasons for connections to dominate and cause bursts. The first three are based on the predominant network protocol, TCP, while the last one blames the heterogeneity in bottleneck bandwidths.

1. *Transient response to re-routing:* Some connections could be rerouted from a high-bandwidth end-to-end path to the (lower-bandwidth) measured link. Since TCP feedback takes at least one round-trip time, we could expect a transient bursty behavior for such connections.
2. *Transient response to start/stop of connections:* When connections sharing a link terminate or are rerouted away from the measured link, other competing connections will grab their share of the bandwidth. This could potentially lead to bursty connections, especially for those connections in TCP slow-start.
3. *TCP slow-start peculiarities:* Some connections could get “lucky” during slow start in that they encounter no packet drops for an unusually long time.
4. *Heterogeneity in bottleneck bandwidths:* This scenario acknowledges the fact that connections are limited in their transmission rates by bottlenecks somewhere in the network, which may or may not be at the measured link. Bottlenecks can occur through limited capacity at a router, through shaping, or through a thin client, to give a partial list. When we have a large pool of connections, we would expect that the high bottleneck connections should dominate over the low-bottleneck connections and could potentially cause bursts. This scenario assumes that there are only a few connections that have high bottleneck bandwidth.

B. Dominance and bursts as end-to-end properties

We argue that the last of the above scenarios, i.e., *heterogeneity in bottleneck bandwidths*, is solely responsible for the bursts observed in our analyses.⁴

Our first and most striking observation is the clustering of alpha connections according to their source-receiver host pairs. We collect all connections with the same source and destination hosts into an *end-to-end group*. If random effects of the networking protocols were causing bursts, then we would find the dominating alpha connections distributed randomly over the end-to-end groups. The above-mentioned clustering indicates that the opposite is true.

More convincingly, our second observation states that if an end-to-end group contains *one* alpha connection, then *all* connections in that group are either alpha, or too short to cause a burst. This demonstrates clearly that bursts are caused by large volume connections over particular (necessarily high-bandwidth) end-to-end paths.

Finally, a third observation based on a more refined analysis gives a potential *physical reason* for the original of alpha bursts. To this end, we compare for each connection its total transfer load with its peak rate, which we compute as the maximum number of bytes sent by the connection during any time period of duration T . Figure 3 displays the total load versus peak rate for all connections within six particular end-to-end groups. The groups in the top row contain at least one alpha connection, the groups in the bottom row none.

For all of the connections in groups 1 and 2, most of the transfer is completed within a single time period of $T = 500$ ms. We conclude that these connections were not limited in their bandwidth consumption. In other words, the end-to-end paths corresponding to groups 1 and 2 have a high available bandwidth. To the contrary, the connections in groups 4–6 are obviously not getting as much bandwidth as they could consume. None of these connections are alpha. Finally, group 3 must have limited bandwidth, but at a level high enough to admit burst causing connections — indeed, all its connections are alpha. With such a consistent picture, we can exclude re-routing as the main cause for bursts, since it would be highly unlikely that all connections in a group would be systematically affected by it.

In summary, we conclude that the majority of burst causing connections are due to *large file transfers over a large bottleneck bandwidth end-to-end path*.

⁴We cannot exclude the possibility that one or more of the other scenarios cause an occasional burst. However, we focus here on the mechanisms that produce the overwhelming majority of bursts.

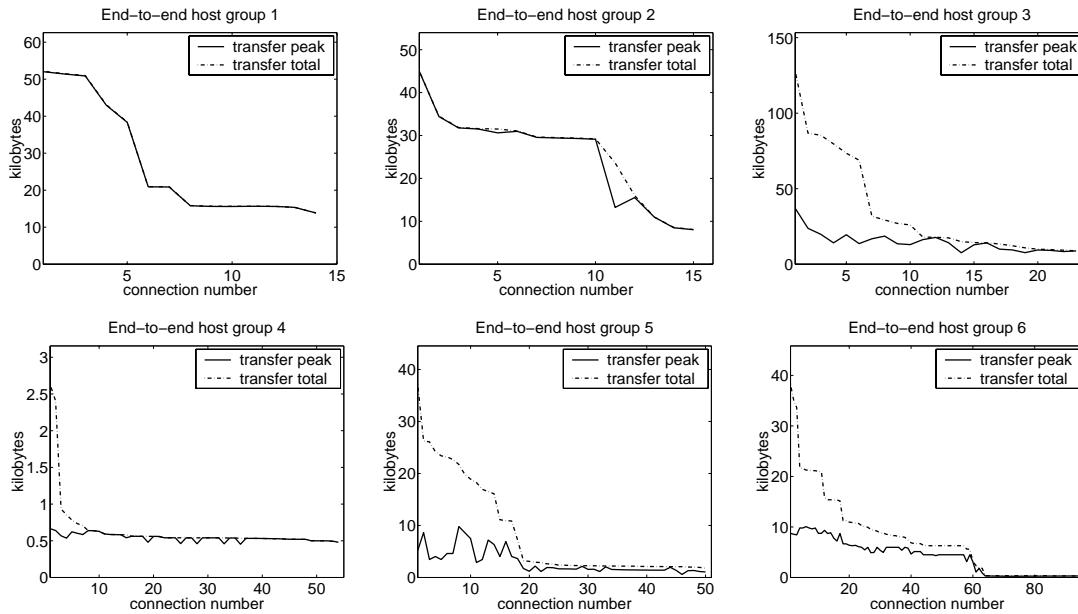


Fig. 3. Plot of peak rate and total transfer for six end-to-end groups (which comprise all connections sharing the same pair of source and destination hosts). Connections are sorted in decreasing order of total transfer. Note the high peak rates in the top row — groups 1–3 contain alpha connections. Note the starvation of connections and low overall peak rates in the bottom row — groups 4–6 contain only beta connections.

IV. CONNECTION-LEVEL TRAFFIC MODELING

Our alpha/beta decomposition technique suggests an intuitive and natural traffic model that takes into account both the network topology and user behavior. A major conclusion of our analysis is that the burstiness in network traffic is caused by the *heterogeneity* in link speeds and computational power within the network (including the networking software/hardware of the clients) and user behavior. The TCP protocols enters in its ability to grab free bandwidth quickly in slow-start.

Consider a simplified taxonomy of a network system. There are roughly two kinds of file sizes: large ones such as jpeg images and small ones such as text emails. There are also roughly two speeds of connections: fast ones such as Ethernet or DSL lines and slow ones such as 56k modems. We argue that *only large files over fast links contribute to alpha traffic* (see Table I). The remaining combinations aggregate together into fGn beta traffic. Therefore, we model traffic as a sum of alpha and beta traffic as in (1) and Figure 2. Such traffic is simple to both analyze and synthesize.

The beta component succinctly collects all the “average” connections and is well modeled as fGn with LRD parameter equal to that of the overall traffic. The alpha component consists of the dominant, burst causing connections and contributes low traffic volume but accounts for all the bursts, which arrive in a pattern according to the capabilities of network and requests of the clients routing

TABLE I

SIMPLE 4-WAY FILE SIZE VS. LINK-SPEED TAXONOMY.

		connection speed	
		slow	fast
file size	large	Beta	Beta
	small	Beta	Alpha

through the given point of measurement. The alpha burst arrivals are not exactly Poisson but can to a first approximation be modeled as such (most likely, they are compound Poisson). Furthermore, the dominant connections are uncorrelated with the overall connection arrivals and thus can be modeled as an independent process.

Our alpha/beta model is predictive, since for a given topology and user behavior, we can determine a priori which connections will aggregate into the fGn background and which connections will cause bursts. Syntheses from our model then closely match real traffic and controlled *ns* simulations [10]. For example, under heavy utilization (when the router at which we take our measurements becomes itself the bottleneck link for all connections) or with a considerably homogeneous clientele, we should and do observe fewer bursts and more Gaussian traffic. We observe exactly the contrary in the LBL data [11], since there

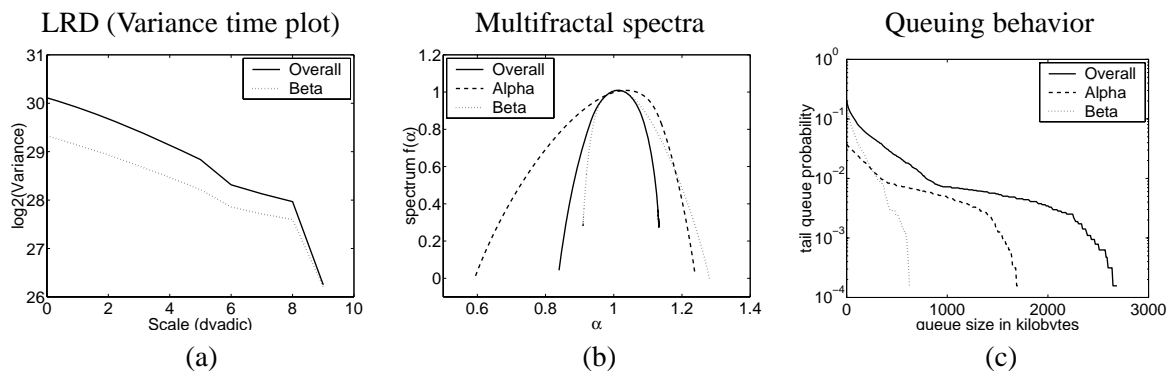


Fig. 4. Statistical properties of the components. (a) The beta component determines the LRD of the overall traffic (same slopes in a variance-time plot). (b) The alpha component is responsible for the bursts (as evidenced by the multifractal spectrum). (c) The alpha/beta components influence queue overflow probability primarily at large/small queue sizes.

are not enough active connections to give a full fGn beta component once the dominant connection is stripped away.

Finally, through queueing simulations, we have demonstrated that the beta component determines the tail queue probability for *small* queue sizes, whereas the alpha component determines the tail queue probability for *large* queue sizes (see Figure 4(c)). More experiments are underway.

V. CONCLUSIONS

We have proposed a new framework for analyzing and modeling network traffic that takes into account the crucial connection-level information that aggregate analysis ignores. Our alpha/beta traffic model incorporates both the *topological* and *user* aspects of networking. The topological variability of the network enters through the distribution of bottlenecks link speeds, which in a real world situation 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. Currently we are analyzing more traces to further test and prove-out the model.

REFERENCES

- [1] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Trans. Networking*, pp. 1–15, 1994.
- [2] R. H. Riedi, M. S. Crouse, V. Ribeiro, 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.
- [3] A. Erramilli, O. Narayan, and W. Willinger, "Experimental queueing analysis with long-range dependent traffic," *IEEE/ACM Trans. Networking*, pp. 209–223, April 1996.
- [4] F. Bricchet, J. Roberts, A. Simonian, and D. Veitch, "Heavy traffic analysis of a fluid queue fed by a superposition of ON/OFF sources," *COST*, vol. 242, 1994.
- [5] M. Taqqu and J. Levy, *Using renewal processes to generate LRD and high variability*. In: Progress in probability and statistics, E. Eberlein and M. Taqqu eds., vol. 11, Birkhaeuser, Boston, 1986, pp 73–89.
- [6] 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.
- [7] 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, no. 1, pp. 71–86, Feb. 1997.
- [8] 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.
- [9] NLANR, "Auckland-II trace archive," <http://moat.nlanr.net/Traces/Kiwitraces>, 2000.
- [10] S. Sarvotham, R. Riedi, and R. Baraniuk, "Connection-level analysis and modeling of network traffic," Tech. Rep., ECE Dept., Rice Univ., July 2001.
- [11] LBL, "Internet traffic archive," <http://ita.ee.lbl.gov/html/traces.html>.
- [12] K. Fall and K. Varadhan, "ns notes and documentation," <http://www-mash.cs.berkeley.edu/ns>, 2000.
- [13] S. McCane and S. Floyd, "ns- network simulator," <http://www-mash.cs.berkeley.edu/ns>.
- [14] J. Lévy Véhel and R. Riedi, "Fractional Brownian motion and data traffic modeling: The other end of the spectrum," *Fractals in Engineering*, pp. 185–202, Springer 1997.