

Network and user driven alpha-beta on–off source model for network traffic

Shriram Sarvotham ^{a,*}, Rudolf Riedi ^b, Richard Baraniuk ^a

^a *Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005, USA*

^b *Department of Statistics, Rice University, Houston, TX 77005, USA*

Available online 8 January 2005

Abstract

We shed light on the effect of network resources and user behavior on network traffic through a physically motivated model. The classical on–off model successfully captures the long-range, second-order correlations of traffic, allowing us to conclude that transport protocol mechanisms have little influence at time scales beyond the round trip time. However, the on–off model fails to capture the short-range spikiness of traffic, where protocols and congestion control mechanisms have greater influence. Based on observations at the connection-level we conclude that small rate sessions can be characterized by independent duration and rate, while large rate sessions have independent file size and rate. In other words, user patience is the limiting factor of small bandwidth connections, while users with large bandwidth freely choose their files. We incorporate these insights into an improved two-component on–off model—which we call the alpha-beta on–off model—comprising an aggressive alpha component (high rate, large transfer) and passive beta component (residual). We analyze the performance of our alpha-beta on–off model and use it to better understand the causes of burstiness and long-range dependence in network traffic. Our analysis yields new insights on Internet traffic dynamics, the effectiveness of congestion control, the performance of potential future network architectures, and the key parameters required for realistic traffic synthesis.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Alpha-beta on–off model; Network traffic; Long-range dependence

1. Introduction

The goal of this work is to develop a physically motivated model for the traffic at some point in a network based on observations at the connection level. By connection-level, we mean all the packets that have the same source and destination IP

* Corresponding author.

E-mail addresses: shri@rice.edu (S. Sarvotham), riedi@rice.edu (R. Riedi), richb@rice.edu (R. Baraniuk).

URL: <http://www.spin.rice.edu> (S. Sarvotham).

addresses, port numbers and protocol. Connection level measurements contain much more information than the simple aggregate; we study the statistics of session parameters such as duration, file size and rate, and their impact on aggregate traffic. We show that traffic can be roughly divided into two types: *alpha* connections (high rate, large transfer) and *beta* connections (remaining).

The most widely used traffic model based on connection-level information is the classical on–off model [1,2]. The on–off model successfully captures the second-order correlations of traffic, in particular their long-range dependence (LRD). The on–off model models traffic as a superposition of a large number of on–off sources, with heavy-tailed on and/or off periods. A conclusion commonly drawn based on the on–off model is that at large timescales (beyond several RTTs), transport protocol mechanisms like TCP do not have a significant impact on the network traffic [3]. Unfortunately, however, the on–off model fails to capture several key aspects of the network traffic, including its spikiness at small timescales.¹

Our goal in this paper is to use a new traffic analysis technique—*alpha-beta analysis* [4]—to develop a physically motivated traffic model that accurately matches both the large and small time scale behavior of network traffic yet retains the simplicity of the on–off model. We call our model the *alpha-beta on–off model*.

The alpha-beta traffic analysis is motivated by the observation that small scale spikiness comes from heterogeneity in network resources. The analysis involves decomposing the network traffic into two components according to the connection bandwidth: the large bandwidth alpha component contributes all the small scale spikes and is highly non-Gaussian, while the small bandwidth beta component contributes most traffic and is Gaussian. This decomposition explains the LRD of network traffic at very large time scales as well as the highly non-Gaussian marginals and multi-fractal scaling on small time scales. The alpha-beta analysis parsimoniously accounts for user behavior, net-

work topology, and the heterogeneous distribution of network band-widths (see Fig. 5).

The alpha-beta analysis indicates that the bursts in traffic are not due to a “conspiracy” of connections [4]. This is a significant finding, because the classical on–off model [1,2] produces bursts only by having a large number of on-sources active at the same time. Hence, the on–off model should be extended to accommodate small-scale spikes.

We will analyze two different traffic data sets (see Appendix A for information on the datasets) to show that the key parameters of the on–off model are different for the alpha and beta components. This stems from the difference in the dynamics that give rise to alpha and beta components. Alpha flows have high bandwidth and small RTTs, and hence typical connection durations are small and not a concern. Beta flows are constrained by a bottleneck bandwidth, and transfer of large files translates to a large latency, which a user may find undesirable. Our alpha-beta on–off model uses this intuition to model the on–off parameters differently for the alpha and beta components.

This paper is organized as follows. In Section 2, we briefly describe the on–off model. Section 3 gives an overview of our previous work on alpha-beta traffic decomposition and summarizes the properties of the alpha and beta components of traffic. In Section 4, we set the stage for our alpha-beta on–off model by identifying the components in the real traffic that make up the on and off periods. Section 5 analyzes the real traffic to verify our claim of how the network and the user behavior affect the on–off parameters. In Section 6, we present our alpha-beta on–off traffic model and use it to synthesize traffic. Section 7 describes the implications of our traffic model in the real Internet. After an overview of the related work in Section 8, we conclude with a summary of our work and future work in Section 9.

2. The on–off model

The on–off model, hinted in [5] and formally introduced in [2] provides a foundation for traffic modeling based on physically verifiable properties. The model posits that network traffic is a super-

¹ By small scale spikiness, we refer to the deviation of the distribution from Gaussianity.

position of a large number of independent on–off sources with the on and off periods drawn from heavy tailed distributions. (A random variable X has a heavy tailed distribution if the asymptotic shape of the distribution is hyperbolic: $P[X > x] \sim x^{-\alpha}$, as $x \rightarrow \infty$, $0 < \alpha < 2$.) This gives rise to self-similarity in the aggregate process—a fractional Gaussian noise (fGn) process [1]—whose LRD is determined by the heavy-tailed nature of the on and off periods. The key parameter in the description of LRD is the *Hurst parameter* H that captures the long-term correlation structure. The case $H = 0.5$ corresponds to the absence of LRD. In the case $0.5 < H < 1$, LRD occurs. For an overview of LRD, refer to [2].

The on–off model has its roots in a certain renewal reward process introduced by Mandelbrot [6] (and further studied in [7]) and provides the theoretical underpinning for much of the recent work on network traffic modeling. Park et al. [8] have shown that the application layer property of heavy-tailed file sizes is preserved by the protocol stack and mapped to approximately heavy-tailed busy periods at the network layer.

The factual basis of heavy tailed on-periods has been shown in [8]. The corresponding empirical basis for heavy tailed file sizes in UNIX file systems, whose transport may be the cause of heavy tailed on-periods in packet trains, is shown in [8]. The on–off model is able to create both fGn and asymptotic second-order self similarity.

3. Alpha-beta traffic analysis

In this section we survey our previous work [4] on alpha-beta traffic analysis.

3.1. Alpha and beta components of traffic

Connection-level information enables us to conduct a refined analysis of traffic bursts. In aggregate traffic models (including the on–off model), traffic bursts can only arise from a large number of connections transmitting bytes or packets simultaneously. That is, bursts stem from a “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, [4] demonstrates that this is not the case. Bursts in the sequence of bytes-per-time-bin generally do not coincide with large values in the sequence of connections-per-time-bin.

In [4] we proposed 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.

1. In each T -second time bin, tally the contribution of bytes coming from each active connection.
2. If the contribution of a connection is greater than a threshold, then label it as an *alpha* connection.² Section 6 presents some suggestions on choosing the threshold. The aggregate of all the alpha connections gives the alpha component of traffic.
3. The *beta* component is obtained by subtracting the alpha component from the total trace.

Our procedure thus decomposes the total traffic trace into alpha traffic + beta traffic, as shown in Fig. 1. The statistical properties of the two components can be summarized as follows:

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

Beta traffic: At time-scales coarser than the round-trip time (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.

See [9] 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.

² By analogy with the dominating *alpha* and submissive *beta* members of the animal kingdom.




Fig. 1. Decomposition of a traffic trace into the sum of an alpha component and beta component.

3.2. Origins of alpha and beta components

We now show that the alpha connections come from large file transfers over large bottleneck bandwidths.

Our first and most striking observation is the grouping of alpha connections according to their source-destination host pairs. We collect all connections with the same source and destination hosts into an *end-to-end (e2e) group*. If random effects of the networking protocols were causing spikes, then we would find the dominating alpha connections distributed randomly over the e2e groups. The above-mentioned clustering indicates that the opposite is true.

More strikingly, our second observation is that if an e2e group contains *one* alpha connection, then *all* connections in that group are either alpha or too short to cause spikes. This demonstrates clearly that spikes are caused by large volume connections over particular (necessarily high-bandwidth) e2e paths.

An observation based on a more refined analysis provides a potential *physical reason* for the origin of alpha spikes. 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 . Fig. 2 displays the total load versus peak rate (for $T = 500$ ms) for all connections within four particular e2e 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 e2e paths corresponding to groups 1 and 2 have a high bottleneck bandwidth. On the contrary, the connections in groups 3 and 4 are obviously not getting as much bandwidth as they could consume. None of these connections are alpha. With such a consistent picture, we can exclude causes such as re-routing as the main cause for spikes, since it would be highly unlikely that all connections in a group would be systematically affected.

It is instructive to compare the elephant-mice traffic decomposition [10] and the alpha-beta traffic decomposition. Elephants and mice are based on the file size—a source property. Alpha and beta are based on the peak rates of connections, which depend both on the user behavior (we need large files to create large peak rates) as well as network characteristics (we require large bandwidths to accommodate large peak rates).

We conclude that the majority of spiky connections are due to *large file transfers over large bottleneck bandwidth e2e paths*. This summarizes our previous work on alpha-beta analysis. In the next sections, we motivate and present our alpha-beta on-off model.

4. Towards an alpha-beta traffic model

To motivate our model design, recall that the mechanisms that control the alpha and beta traffic components are different. In particular, *beta connections are rate-limited, whereas the alpha connec-*

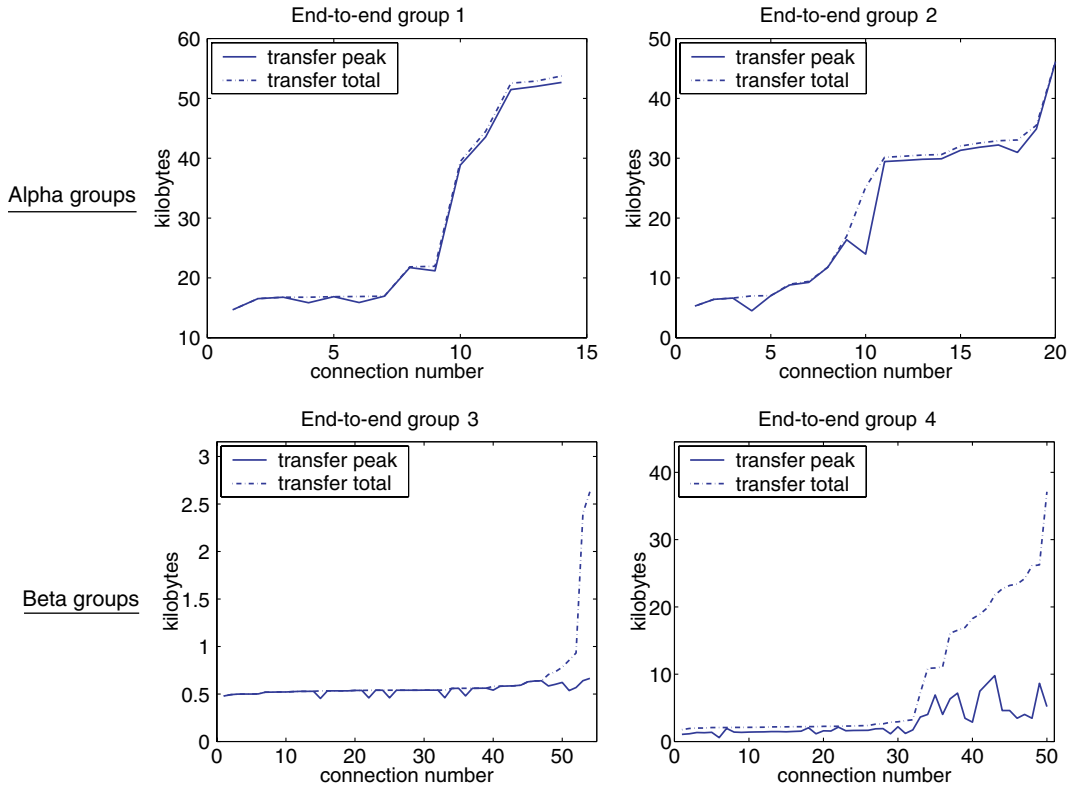


Fig. 2. Plot of peak rate (at 500 ms) and total transfer for all connections which share the same pair of source and destination hosts. The connections are ordered in non-decreasing values of total transfer. Note the high peak rates in the top row (alpha connections). Note the starvation of connections in the groups in the bottom row and the overall low peak rate (beta connections).

tions are not (as seen from Fig. 2). The goal of this section is to understand the key parameters for the on-off model for the alpha and beta components.

Aiming at a physically motivated model of the on-off type, we first search for the most appropriate real world equivalent of an on-period in the model. The choices for an on-period are enumerated below:

- *Connection*: The byte-stream belonging to the same connection (5-tuple of source and destination IP addresses and port numbers and the protocol),
- *e2e stream*: The byte-stream belonging to the same source and destination IP addresses, and
- *e2e session*: We define an e2e session as a cluster of bytes belonging to the same e2e stream, with the following constraints: the delay between two successive packets in an e2e session is always

less than a threshold, and the delay between packets belonging to different e2e sessions is greater than the threshold. For our analysis and modeling, we set this threshold to 2 s. This definition is motivated by [2], where the authors identify off-periods in an e2e stream by looking at time instants $t \geq 2$ s that does not contain any packet. An e2e session can have packets from many connections. We assume that the session consists of several connections on the same e2e path separated only by a time interval caused by the transport and lower-layer protocols and the network, not the user. Sessions are separated by longer, human “think” times.

In this section, we show that the most appropriate mapping for an on-period is the e2e session.

In order to identify the most appropriate on-period, we first consider the independence

property of the on–off sources. In the on–off model, the sources are assumed to be independent, and the correlation in the aggregate traffic comes only from the distribution of the on-times.³ Further, the on and off periods are independent (Table 1).

To check the independence assumption, we shuffle the start-times of the presumed on-periods and check whether the modified trace has the same correlation structure. The start-times of the groups of bytes are changed and assigned a random variable chosen uniformly over the duration of the trace, with wrap-around if the on-period extends beyond the end of trace. Shuffling eliminates any correlation in the trace that appears due to correlated start times of the sources.⁴

Table 4 shows how shuffling the components of the trace affects the Hurst parameter of the aggregate traffic. If the start times were independent, then the resulting trace would have the same Hurst parameter. In our analysis, we use the variance-time plot to infer the Hurst parameter. From Table 4, we note that the connections have correlated start times whereas the e2e streams and e2e sessions do not. This correlation could come from the fact that each Web session can spawn many connections whose start times are close together. This indicates that the connections are not good candidates for on-periods.

Let us consider the choice of e2e streams as on-periods. An on-period that is based on an e2e stream will also cover large time durations when no packets were transmitted between the constituent e2e sessions. This results in a reduced instantaneous rate for the on-period and reduces the variance of the aggregate trace sharply. To illustrate this point, we average out the bytes in each on-period, spreading them evenly across the duration of the on-period. Table 2 shows the results of such a modification. Clearly, an e2e stream is not a good choice for the on-period, since the variance (as depicted by the ratio σ/m , where σ is the stan-

Table 1

Changes in the Hurst parameter when we shuffle the on-periods for connections, e2e streams, and e2e sessions

Trace	Auck (H)	Bell (H)
Real trace	0.82	0.79
Shuffled connections	0.66	0.68
Shuffled e2e streams	0.82	0.78
Shuffled e2e sessions	0.82	0.78

Clearly, the connections have correlated start times, as the reduction in the Hurst parameter reveals.

dard deviation of the aggregate trace and m is the mean) of the resulting trace is drastically reduced.

Let us consider the choice of e2e sessions as on-periods. The resulting variance (as seen from Table 2) for the e2e session more closely matches the real trace, and hence is most appropriate to model the on-period. In the remainder of this paper, we assume that the on-period corresponds to an e2e session.

For an on-period we have file size = rate \times duration. Hence we have only two degrees of freedom; once we specify two of the parameters, the third is automatically determined. We now consider the likely independent parameters for alpha and beta connections.

Large file transfers via beta connections are of long duration due to the ceiling on the transmission rate (as shown in Fig. 2). Hence the user's latency tolerance creates a bias against large file transfers. For alpha connections, the latency is usually small for most file sizes, and hence there is no user-bias against large files. Hence we argue that for the beta component, the key parameters are the rate (network bottleneck rate) and the duration (determined by the user behavior); the file size is set by the available rate and the duration.

Table 2

Changes in the standard deviation when we average out (flatten) the on-periods for connections, e2e streams, and e2e sessions

Trace	Auck (σ/m)	Bell (σ/m)
Real trace	0.38	0.72
Flattened connections	0.21	0.49
Flattened e2e streams	0.095	0.22
Flattened e2e sessions	0.32	0.65

It can be seen that the resulting standard deviation for the e2e sessions matches that of the resulting trace best.

³ LRD in the aggregate traffic can be generated using heavy-tailed off-periods as well.

⁴ Shuffling also removes any correlation in the byte arrival process due to sources sharing a common bottleneck link. However, we ignore the effect of this phenomenon in this paper.

On the other hand, it is more appropriate to model the alpha component using independent file-size and the rate, since the duration is not of concern to users, as it is usually small. The duration gets set by the other two parameters.

We now strengthen the above arguments by performing a pairwise correlation analysis on the file size, duration, and rate of the connections for some real world traffic data sets.

5. Joint distributions of size, rate, and duration

For our analysis and modeling, we use two real traffic traces—the Auck trace and the Bell trace. Please refer to the [Appendix A](#) for a summary of the traces.

First we consider the alpha component. [Figs. 3 and 4](#) show pairwise scatter-plots of the size, rate and the duration of the sessions in a log scale (we use the log scale because each of these parameters spans several orders of magnitude). For this analysis, we ignore sessions with duration less than 100 ms, since the rate of a very small session is ill-defined and problematic to handle (see also [\[11\]](#)). As observed in previous work [\[11,12\]](#), there exists a strong correlation between the rate and the session-size for the overall traffic. This indicates that the users choose the file sizes to download based on the speed of the connection. However, the correlation between session-size and the rate is greatly reduced when we look only at the alpha sessions. Hence we conclude that for the alpha sessions, the download time, which is small, is not a significant factor that limits a user's behavior. In other words, for alpha sessions, we may assume that the session rate and the session size are independent for the purpose of simplified modeling and synthesis.

Next we consider the beta component and the overall traffic. Looking at [Table 4](#), we note that at first sight, all of the three parameters are correlated pairwise. We argue that the correlation between the rate and duration comes from the constraint that the packet sizes are at least 40 bytes long (consisting of the TCP/UDP and IP headers). Looking at the scatter-plots of the rate and duration, we see that there exist no samples in the low-rate, low-duration regime (lower triangle).

This is because the minimum session size is 80 bytes (corresponding to a session of just two minimum-sized packets), and hence a combination of low rate and low duration cannot occur together in reality.

In a gedanken experiment we could lump sessions with duration less than 2 s together into a third component. This component would constitute simple noise due to the small size of these sessions. For the Bell trace, this component would contribute only 23% of the bytes of the traffic (18% for the Auck trace). More importantly, the remaining beta component would show a convincingly small correlation between rate and duration of -0.12 for the Bell trace (-0.05 for the Auck trace). We argue that the concepts of rate and duration do not apply to these small sessions and that some other mechanism than user behavior shapes them. However modeling this part of traffic as another component adds undesirable model complexity; incorporating them into the beta component appears to be appropriate within our desired model accuracy as our simulations show. We leave the study of the parameters of small sessions for future work.

To summarize, our observations are consistent with the view that for beta sessions the size of its connections depends on the bottleneck bandwidth (rate) and the duration of the download that the user is willing to tolerate. In our synthesis, we will thus assume that the rate and duration are independent for the beta sessions. For alpha sessions, we assume that the file size and rate are independent.

6. The alpha-beta on-off model

We now exploit the results of the last section to define a new two-component on-off model aligned with the alpha-beta analysis. We call our model the *alpha-beta on-off model*.

6.1. Alpha-beta separation

First we describe the procedure to classify connections as alpha and beta. The alpha-beta principle is that a few outliers (in terms of sessions with large peak-rates) are responsible for a very visible phenomenon, namely non-Gaussianity of

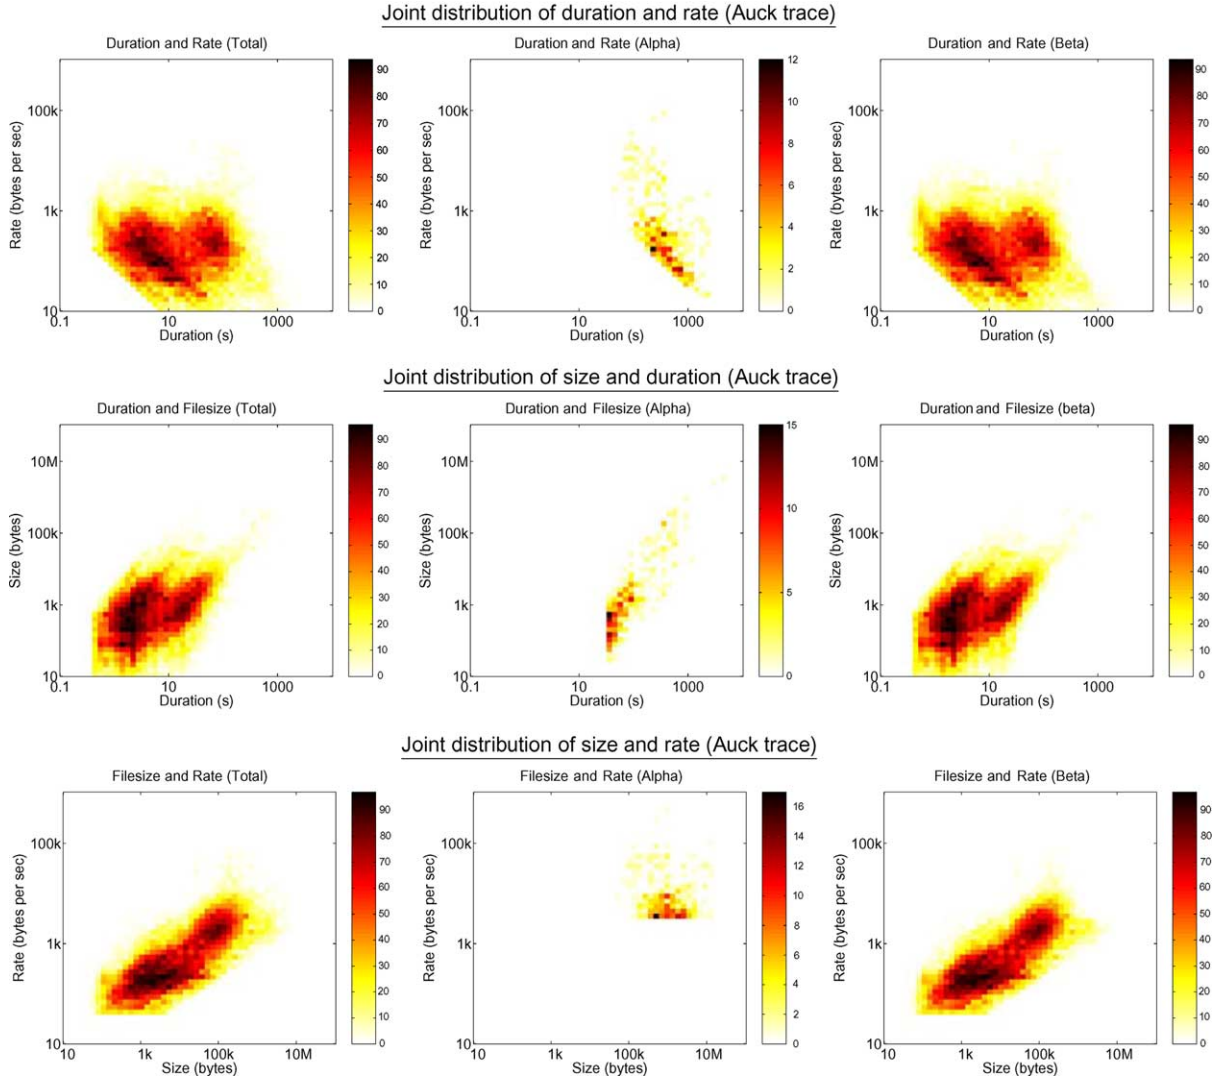


Fig. 3. Scatter-plot depicting the pairwise joint distribution of session rate, duration, and size for the Auck trace.

the traffic trace at certain small scales. There are many possible schemes to define alpha sessions and separate these outliers.

As described in Section 3, in one such scheme we compute the maximum contribution of bytes from every session in a time window of T seconds. We classify as alpha sessions those that contribute more bytes than the *alpha-threshold* in any time window of T seconds. The remaining sessions are classified as beta. This scheme relies on the peak-rate of the sessions. A second, robust separation scheme sets

the alpha-threshold to remove the sessions with largest peak-rates until the distribution of the peak-rates of the remaining sessions (excluding mice sessions [10]) resembles a Gaussian distribution. In many experiments, we have found that the results are not sensitive to the exact alpha-threshold.

6.2. Alpha-beta modeling

Alpha-beta on-off model consists of the sum of two on-off models: one for the alpha component

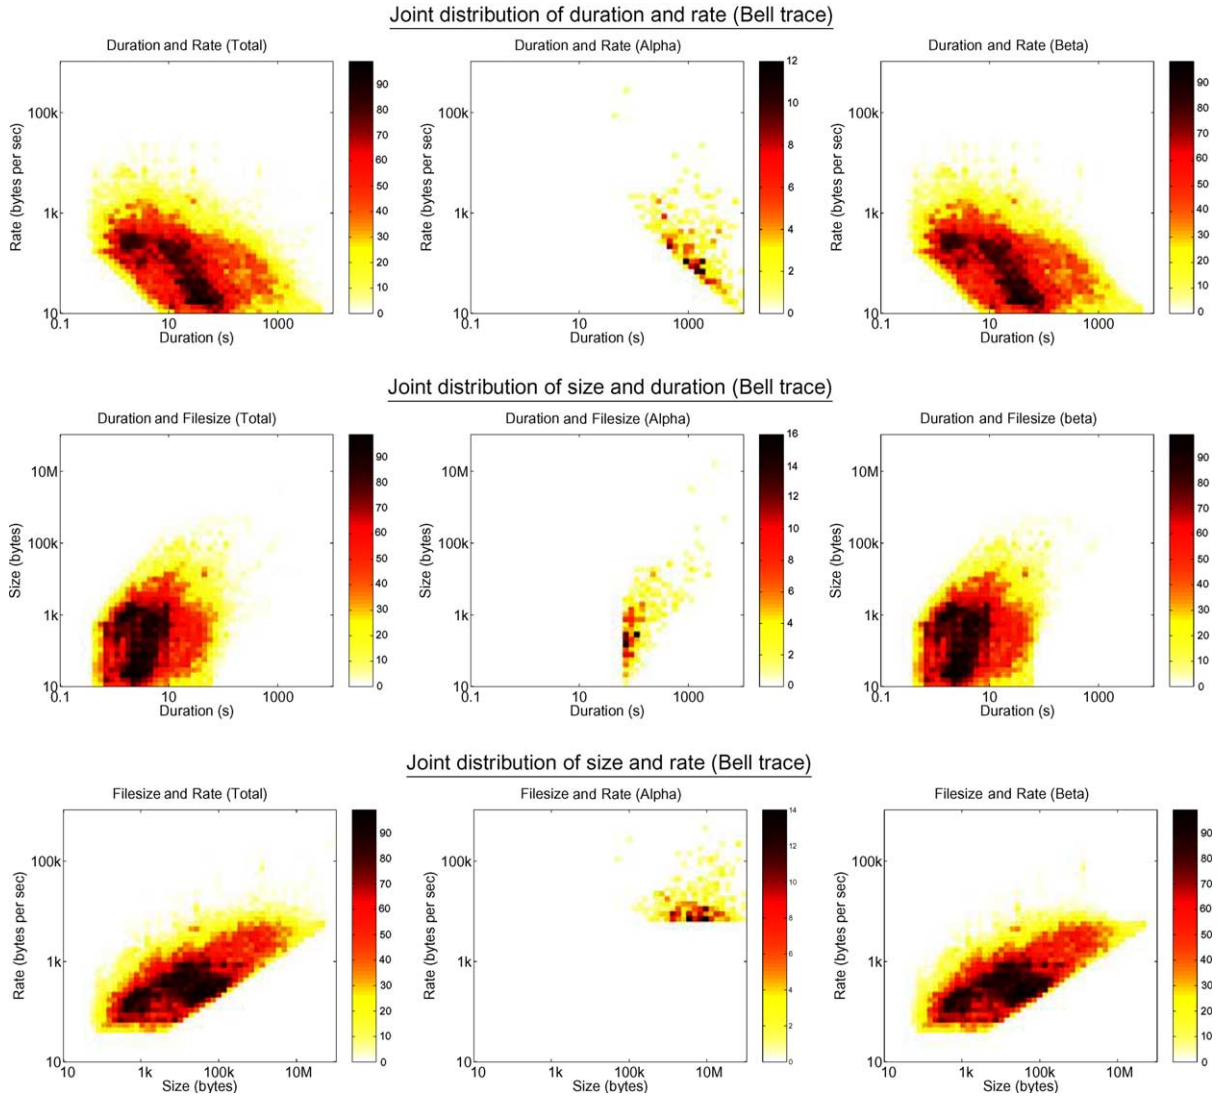


Fig. 4. Scatter-plot depicting the pairwise joint distribution of session rate, duration and size for the Bell trace.

Table 3
Correlation coefficient between the on-off parameters (Auck trace)

Component	Duration and rate	Duration and size	Size and rate
Total	−0.11	0.52	0.79
Alpha	−0.61	0.87	−0.13
Beta	−0.12	0.52	0.78

Table 4
Correlation coefficient between the on-off parameters (Bell trace)

Component	Duration and rate	Duration and size	Size and rate
Total	−0.46	0.29	0.72
Alpha	−0.61	0.73	0.10
Beta	−0.46	0.28	0.71

and one for the beta component. Interestingly, the three controlling parameters duration = file size/rate are used in different combinations for each of the two models.

We model the alpha sessions as having on-periods with independent size and rate. Random sizes and rates are chosen independently from their respective empirical marginal distributions in the alpha connections of the real trace. The duration of the on-period is determined by duration = file size/rate.

We model the beta sessions as having on-periods with independent duration and rate. Random durations and rates are chosen independently from their respective empirical marginal distributions in the beta connections of the real trace. The file size of the on-period is determined by file size = rate \times duration.

We choose the number of alpha and beta sessions to match that of the real trace. The starting times of the alpha and beta on-periods are chosen uniformly over the duration of the trace, with wrap-around if the on-period exceeds the total time of the trace. The synthetic alpha trace is the superposition of the alpha on-periods and the synthetic beta trace is the superposition of the beta on-periods. The total trace is the sum of these two superpositions.

6.3. Justification

We now justify our choices for the controlling parameters of the alpha and beta on-off models by matching real traffic traces.

For each of the Bell and Auck traces, we first computed the (1-d) empirical distribution of the log of the three parameters size, rate and duration for each of the three traffic components: alpha, beta and the total traffic. (We used histograms of 2000 bins; the results were not sensitive to the exact number.) We synthesized on-periods by choosing two of the parameters as independent from among the three possible schemes: scheme RD assumes that the rate and the duration are independent; scheme SD assumes that the size and duration are independent; and scheme SR assumes that the size and rate are independent. The three schemes are summarized in Table 5. We then

superpose these on-periods to generate the synthetic traffic. We choose the number of on-periods in the synthesis to be exactly the same as the number of e2e sessions in the real trace.

The results of the synthesis are summarized in Tables 6 and 7. The synthetic traces for the Bell trace are shown in Fig. 5 for a visual comparison. Note that the beta component of traffic is well matched by scheme RD and the alpha component by Scheme SR, as predicted. Also it is clear that none of the three schemes match the total traffic, which is not surprising when we note that the three parameters are mutually dependent for the overall trace (Tables 3 and 4). That is, no single-component on-off model is capable of matching the total traffic. However, the alpha-beta on-off model matches the real trace well.

Fig. 6 shows the synthesized total traffic using the alpha-beta on-off model. Here we synthesize the total traffic as the sum of the alpha component generated using scheme SR and the beta compo-

Table 5
Independent parameters in the schemes to generate traffic

Scheme name	Independent parameters
Scheme RD	Rate and duration
Scheme SD	Size and duration
Scheme SR	Size and rate

Table 6
Auck data synthesis summary

Time series under consideration	Ratio σ/mean	Hurst parameter H	Kurtosis
Real total	0.38	0.82	5.79
Real alpha	1.05	0.79	15.07
Real beta	0.36	0.85	2.997
Scheme RD total	0.38	0.78	4.85
Scheme RD alpha	0.79	0.83	4.96
Scheme RD beta	0.36	0.78	4.24
Scheme SD total	3.64	0.51	1691
Scheme SD alpha	10.34	0.58	1981
Scheme SD beta	2.41	0.52	742
Scheme SR total	0.19	0.83	11.20
Scheme SR alpha	0.81	0.83	7.61
Scheme SR beta	0.19	0.84	7.00
Alpha-beta on-off model (RD beta + SR alpha)	0.42	0.83	5.80

The table shows the comparison of the statistics of the real and synthetic traces at 500 ms timescale.

Table 7
Bell data synthesis summary

Time series under consideration	Ratio σ/mean	Hurst parameter H	Kurtosis
Real total	0.71	0.79	10.57
Real alpha	2.06	0.81	15.74
Real beta	0.47	0.71	3.96
Scheme RD total	0.44	0.84	4.81
Scheme RD alpha	1.04	0.86	8.26
Scheme RD beta	0.44	0.85	4.73
Scheme SD total	4.58	0.49	1793
Scheme SD alpha	11.65	0.53	3476
Scheme SD beta	3.02	0.51	608
Scheme SR total	0.40	0.73	14.93
Scheme SR alpha	1.33	0.82	14.89
Scheme SR beta	0.43	0.66	16.94
Alpha-beta on-off model (RD beta + SR alpha)	0.55	0.84	10.43

The table shows the comparison of the statistics of the real and synthetic traces at 500 ms timescale.

nent generated by scheme RD. Observe that the synthetic trace generated by this model visually resembles the real trace. The closeness of the energy plots shown in Fig. 8 for the real and synthetic trace further indicate that the synthetic trace is a good fit.

Another quantitative metric to compare the real and synthesized traces is the queue length distribution as shown in Fig. 7. In this analysis, we feed the trace into an infinite queue with service rate equal to the link capacity of the link where traffic was recorded (9 Mbps for the Bell trace) Fig. 7(a) shows the complementary queue distribution of the real trace and the synthetic trace generated by our alpha-beta on-off model. The synthetic trace based on the alpha-beta on-off model is generated by superposing the alpha component generated by scheme SR and the beta component generated by scheme RD (same as the trace is shown in Fig. 6(b)). Note that the synthetic trace from our model matches the queuing behavior of the real trace well, with increasing error for larger queue sizes. However, when we model the total traffic using a single-component on-off model with the schemes RD, SD and SR, none of the synthetic traces match the queuing behavior of the real trace, as seen from Fig. 7(b). This convincingly shows the value of alpha-beta decomposition, and the importance of treating the alpha and beta components separately in the on-off model.

7. Discussion

The implications of our alpha-beta on-off model to congestion control and what-if-scenarios are relevant and even somewhat surprising.

The network-driven user. Let us start by drawing a telling relationship between network and the user layer. Our evidence suggests separating user behavior into two regimes. We base the discussion on the three basic parameters of file size, sending rate, and duration.

The beta regime, on the one hand, is characterized in the network layer by rate-limitation, which results in the user layer as *patience* being the apparent parameter ruling behavior. For the network-user-driven on-off model this translates into approximating the joint distribution of file, size and rate as being characterized by independent rate and duration.

The alpha regime, on the other hand, corresponds in the network layer to high available bandwidth, which translates in the user layer into a free choice of file transfers according to *interest*. For the model this means that file size and rate are independent.

What-if? scenarios. In earlier work [9], we showed that most of the alpha sessions come from short RTT connections. By looking at the histogram of the connection RTT, a network operator can immediately comment on the relative strengths of the alpha and beta components. The relationship between the connection RTTs and the strength of the alpha component has some important implications when we consider Content Distribution Networks (CDNs). CDNs change the RTT distributions of sessions, and we expect more flows to have shorter RTTs. In related earlier work [4], we showed that alpha connections tend to build up large queues. This points to a future need for a mechanism at routers (AQM) to detect and control alpha connections so that other connections are not unduly affected. This is particularly relevant as we expect more alpha flows in the future, due to higher available bandwidth to users, thanks to digital subscriber line (DSL), cable modem, and fiber-to-the-home technologies.

Another interesting aspect is the potential of alpha-DOS attacks where hosts close to a server

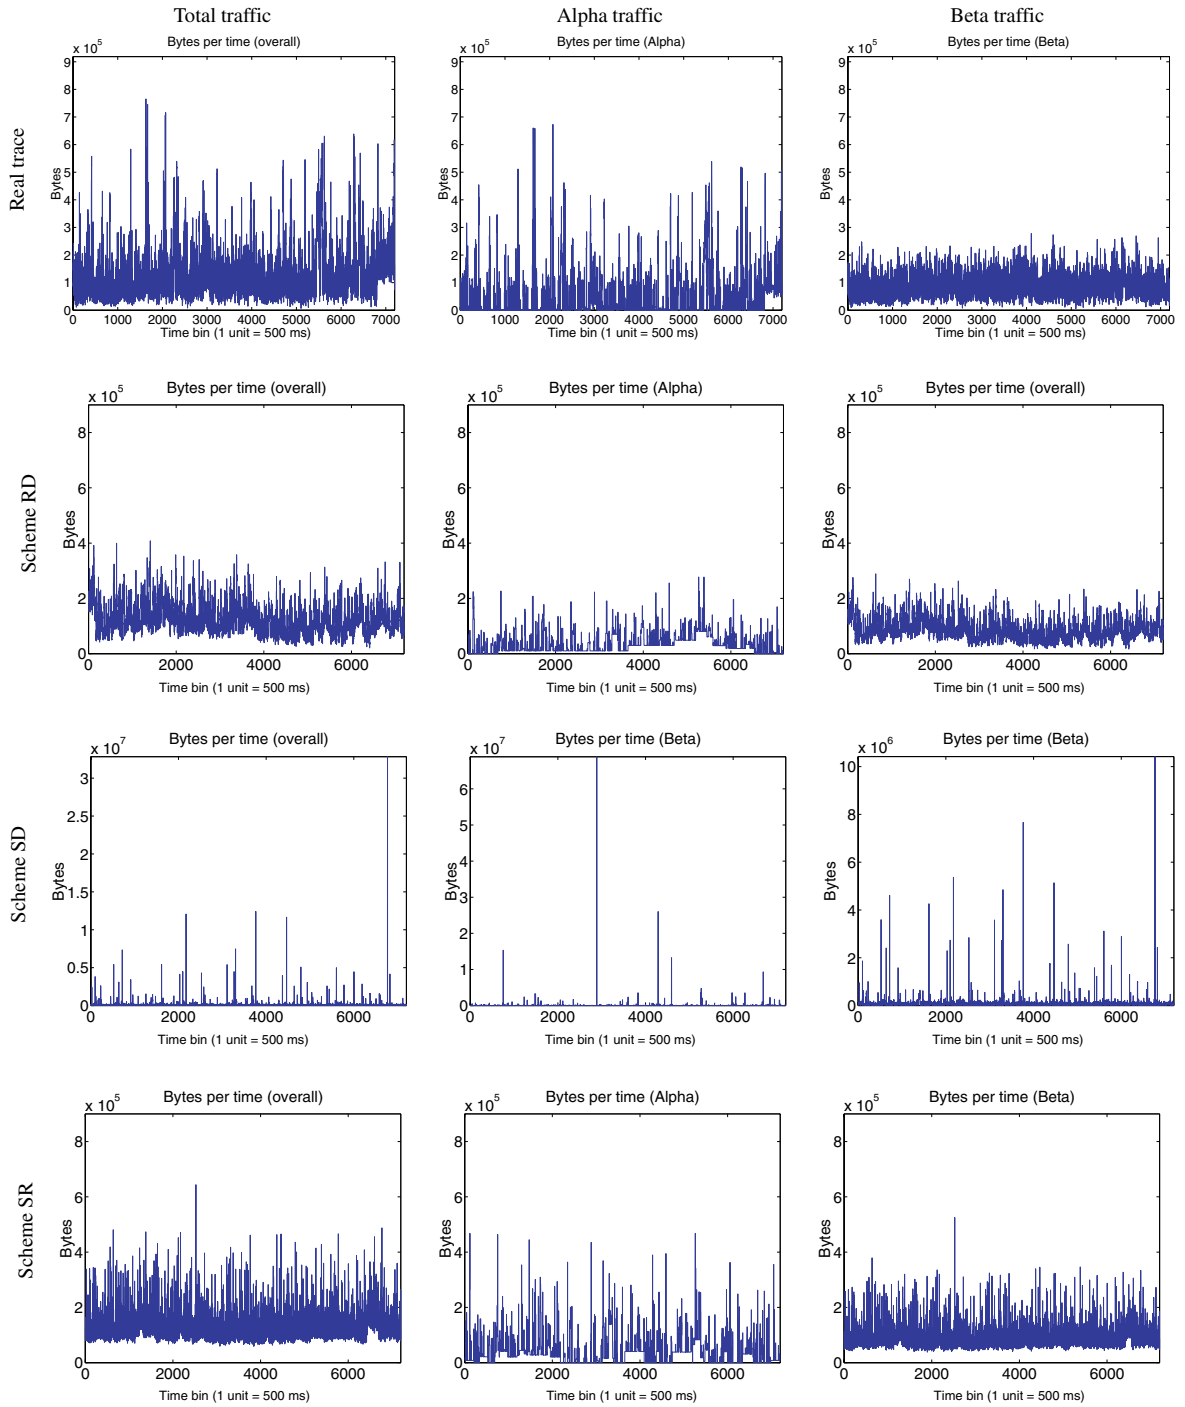


Fig. 5. Synthesis of total, beta and alpha components of the Bell trace. The first row depicts the real trace. The plots in the remaining rows show the synthetic traces obtained by using Scheme RD, Scheme SD and Scheme SR. Note that the total traffic is not captured well by any of the schemes. The beta traffic is captured well by scheme RD, while Scheme SR captures the alpha traffic well. This observation is reinforced in the Table 7.

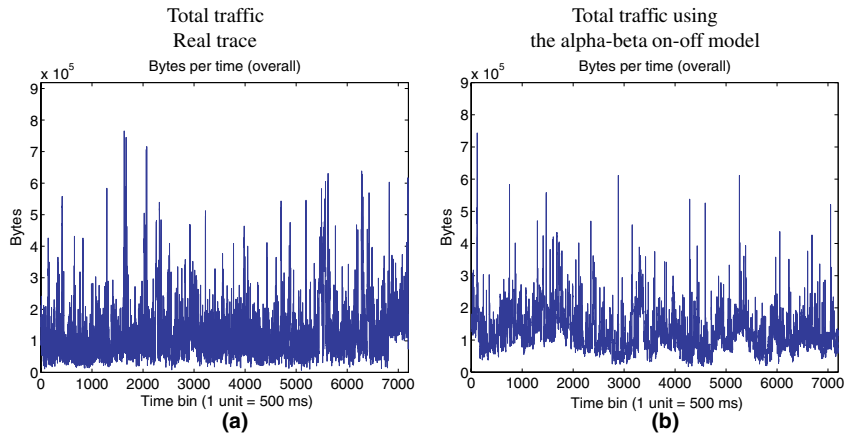


Fig. 6. Synthesis of total traffic using our alpha-beta on-off model. The alpha component is generated using scheme SR and the beta component is generated with scheme RD. The alpha and beta components generated are then superposed. The figures show that the synthetic trace resembles the real trace. Note that none of the single-component on-off schemes shown in Fig. 5 match the total traffic visually.

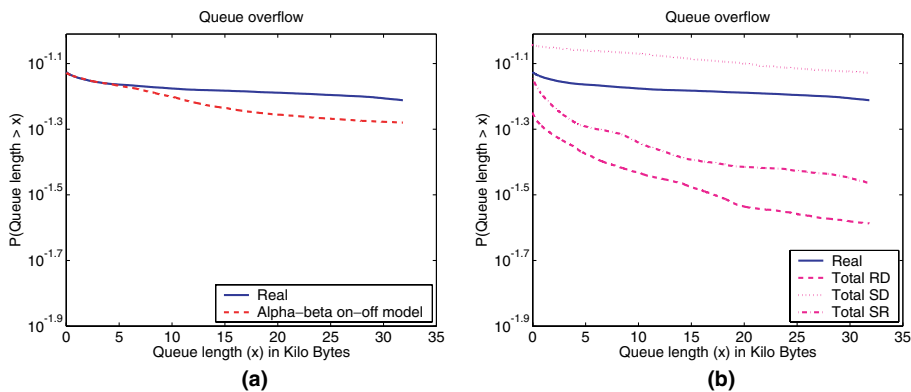


Fig. 7. Figure (a) shows the complementary queue distribution of the real trace and the synthetic trace that is generated using our alpha-beta on-off model. Note that the synthetic trace matches the queuing behavior of the real trace well. Figure (b) shows the complementary queue distribution of the real trace and the synthetic traces generated by schemes RD, SD and SR on the real trace, without doing an alpha-beta decomposition. Note that none of the three synthetic traces in (b) match the queuing behavior of the real trace accurately. This reveals the importance of the alpha-beta separation and the need to treat alpha and beta components differently for the on-off model.

leverage small RTTs to drain the network resources (bandwidth), depriving other clients of service. One way to detect such an attack would be to identify transfers in the alpha regime that persist beyond the length of time a typical human user would tolerate.

Effectiveness of congestion control. Among the two components, only the beta component appears to be effectively controlled by TCP. Alpha connections profit greatly from small RTTs to

obtain large bandwidth. Operating as it does with small TCP windows, loss-based congestion control does not restrict alpha connection since they recover much more quickly to drops than the average beta user. This calls for alternative methods of congestion control than that provided by TCP.

Predicting traffic patterns in a high-speed TCP environment. An interesting question is to predict a priori the expected volumes of alpha and beta components of traffic in a high speed TCP

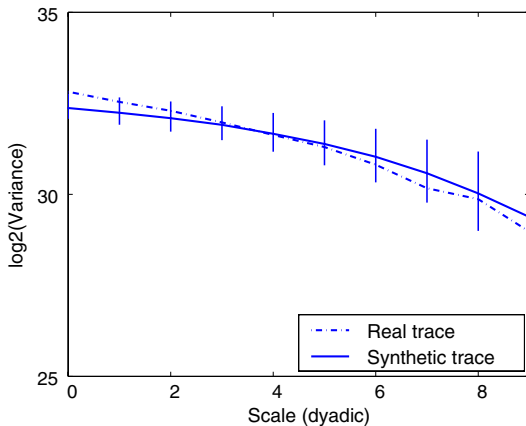


Fig. 8. Variance–time plot comparison of the real and synthetic traces.

environment. Many high-speed TCP schemes have been proposed recently to adapt to high bandwidth-delay product environments such as HSTCP [13], STCP [14], XCP [15], FAST-TCP [16] and BI-TCP [17]. For TCP in a high bandwidth-delay environment, we predict that the strengths of the alpha and beta components will depend on TCP's bias toward short RTT connections. In the current TCP-Reno, alpha connections cannot arise from large RTT connections because they lose out to short RTT flows due to TCP bias. [17] presents a detailed analysis of the RTT bias for several high-speed TCP schemes. Based on these results we expect that the alpha-beta differentiation between flows will be more pronounced for STCP and HSTCP (which have large RTT bias) than for BI-TCP (which has RTT bias comparable to TCP-Reno).

8. Related work

Zhang et al. [11] have investigated the origins of Internet flow rates, and studied their relationship with the size and duration of the flow. They report the correlation between the three parameters (for the overall set of connections), and their results are consistent with our findings for the Auck and Bell traces. In particular, they observe a strong correlation between size and rate. The authors conclude that the users choose the size of their

transfer based strongly on the available bandwidth, which is in agreement with what we see for beta sessions (recall that beta sessions constitute the bulk of the traffic load).

Lan et al. [12] have explained the strong correlation between the file size and the rate as arising due to protocol interactions for small and medium flows. Their work also analyzes the relationships between different classifications of network traffic connections, such as elephant–mice [10], tortoise–dragonfly [18], and alpha-beta [4].

Campos et al. [19] have applied multivariate extreme value theory to show that the file size and throughput (the inferred rate at which the file is transferred) exhibit extremal independence. In other words, for large file transfers, the throughput tends to be more closely related to time, and essentially independent of response size. This result agrees with our analysis that for alpha sessions, the rate and file size are essentially independent (note that the alpha flows have both large file sizes and high rates). The authors conclude that a significant percentage of the longest durations are due to low throughput rates rather than large file transfers.

In [20], Campos et al. address and study the heterogeneity of Internet traffic by combining an abstract model of application-level communication with statistical cluster analysis. Their model captures the pattern of data exchange in an abstract manner for a wide range of applications. The main idea is to cluster the connections together that are dependent if they come from the same session.

9. Conclusions and ongoing work

We have demonstrated the impact of the relationship between the session parameters (duration, rate and size) on network traffic. The user-behavior (how much delay a user is willing to tolerate) and the network characteristics (bottleneck bandwidth) determine the relationship between the parameters. The beta component of traffic is influenced by the rate and duration, since the beta sessions are rate-limited and users cannot tolerate large durations in download times. The alpha component is governed primarily by independent session size and rate, since its duration is negligible. Based on these

observations we have proposed the alpha-beta on–off model for Internet traffic. The synthetic traffic generated using our model captures the characteristics of the real trace in terms of the queuing behavior and the energy plots.

For future work, we plan to validate our alpha-beta on–off model on a wider variety of traces. We will perform theoretical analysis on the distribution of parameters using techniques such as Massey and Whitt’s generalization of Taylor’s theorem to probability distributions [21]. We are also working on finding parametric models for the alpha and beta components.

Acknowledgement

This work was supported by NSF grant ANI-0099148, DARPA/AFRL grant F30602-00-2-0557, DOE SciDAC grant DE-FC02-01ER25462, and Texas Instruments Leadership University Program.

Appendix A. Real traffic traces

The alpha-beta on–off model developed in this paper has been successfully tested on many real network traffic traces. However, for the purpose of illustration, we present our analysis on two representative traces. The summary of the traces are given in Table 8.

The first set of traces were recorded at the University of Auckland [22]. The traces are part of a large collection of GPS synchronized IP header traces captured with a DAG2 system at the

University of Auckland Internet uplink by the WAND (Waikato Applied Network Dynamics) research group, which is based in the University of Waikato Computer Science Department. The University of Auckland ITSS department operates an OC3 ATM link to carry a wide variety of services off the main campus. A single ATM channel is used to connect the university to the global Internet, and since it is the only connection, all packets for all external connections pass the measurement point. The connection has a packet peak rate of 2 Mbps in each direction.

The second trace is from the Bell Labs-I dataset [23], from a joint project between NLNR PMA and the Internet Traffic Research group lead by Bill Cleveland. The data set was collected with a Dag3.2E 10/100 MBit/s Ethernet card at the outside of the firewall servicing researchers at Bell Labs via a 9 Mbps link to the Internet.

References

- [1] M. Crovella, A. Bestavros, Self-similarity in World Wide Web traffic. Evidence and possible causes, *IEEE/ACM Trans. Network.* 5 (6) (1997) 835–846.
- [2] W. Willinger, M. Taqqu, R. Sherman, D. Wilson, Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level, *IEEE/ACM Trans. Network.* (Extended Version) 5 (1) (1997) 71–86.
- [3] W. Willinger, V. Paxson, R. Riedi, M. Taqqu, Long Range Dependence: Theory and Applications, Wiley, New York, 2002, chapter Long range dependence and Data Network Traffic.
- [4] S. Sarvotham, R. Riedi, R. Baraniuk, Connection-level analysis and modeling of network traffic, in: *Proceedings of the IEEE/ACM Internet Measurement Workshop*, 2001.
- [5] W. Leland, M. Taqqu, W. Willinger, D. Wilson, On the self-similar nature of Ethernet traffic, *IEEE/ACM Trans. Network.* 2 (1) (1994) 1–15.
- [6] B.B. Mandelbrot, Long-run linearity, locally Gaussian processes, H -spectra and infinite variances, *Int. Econ. Rev.* 10 (1969) 82–113.
- [7] M. Taqqu, V. Teverovsky, W. Willinger, Estimators for long-range dependence: an empirical study, *Fractals* 3 (1995) 785–798.
- [8] K. Park, G. Kim, M. Crovella, On the relationship between file sizes, transport protocols, and self-similar network traffic, in: *Proceedings of the IEEE International Conference on Network Protocols*, 1996, pp. 171–180.
- [9] S. Sarvotham, R. Riedi, R. Baraniuk, Connection-level analysis and modeling of network traffic, Tech. Rep., ECE Department, Rice University, July 2001.

Table 8
Network trace files used in this paper

Dataset	Auck	Bell
Filename	19991207-125019-0	20020520-110000
Link bandwidth	2 Mbps	9 Mbps
Duration (s)	3182	3600
Number of packets	1,000,000	2,229,864
Number of bytes	571 MB	870 MB
Number of connections	64,090	187,310

- [10] K. Thompson, G. Miller, R. Wilder, Wide area Internet traffic patterns and characteristics, *IEEE Network* 11 (6) (1997) 10–23.
- [11] Y. Zhang, L. Breslau, V. Paxson, S. Shenker, On the characteristics and origins of internet flow rates, *SIGCOMM*, August 2000.
- [12] K. Lan, J. Heidemann, On the correlation of Internet flow characteristics, *Tech. Rep. ISI-TR-574*, USC/Information Sciences Institute, July 2003.
- [13] S. Floyd, High-speed TCP for large congestion windows, *RFC 3649*, Experimental, Dec. 2003.
- [14] T. Kelly, Scalable TCP: improving performance in high-speed wide area networks, *Comput. Commun. Rev.* 32 (2) (2003).
- [15] D. Katabi, M. Handley, C. Rohrs, Internet congestion control for future high bandwidth-delay product environments, *ACM SIGCOMM*, 2002.
- [16] C. Jin, X. Wei, S.H. Low, FAST TCP: motivation, architecture, algorithms and performance, *IEEE Infocom*, 2004.
- [17] L. Xu, K. Harfoush, I. Rhee, Binary increase congestion control for fast long-distance networks, *IEEE Infocom*, 2004.
- [18] N. Brownlee, K.C. Claffy, Understanding Internet traffic streams, *IEEE Commun.* 40 (10) (2002) 110–117.
- [19] F.H. Campos, J.S. Marron, S.I. Resnick, K. Jeffay, “Extremal dependence: Internet Traffic Applications, *Tech. Rep. 1352*, School of Operations Research and Industrial Engineering, Cornell University, 2004.
- [20] F.H. Campos, A.B. Nobel, F.D. Smith, K. Jeffay, “Statistical clustering of Internet communication patterns, *Proceedings of the 35th Symposium on the Interface of Computing Science and Statistics*, July 2003.
- [21] W.A. Massey, W. Whitt, A probabilistic generalization of Taylor’s theorem, *Stat. Probab. Lett.* 16 (1) (1993) 51–54.
- [22] NLNR MOAT: Auckland-II trace archive, <http://pma.nlanr.net/traces/long/auck2.html>.
- [23] NLNR MOAT: Bell Labs-I data set, <http://pma.nlanr.net/traces/long/bell1.html>.



Shriram Sarvotham received the B.Tech degree in Electrical Engineering from the Indian Institute of Technology, Madras, India, in 1997. He received the MS degree, also in Electrical Engineering, from Rice University, Houston, Texas, in 2001. From 1997 to 1999, he served as a systems software engineer with HP-Verifone in Bangalore, India. In the summer of 2001, he was a research intern with HP labs,

Palo Alto, California, with the Storage Systems Program (SSP) group. Shriram Sarvotham is currently working towards a PhD degree in Electrical Engineering at Rice University. His research interests include Internet Traffic modeling, Signal Processing in sensor networks, and non-asymptotic coding in Information theory.



Rudolf H. Riedi received the MSc degree in 1986 and the PhD degree in 1993, both from the Federal Institute of Technology ETH Zurich (Switzerland) in Mathematics. From 1993–1995 he was with B. Mandelbrot at the Mathematics Department of Yale University in New Haven, Connecticut. After spending 1995–1997 with the National Research Institute in Automation and Computing INRIA in Paris, France, he

joined the Electrical and Computer Engineering Department at Rice University in Houston, Texas. From there he moved to the Dept. of Statistics at Rice University where he is currently an assistant professor. His research interests lie in the theory and practice of multiscale stochastic modeling and statistical inference, especially for networks in communication and economics. Dr. Riedi won the ETHZ Polya prize in 1986 and received a postdoctoral fellowship from the National Science Foundation of Switzerland in 1993. At Rice he has won research grants from NSF (2001, co-PI), DoE (2001, co-PI), DARPA (2000, acting PI), NSF (2003, PI) and ATP (2003, co-PI), and consulted with AT&T Labs. He served as a guest-editor for the *IEEE Transactions on Signal Processing* in 2002 and is on the technical program committee of the *IEEE Infocom* conference since 2003.



Richard G. Baraniuk received the B.Sc. degree in 1987 from the University of Manitoba, the M.Sc. degree in 1988 from the University of Wisconsin-Madison, and the Ph.D. degree in 1992 from the University of Illinois at Urbana-Champaign, all in Electrical Engineering. After spending 1992–1993 at Ecole Normale Supérieure in Lyon, France, he joined Rice University in Houston, Texas, where he is currently

the Victor E. Cameron Professor. His research interests lie in the areas of statistical and distributed signal processing and communications networks. Dr. Baraniuk received a NATO postdoctoral fellowship from NSERC in 1992, the National Young Investigator award from NSF in 1994, a Young Investigator Award from ONR in 1995, the Rosenbaum Fellowship from the Isaac Newton Institute of Cambridge University in 1998, the C. Holmes MacDonald National Outstanding Teaching Award from Eta Kappa Nu in 1999, the Charles Duncan Junior Faculty Achievement Award from Rice in 2000, the ECE Young Alumni Achievement Award from the University of Illinois in 2000, and the George R. Brown Award for Superior Teaching at Rice in 2001 and 2003. He was co-author on a paper with Matthew Crouse and Robert Nowak that won the *IEEE Signal Processing Society Junior Paper Award* in 2001 and another with Vinay Ribeiro and Rolf Riedi that won the *Passive and Active Measurement (PAM) Workshop Best Student Paper Award* in 2003. He was elected a Fellow of the IEEE in 2001 and a Plus Member of AAA in 1986.