A Comparative Study of TCP/IP Traffic Behavior in Broadband Access Networks ⋆

Amit Sinha, Kenneth Mitchell *

School of Computing and Engineering, University of Missouri-Kansas City

Abstract

In recent years, broadband access has become the norm in the Internet world. As of now, not much has been said about the behavior of traffic on these access networks. Also, the common assumption has been that the traffic behavior on one broadband access network is not very different from another. In this paper, we present an analysis of TCP/IP traffic data that has been collected from two broadband access networks: Digital Subscriber Line (DSL) and Broadband Fixed Wireless (BFW). We investigate the TCP session behavior in terms of duration and number of bytes transferred for both access schemes, including that of several dominating applications, namely, *kazaa, gnutella* (peer-to-peer (P2P) applications) and *http*. We also investigate the TCP connection interarrival time process in the two access schemes. We observe that although the TCP session behavior is similar in both access networks, the behavior of the connection interarrival time process in BFW is influenced by the underlying MAC (Medium Access Control) protocol.

Key words: Traffic Behavior; Access Networks; DSL; Broadband Fixed Wireless; Traffic Modeling; Heavy-tailed Distribution

1 Introduction

The evolution of broadband solutions in "last mile" (access) networks has enabled users to use applications that predominantly require large bandwidth along with low delay and jitter. These applications include streaming audio/video, online games, Internet Protocol (IP) telephony, peer-to-peer (P2P)

⋆ This work has been supported in part by a grant from the National Science Foundation NSF ANI-0106640

* Corresponding Author
Email address: mitchellke@umkc.edu (Kenneth Mitchell).

applications, etc. The advent of technologies such as Digital Subscriber Line (DSL), Broadband Fixed Wireless (BFW), and cable modems have made this possible. Therefore, understanding the behavior of traffic on these access networks is desired in order to analyze their performance and facilitate network design.

Various studies have been conducted which delve into the nature of traffic observed on access networks. Most of the work in this area is empirical in nature. Vicari et al. [34] provide some information about Asymmetrical Digital Subscriber Line (ADSL) traffic based on field trials. Their work compares user traffic behavior with varying access speed, though no theoretical model is provided. Cano et al. [3] describe the traffic observed on Ciez@net, which provides Internet access to users through N-ISDN technology. Oliver and Benameur [24] present a flow-level IP traffic characterization for traffic on the backbone and the ADSL *downstream* access link (traffic from the ADSL service provider to the end-users). This work is aimed at characterizing the flow-level IP traffic and not traffic specific to the access link. Farber et al. [8] present a model for Internet traffic collected from dial-up access networks. One attempt at characterizing the traffic on access networks has been made by Kilpi and Norros [16], where they present modeling aspects of the session arrival process for ISDN and dial-up user traffic. In a recent work, Roughan and Kalmanek [28] present a modeling approach for broadband access traffic based on coarse SNMP data collected from a cable access network. As of now, we are not aware of any comparative study that provides models for traffic based on actual measurements from different broadband access technologies.

In this paper, we attempt to characterize traffic observed in two different broadband access networks, namely, BFW and DSL. Note that one important distinction between BFW and DSL is that in BFW, the upstream access is shared (like cable modems), while in DSL, the upstream bandwidth is dedicated per end-user. This paper is an extension of the preliminary work presented in [30] that discusses flow-level traffic behavior on BFW and DSL. Our aim here is as much to understand the traffic behavior in the broadband access network, as it is to verify the assumption that the behavior remains same for different access schemes. We specifically look into the Transmission Control Protocol (TCP) session behavior for aggregate traffic in terms of duration and the volume of data carried. We do the same for certain dominant applications found in these accesses [30], namely, *http*, *kazaa* and *gnutella*, as they use TCP as their transport layer protocol. We also characterize the TCP connection interarrival time process by investigating various models that have been applied in the past and available in the literature and we perform a statistical goodness-of-fit test in order to find models that are reasonably close. We then use these models for performance analysis of the access networks and weigh the outcome against trace driven simulation results.

We emphasize that the results shown in this paper are from services deployed by a large ISP (Internet Service Provider) in specific markets; thus, it is important to recognize that this may not be representative of traffic observed in every broadband access environment. Furthermore, the traffic pattern can change over time due to availability of different applications over time. Nevertheless, our interest is in the comparative study of BFW and DSL, with the hope that this may be of benefit to other network service providers and designers.

The paper is organized as follows: Section 2 describes the broadband access technique and the data set collected. Section 3 consists of the characterization of the TCP traffic mix and session behavior. Section 4 presents the analysis and modeling of the TCP connection interarrival time processes on the access networks. In section 5, we present the performance analysis of the models and compare the results with trace driven simulations. Finally, we conclude this paper in section 6.

2 Access Technology and Data Description

The Broadband Fixed Wireless (BFW) service considered here uses the Multi-Channel Multi-Point Distribution Service (MMDS) as the wireless transport from customer premises back to the head end. MMDS operates in the 2.1 GHz to 2.7 GHz band and can support distances up to 30 miles between sites. Fig. 1 presents the typical architecture of the BFW environment and DSL service structure. This paper considers measurements from two different markets where the two broadband data services are deployed, consisting primarily of best-effort Internet applications. Customer bases include both residential and business customers. The markets considered in this study allow a downstream speed of up to 10 Mbps and an upstream speed of up to 256 Kbps in the case of BFW, and a downstream speed of up to 3 Mbps with an upstream speed ranging from 96 Kbps to 640 Kbps in DSL.

Flow-level data has been collected using Cisco's NetFlow from the router at the head-end for BFW and at the central office for DSL at the points indicated in Fig 1. NetFlow provides more fine grained data than SNMP, but is not as detailed and high volume as packet sniffers. According to Cisco's documentation [4], a NetFlow identifies a flow based on the tuple (*source address*, *source port*, *source interface*, *destination address*, *destination port*, *IP protocol number*, *IP type-of-service*). Any undefined fields are set to 0 for each flow. Table 1 summarizes the various fields in a typical NetFlow record. A NetFlow enabled router exports aggregated flows to some predefined destination for collection using UDP. Cisco's collection of flow level data is based on the definition given by Claffy *et al.* in [5]. According to their definition, "a flow is active as long

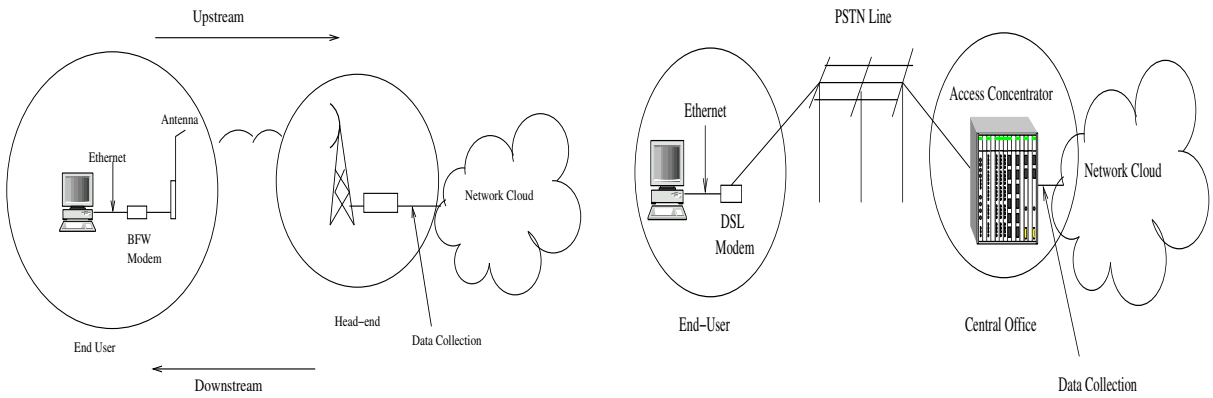

(a) BFW (b) DSL

Fig. 1. Typical service structures

Table 1
Subset of fields in a typical NetFlow record

Field Name	Description	Field Name	Description
Start	Start time of the flow	*DIF*	Destination interface
End	End time of the flow	*DstIPaddress*	Destination IP address
Sif	Source interface	*DstP*	Destination Port
SrcIPaddress	Source IP address	*P*	Protocol Number
SrcP	Source Port	*Fl*	TCP control segment Flag
Pkts	Number of IP packets	*Octets*	Number of octets
tos	IP type-of-service		

as observed packets that are meeting the flow specifications are observed separated in time by less than a specified amount of time". Cisco's additional constraints on this definition for the expiration of a flow are as follows:

- The flow exceeds a maximum duration that has been pre-defined.
- The flow contains a *FIN* or *RST* TCP flag.
- The router's cache where the flow is temporarily stored before exporting gets exhausted.

The ISP provided one hour of flow-level data for the BFW access network and another hour of data for the DSL access network. [1] Since data are collected at the IP level, it is possible to identify and characterize different applications. Moreover we used Fullmer's *flow-tools* [11] to process the data collected from NetFlow.

[1] The actual trace had a preceding and following half-hour of data, which was filtered out before the analysis.

3 TCP Session Behavior

A TCP session is defined as a bidirectional exchange of TCP/IP packets between a pair of IP host addresses using a specific source and destination port for a connection. A TCP session is established via a three-way handshake using the "*SYN*" (Synchronize) flag in the TCP packet header. The termination is done using the "*FIN*" flag [6]. A complete TCP session starts with the source sending a *SYN* flag to the destination, and ends when the source receives a TCP packet with the *FIN* flag set from the destination. Furthermore, a session can be terminated via a "*RST*" (Reset) flag, which can be sent by either the source or the destination. In order to analyze the performance of the access technologies, it is important to know the characteristics of the session duration as well as the volume of traffic exchanged on the access network during the sessions. In this section, we look into these behaviors, starting with the regeneration of TCP sessions from the NetFlow data.

3.1 Regenerating TCP Sessions from NetFlow

According to Cisco's NetFlow definition, a flow record may or may not consist of complete TCP sessions. Therefore, we have built a post-processing tool to regenerate TCP sessions from the NetFlow flow records. This tool is based on the logic used by Sommer and Feldmann to build FLOW-REDUCE [31], a post-processing tool used to recreate TCP summaries from NetFlow records. Sommer observed that their tool could produce close approximations for over 90% of all connections. To be on the conservative side, we assume that our tool is at least as reliable as FLOW-REDUCE, since certain anomalies observed by Sommer in their NetFlow records are not found in the trace that we are investigating. One important difference between FLOW-REDUCE and our tool is that we do not have to incorporate the logic for finding the origin of the sessions since we already have the origin information (i.e. the subscribers' IP addresses). Note that in certain cases, the origin may or may not be an internal user (subscriber), since many peer-to-peer (P2P) applications can be initiated by an external user. Our tool takes care of this situation by carefully looking at the source and destination port, and determines if either matches those used by some popular P2P applications, namely *kazaa* and *gnutella*. If a match occurs, we check the other port number to see if it is a standard port (i.e. port numbers less than 1024), else the flow record belongs to a P2P session irrespective of whether the originator is a subscriber or an external user.

The output generated by the tool is a file that contains only complete TCP

Table 2
TCP session information

Dataset	Complete	Incomplete	Application
DSL	89.70%	10.30%	P2P, *smtp*, *http*
BFW	80.00%	20.00%	P2P, *http*

sessions [2] that includes *start time*, *session duration*, *source port*, *destination port*, *bytes upstream*, *bytes downstream*, *source address*, *destination address*, *TCP flags seen in upstream flows*, and *TCP flags seen in downstream flows*. This is in accordance with the TCP connection summaries suggested in [31]. To avoid edge effects due to finite data sets, we did not include the first ten minutes of the flows that contained a destination *FIN* with no corresponding source *SYN*, and also the last ten minutes of the record that had a source *SYN* with no corresponding destination *FIN*. This is done with the assumption that the first ten minutes of incomplete NetFlow record would have started before the start of the data collection, while the last ten minutes of the flow would have continued even after the data collection stopped. Table 2 gives the detail of the TCP session data generated from the NetFlow records. The first column gives the access type to which the dataset belongs, while the second column gives the percentage of complete TCP session information that the tool generates from the dataset. The third column gives the percentage of incomplete TCP sessions and the fourth column gives the application that dominates the incomplete information. The P2P applications were expected to dominate because sessions in these applications have been observed to last well over a week [14], and hence it would be impossible to capture entire sessions with an hour's worth of data. Note that the main purpose of this paper is to draw a comparison between the traffic behavior in the two broadband access networks, and therefore looking into the available TCP information would help us by providing some good insights.

3.2 Traffic Overview

We now present the statistics related to the two datasets to highlight the overall similarities and differences in the traffic mix. Table 3 shows the connection mix, total TCP traffic (byte) mix, and the upstream TCP traffic (byte) mix. The first column indicates the access network to which the dataset belongs, and the remaining columns give the percentage contribution of various TCP based protocols to the total flow. The ratio of the total number of complete TCP connections for DSL to BFW is 1:14, while that for the total TCP traffic (bytes) is 1:13, and that for the total upstream traffic (bytes) is 1:15. Note

[2] By complete we mean, the TCP session which had a source *SYN* and destination *FIN*, or a source *SYN* and a source/destination *RST*

Table 3
Percentage connection and byte mixes for various TCP based application

TCP Connections								
Dataset	*http*	*kazaa*	*gnutella*	*smtp*	*https*	*pop3*	*nntp*	others
DSL	64.65%	1.43%	4.30%	5.46%	3.07%	11.28%	0.00%	9.81%
BFW	68.90%	7.17%	3.68%	0.94%	3.36%	8.59%	0.05%	7.25%
Total TCP Traffic (Bytes)								
DSL	31.21%	18.37%	7.81%	2.90%	1.80%	1.48%	0.00%	36.43%
BFW	42.35%	21.08%	6.78%	3.27%	1.57%	2.84%	2.51%	19.70%
Upstream TCP Traffic (Bytes)								
DSL	12.47%	52.15%	14.70%	11.98%	1.14%	0.14%	0.00%	7.52%
BFW	12.32%	39.65%	11.88%	12.56%	0.98%	0.11%	0.01%	22.50%

that these traces belong to different markets and were collected at different times of the day. Therefore, the difference here should not come as a surprise.

To start with, we observe that there is a slight variation in the TCP connection mix in the two access networks. For example, the relative contribution to the number of TCP connections from *smtp* is higher in DSL as compared to BFW. On the other hand, the relative number of *kazaa* based connections shoots up for BFW.

We see variations in the contribution from applications to the total TCP traffic mix (bi-directional traffic) in the two access networks. We observe that the traffic generated by BFW-*http* has higher relative contribution compared to DSL-*http* traffic to the total volume. The other significant variation is that we find *nntp* traffic in BFW access, which is absent in the DSL dataset. There are high variations in the contribution from *other* TCP applications in DSL when compared to BFW. On further investigation we find that 70% of this *other* traffic belongs to unknown applications (i.e. applications using port numbers greater than 1023). So we check the destination address and find that it belongs to an organization that helps establish the best optimized route for applications like video-conferences by initially routing data through its own network.

It is well established that for any video transmission UDP is the transport protocol to use, so we calculate the upstream to downstream ratio of the data transferred and find it to be 1:85 (i.e. for every 1 byte sent upstream, there is 85 bytes of payload sent downstream), which is highly asymmetric. This indicates that the users are involved in heavy downloading, rather than two way conversations. This though is not a big concern, as the access schemes are

asymmetric. Therefore, we also find it important to look at the traffic contribution (bytes) by various TCP based applications in the upstream direction. This would reflect those protocols responsible for high bandwidth consumption in the upstream direction.

The last two rows of Table 3 give the percentage contribution of various TCP based protocols to the total upstream traffic (bytes). It is clear that the DSL access network carries a relatively higher contribution from P2P applications (66.85%) as compared to BFW (55.53%). In BFW upstream traffic we find a significant contribution from *other* traffic as compared to DSL. Significant among these were *napster* (a P2P application), *ftp* and some unknown applications. The important point to notice here is that the contribution from P2P applications in both networks is more than four times the next highest contributor.

Now that we have an idea about various applications contributing to the traffic traces collected, we are ready to characterize the TCP behavior and look further into specific protocols which have high impact on the access network both in terms of session duration and the amount of bytes being exchanged. The overall TCP behavior is naturally influenced by these protocols.

3.3 File Size Characteristics

The session size or the file size is defined as the total number of bytes transacted during a TCP session that starts with a source *SYN* and ends with a destination *FIN*. Fig. 2 shows the log-log plot for the complementary cumulative probability distribution function (CCDF) of the file size for the TCP sessions in DSL and BFW. Here, we observe that there is a striking similarity between the distributions observed for the different access schemes. Moreover, the linear behavior of the plots for the tail of the file size distribution indicates heavy-tail characteristics. On further investigation into the dominating applications, we find that both *http* and P2P file size distributions exhibit

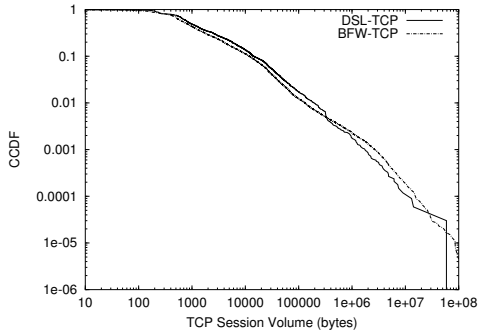

Fig. 2. TCP connection volume (bytes) distribution

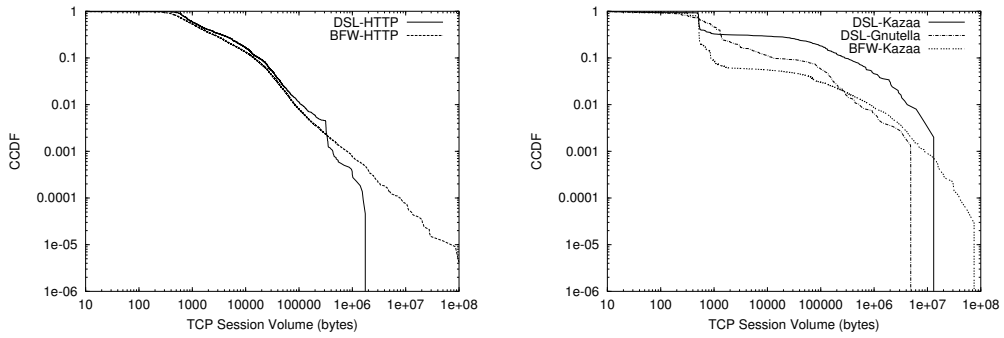

(a) HTTP (b) P2P

Fig. 3. TCP connection volume (bytes) distribution

heavy-tails (see Fig. 3). This is in agreement with the observations made by Crovella *et al.* [7] for *http* traffic. It is interesting to note that all P2P applications in Fig. 3(b) have similar behavior. There are some small transactions, which consist of requests or queries sent to find the location of a particular file the user is looking for, and then the actual file transfer takes place. This is reflected by the bimodal nature of the distribution. Also, note that the probability of exchanging large files through *kazaa* in BFW access is less than that in DSL. We speculate the possible causes for this difference are either(a) different user behavior in different markets, (b) the difference in the time-of-the-day of the trace, (c) the different access speed in the two schemes, or (d) some combination of the above.

3.4 Session Duration Characteristics

The proper understanding of a TCP session duration or holding time is important in characterizing the workload on the access network. Fig. 4 depicts log-log plots that present the probability density function (PDF) for TCP session durations in DSL and BFW access networks. We have plotted the graphs on the same scale so that the similarities and differences between the two become evident. The PDF for BFW is more spread out as compared to DSL, which tends to have negligible session durations below 20 milliseconds, but has a higher probability of session durations greater than 10 seconds. The sessions are generally short-lived with a peak occurring at approximately 200 milliseconds for both access networks (204 milliseconds for DSL and 208 milliseconds for BFW). On further investigation, we find that for the DSL access network, the peak session duration is dominated by an application called *GoToMyPC* (75% of sessions with peak duration), that supports remote desktop access. These sessions for *GoToMyPC* are generated by the HTTP "pings" that a desktop host has to send to the broker to check for new connection requests [12]. Of these sessions, 21.70% belong to *http*. In case of BFW access, we found that 92.70% of the peak sessions belong to *http*, while 2.52% of these

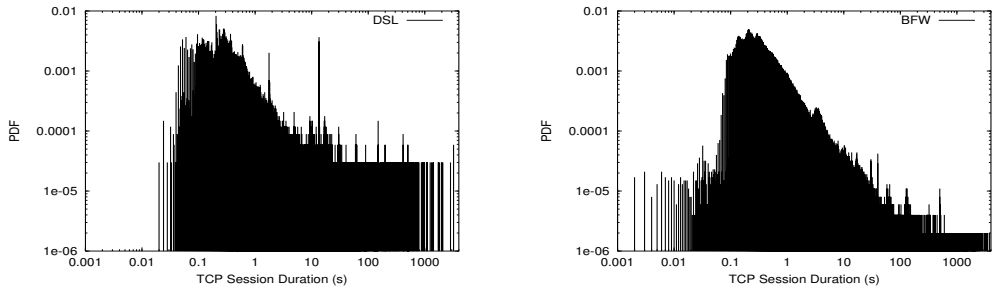

(a) DSL (b) BFW

Fig. 4. TCP session duration distribution in broadband access

belong to *GoToMyPC*. The mean session duration for DSL is 8.055 seconds and that for BFW is 10.807 seconds.

Next, we consider the tail behavior of the TCP session duration for both access schemes. Fig. 5 shows the log-log plot for the CCDF of session durations in both access networks. There is a striking resemblance in the tail behavior. Moreover, the linear plots indicate that the session durations are heavy-tailed. On further analysis, we found that the two dominant applications (*http* and *kazaa*) are responsible for such characteristics. Fig. 6 illustrates the tail of the session length distribution of the two dominating applications found on each access network. In case of *http*, the most dominant application found in both access schemes, we observe the same similarities. The heavy-tail behavior of the *http* transmission time is consistent with the observation in [7]. It is evident that the tail for P2P sessions is heavier than that of *http* applications for the respective access schemes (the scale of the y-axis in Fig. 6(a) and 6(b) are different). This is expected, as P2P applications involve the exchange of large files in both the upstream and downstream directions, thus generating flows that last longer, even spanning several days [14]. Fig. 6(b) is only indicative of the one-hour data trace, with the intention of drawing a comparison between the two broadband schemes. The tails exhibited by P2P applications differ for the different access networks, i.e, the tail for DSL-*kazaa* session duration is heavier than BFW-*kazaa*. This behavior depends on the file size being transacted during the session and the access speed.

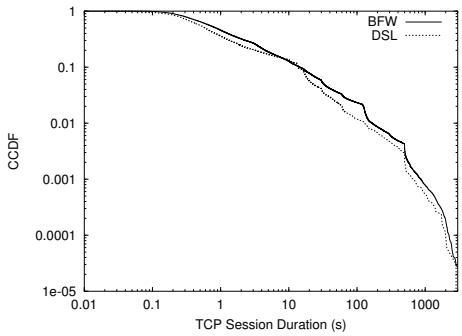

Fig. 5. Tail behavior for TCP session duration distribution

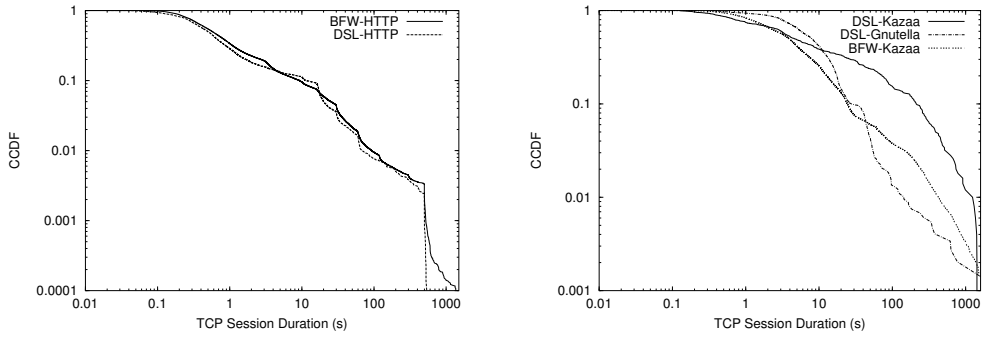

(a) HTTP (b) P2P

Fig. 6. TCP session duration distribution

4 TCP Connection Interarrival Time Process

So far, we have analyzed and compared the TCP session behavior for DSL and BFW. We now investigate the TCP connection interarrival time process to see if this behavior also remains untouched by the underlying access protocols. Since the arrival of a new TCP connection uses a three-way handshake identified by *SYN* (Synchronize) flag in the TCP packet header, we filter all the upstream TCP flows that have the *SYN* flag set. Moreover, since the granularity of the measured data is 1 millisecond, an arrival indicated at time t milliseconds actually takes place somewhere between t and $t+1$ milliseconds, so we make the assumption that the arrivals are evenly distributed in the interval $[t, t+1)$.

The mean TCP connection interarrival time for BFW (0.004s) is smaller than DSL (0.0698s). Moreover, the squared coefficient of variation indicates that the arrival process to the network access point in BFW (c^2=3.598) is relatively bursty compared to that of DSL (c^2=1.769). Does this mean that all broadband access protocols may not have the same behavior, i.e., would it be questionable to design a network with one access type based on observations of another access? More importantly, is there a single theoretical model that can be applied to all broadband access schemes? In order to answer these questions, it is desired to see how well the various known theoretical source models approximate the TCP connection interarrival time process in DSL and BFW access networks. But, first we need to test the applicability of the existing models to the data set.

4.1 Scaling and Stationarity Test

Other researchers have observed that the TCP connection arrival process shows self-similar behavior for WAN (Wide Area Network) traffic [9,10]. There-

fore, we first investigate to see if scaling behavior in the connection arrival process exists. We follow the work by Veitch and Abry [1,32], which use wavelet transforms to analyze long-range dependencies in traffic over a certain range of timescales. These timescales, starting from the reference scale t_0, become coarser according to the relation t_j=$2^j t_0$ for increasing j, and partition the trace in consecutive non-overlapping time intervals. The linear relation between the logarithm of energy (Y_j) and the scale j indicates scale invariance. The slope of the plot between the two is referred to as the scaling exponent (α), and is related to the Hurst parameter H by the expression α=2H-1.

Fig. 7, created using Daubechies wavelets with 3 vanishing moments, shows the scaling analysis for the number of TCP connections arriving during 500 millisecond intervals for the DSL and BFW traces. The plot for DSL shows a clear non-horizontal scaling region even for j=1, corresponding to 1 second, whereas the plot for BFW indicates scaling starting at the 5th octet (corresponding to 16 seconds). The straight line, constructed using weighted regression over the selected range of scales (j_1,j_2), exhibits good alignment and is used to provide a good estimate of slope ($\hat{\alpha}$). The vertical line at each point indicates the confidence intervals and the outcome of the test with a significance level of 95%. In Fig. 7(b), it is interesting to observe a dip at j=2 (2 seconds), which indicates the existence of strong periodicity at j=1 (1 second). This is because the energy at scale j depends on the variation in the process at scale j-1 and the periodicity reduces the variability of the process at that timescale.

Now that we know that the TCP connection arrivals for the traces show scaling behavior (at different timescales for the two access schemes), it is important to determine if the process is stationary to ensure that these scaling behaviors are not due to non-stationarity. The experiment is based on another work by

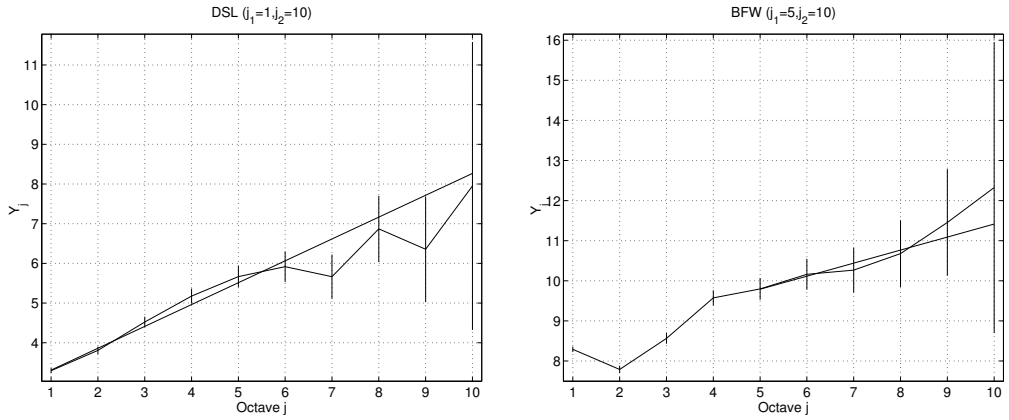

(a) DSL (b) BFW

Fig. 7. Logscale diagram of the whole time series using Daubechies wavelets with 3 vanishing moments

Veitch and Abry [33], which uses wavelet-based estimators to determine the time constancy of the scaling exponent of the long-range dependent process under investigation. We start with the hypothesis H_0 that the mean of scaling exponents are identical. Accordingly, each time series is split into $m = 8$ blocks, and the scaling exponent, $\hat{\alpha}$ is computed for each block using Daubechies wavelets with 3 vanishing moments. Fig. 8 presents the plot for $\hat{\alpha}$ for each block, along with the vertical line at each point, which indicates the confidence interval and the outcome of the test with a significance level of 95%. The plots also consist of two horizontal lines. The solid line is $\hat{\alpha}$ computed for the overall trace, while the dashed line indicates $\hat{\alpha}$ averaged over the 8 blocks. For both DSL and BFW the hypothesis H_0 was not rejected. This indicates that the TCP connection arrival processes in the access schemes are probably long-range dependent and stationary.

4.2 Empirical Interarrival Distribution

We plot the probability density functions (PDF) of the interarrival times for DSL in Fig. 9(a) and for BFW in Fig. 9(b) as estimated from the trace. The significant difference is that the PDF in the case of BFW consists of two peaks, while that in DSL is a slow decay. The arrival process in BFW is influenced by the underlying *proprietary* MAC (Medium Access Control) protocol along with the error control mechanism that runs between the end-user and the head-end. This protocol uses contention and polling to decide when a flow is accepted from any particular end-user. This causes the two peaks that can be observed in the TCP connection interarrival time distribution. We further look into the lag coefficient of autocorrelation function for this process.

In Fig. 10 we see that the correlation coefficient in the TCP connection interarrival time appears insignificant for BFW, while we find some short-term

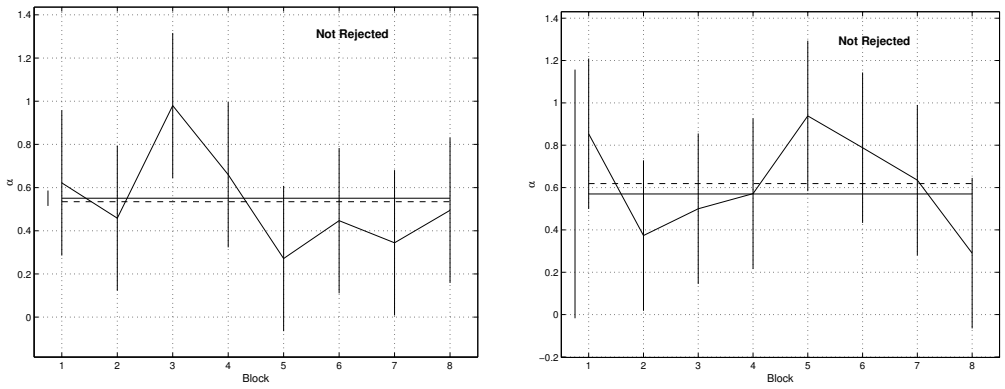

(a) DSL (b) BFW

Fig. 8. Time constancy of scaling exponents

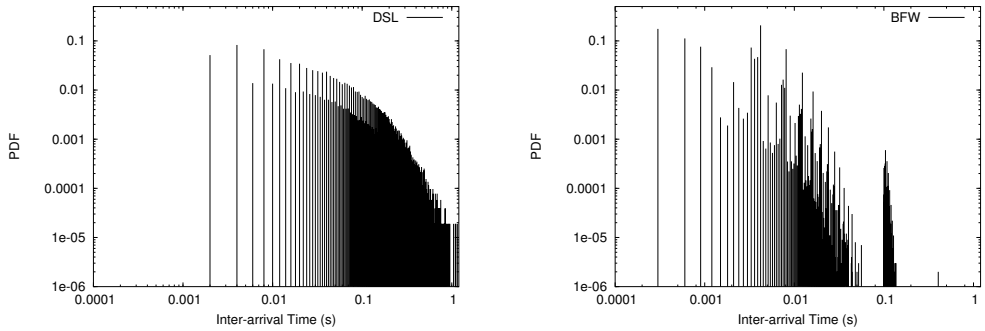

(a) DSL (b) BFW

Fig. 9. Probability density function for TCP connection interarrival time process

correlations for DSL. We then plot the lag coefficient of autocorrelation for the TCP connection arrival count process for a bin size of 500 milliseconds in DSL and for BFW in Fig. 11. The BFW arrival count process exhibits strong oscillatory behavior (alternating between positive and negative values for every 500 milliseconds lag, nearly symmetric about $y = 0$) and an extremely slow decay of correlation, which is not captured in the interarrival time analysis. This is indicative of periodicity in the BFW data for the count process which was also realized during the scaling analysis (see the dip at scale 2 in Fig. 7(a)).

In order to verify the existence of cyclical behavior and find the dominant frequencies, we calculate the Fourier transform of the count process with a small bin size of 60 milliseconds (to avoid possible aliasing [2]). We then compute the periodogram and display the plot in Fig. 12. We observe that there are multiple periods (200, 250, 333, 500 and 1000 milliseconds/cycle) and that the 500 millisecond period is the most dominant one. To sum it up, the TCP connection arrival count process in BFW appears to contain deterministic components and exhibits cyclical behavior at small timescales. The implica-

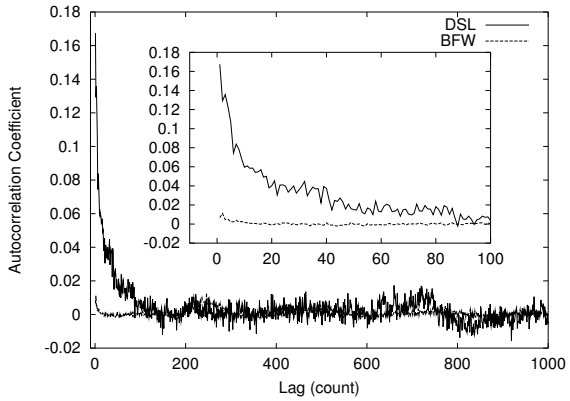

Fig. 10. Autocorrelation coefficient versus lag for TCP connection interarrival time process

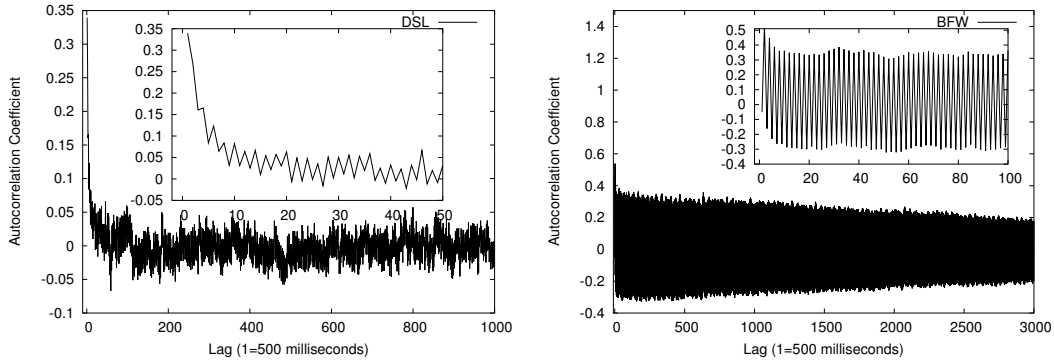

(a) DSL (b) BFW

Fig. 11. Autocorrelation coefficient versus lag for TCP connection arrival count process

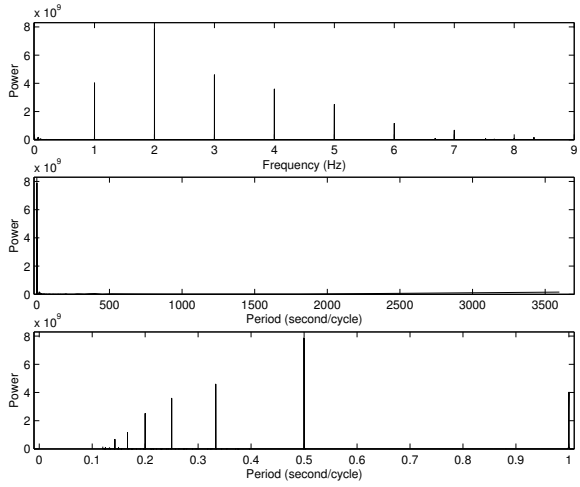

Fig. 12. Periodogram

tion of the observations in the TCP session arrival count process for the BFW access network cannot be underestimated and needs a thorough investigation. We believe that this anomaly is an outcome of the effect of the MAC protocol of the BFW access, though a thorough analysis is outside the scope of the present paper and will be addressed separately. Our aim here is not to create new models for the TCP arrival process in the access schemes, but to differentiate behavior based on their proximity to existing models.

4.3 Theoretical Models and Statistical Tests

The choice of the models is based on previous work done in the area of finding suitable distributions to represent TCP connection arrivals. We consider exponential, Weibull, lognormal, Pareto, and *truncated-power-tail* (TPT) distributions for our analysis. Table 4 provides the definition of all the probability

distributions in consideration. Paxson and Floyd [27] found that for wide-area traffic, Poisson distributions were only valid for modeling user-level session arrivals and failed to model other arrivals. Paxson [25,26] observed that log-normal, Pareto, and extreme distributions provide statistically better models for the tails. Feldmann, in her work [9], noted that the Weibull and lognormal distributions were better suited to represent TCP connection interarrival time processes. Greiner, Jobmann and Lipsky [13], presented an analytic class of well-behaved distributions (a sub-class of which are *phase* distributions) that have *truncated-power-tails* (TPT), and in the limit are power-tail distributions. One such class was used by Hatem, Lipsky and Fiorini [15] for comparison of buffer usage with multiple servers. The reliability function for this class can be represented by:

$$R_M(x) = \frac{1-\theta}{1-\theta^M} \sum_{n=0}^{M-1} \theta^n e^{\left(-\frac{\mu x}{\gamma^n}\right)}, \tag{1}$$

where θ and γ are parameters satisfying the inequalities: $0 < \theta < 1$ and $\gamma > 1$. The l_{th} moment is given by:

$$E(X_M^l) = l! \frac{1-\theta}{1-\theta^M} \cdot \frac{1-(\theta\gamma^l)^M}{1-\theta\gamma^l} \cdot \frac{1}{\mu}. \tag{2}$$

The TPT distributions are M-dimensional *phase* distributions whose vector-matrix representations, $< \boldsymbol{p}, \boldsymbol{B}, \boldsymbol{e'} >$ are given by (see [17]):

Table 4
Distribution models, estimated parameters, and range ($\hat{\lambda}^2 \pm \sqrt{\hat{\nu}\left(\hat{\lambda}^2\right)}$)

Model	Definition	DSL		BFW	
		Est. Para.	Range	Est. Para.	Range
Weibull	$f(x)=\frac{1}{a}\left(-\frac{x}{a}\right)^{c-1} e^{-\left(\frac{x}{a}\right)^c}$	$a = 0.061$, $c = 0.793$	0.000-0.001	$a = 0.003$, $c = 0.764$	2.45-2.76
TPT(M=4)	$R(x)=\frac{1-\theta}{1-\theta^4} \sum_{n=0}^{3} \theta^n e^{\left(-\frac{\mu x}{\gamma^n}\right)}$	$\theta = 0.208$, $\mu = 18.682$, $\gamma = 1.967$	0.023-0.026	$\theta = 0.213$, $\mu = 395.91$, $\gamma = 2.597$	15.66-15.75
TPT(M=2)	$R(x)=\frac{1-\theta}{1-\theta^2} \sum_{n=0}^{1} \theta^n e^{\left(-\frac{\mu x}{\gamma^n}\right)}$	$\theta = 0.045$, $\mu = 16.481$, $\gamma = 4.522$	0.036-0.040	$\theta = 0.027$, $\mu = 305.35$, $\gamma = 9.711$	15.46-15.56
Lognormal	$f(x)=\frac{1}{x\sqrt{2\pi\sigma}} e^{-\frac{(\log(x)-\zeta)^2}{2\sigma^2}}$	$\zeta = -3.514$, $\sigma = 1.569$	0.057-0.061	$\zeta = -6.417$, $\sigma = 1.581$	7.94-7.98
Exponential	$f(x)=\frac{1}{\rho} e^{-\left(\frac{x}{\rho}\right)}$	$\rho = 0.0698$	0.073-0.101	$\rho = 0.004$	7x10^5-8x10^5
Pareto	$f(x)=\frac{ak^a}{x^{a+1}}$	$a = 2.47$, $k = 0.048$	1.697-1.734	$a = 3.752$, $k = 0.008$	2.033-2.042

$$\boldsymbol{p} = \frac{1-\theta}{1-\theta^M}\left[1\ \theta\ \theta^2\ \cdots\ \theta^{M-1}\right], \boldsymbol{B} = \mu\begin{bmatrix} 1 & 0 & 0 & \vdots & 0 \\ 0 & \gamma^{-1} & 0 & \vdots & 0 \\ 0 & 0 & \gamma^{-2} & \vdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \vdots & \gamma^{-(M-1)} \end{bmatrix}, \tag{3}$$

and $\boldsymbol{e'}$ is a column vector with all 1's.

This approach of representing processes with matrix operators is called *Linear Algebraic Queueing Theory* [17] and the distribution is referred to as the matrix exponential distribution (see the Appendix for background information on matrix exponential distributions). The TPT distribution in equation (1) can now be represented in matrix exponential form as follows:

$$R_M(x) = \boldsymbol{p}\exp(-x\boldsymbol{B})\boldsymbol{e'}. \tag{4}$$

The variables θ, γ, and μ can be estimated by using the first three moments of the interarrival time data in equation (2). The parameters of the other distributions in contention are estimated using the Maximum Likelihood Estimator (MLE) technique to fit the data set to the respective distributions (see Table-4 of [9]).

We use statistical goodness-of-fit measures suggested by Paxson [26,27] in order to judge if one model fits better than the other to the empirical data. We use the discrepancy metric $\hat{\lambda}^2$ along with its estimated variance $\hat{\nu}(\hat{\lambda}^2)$ [26]. $\hat{\lambda}^2$ measures the magnitude of the departure of the data from the model, i.e., the larger the value of $\hat{\lambda}^2$, the greater the discrepancy between the data and the model. The discrepancy metric $\hat{\lambda}^2$ is computed by dividing the interarrival time into bins and estimating the number of arrivals in each bin from the experimental distribution and comparing it to the actual one. The bins are spaced logarithmically and the bin-width is given by $w = 3.49\hat{\sigma}n^{-4/9}$, where n is the number of instances in the distribution [29,26,9]. The adjacent bins are combined when required to ensure at least 5 observations per bin. Now, if $\hat{\lambda}^2 - \sqrt{\hat{\nu}\left(\hat{\lambda}^2\right)}$ for a model is greater than $\hat{\lambda}^2 + \sqrt{\hat{\nu}\left(\hat{\lambda}^2\right)}$ of the other model, then the latter model is supposed to be a better fit to the actual model.

The empirical distribution for the TCP connection interarrival process in BFW has two peaks and is not smooth. Such data would cause large $\hat{\lambda}^2$ values, since it is difficult to fit it to any well-known analytical distribution. Yet, the analysis here has two-fold goal: (a) to find the model that fits the best for a

particular access scheme, (b) to see if this model does well in the other access scheme. Therefore, we accept the best fit for the limited number of candidate distributions we have.

Table 4 presents the estimated parameters and the range for the discrepancy estimate, $\hat{\lambda}^2 \pm \sqrt{\hat{\nu}\left(\hat{\lambda}^2\right)}$, for various distributions. The Weibull distribution yields statistically better models for the TCP interarrival time in the DSL access followed by the TPT distributions. In the case of the BFW access network, there is no good statistical fit as the discrepancy estimates have relatively high values for all the distributions. Therefore, it indicates that the assumption that the traffic behavior in different broadband access schemes is similar is misleading. There is no single theoretical model that can capture the behavior in different access schemes.

5 Performance Analysis

We have found the traffic characteristics in the two broadband access networks to be close except for the TCP connection interarrival time processes. It is the nature of the arrival process that determines the network resources that need to be allocated to ensure smooth running of the ISP. It is therefore necessary to verify how closely the theoretical models replicate the queueing behavior of the actual trace. In this section, we see how the queueing and call blocking behavior vary on the two access schemes. We also see the suitability of the theoretical models for designing future access networks by comparing their performance to that of trace driven simulations.

5.1 Queue Statistics for Renewal Processes

We test the applicability of various source models by comparing the queueing performance of the marginal distributions to that of trace driven simulations. In order to find the queue statistics of the trace under the assumption that the connection arrivals are renewal, the trace should have no correlation while keeping the marginals intact. Therefore, we shuffle the interarrival time data for DSL and BFW respectively until the correlations become negligible. We then use the shuffled interarrival time data as the input to a queue with exponentially distributed service times, 80% system utilization, and an infinite buffer. This simulated *GI/M/1* model was created using CSIM [19]. Table 5 provides results from the simulation, as well as from the constructed source models. The queue length statistics confirm that the connection arrival process in the BFW is more bursty than in DSL.

The queue statistics from Weibull, lognormal, and TPT distributions are in close agreement with the statistics of the shuffled data for DSL. In the case of BFW, all the models either overestimate or underestimate the variance in the queue length, although TPT(M=2) seems to most closely approximate the queue statistics.

5.2 Queue Statistics for Non-renewal Processes

The next step is to analyze the non-renewal process, i.e, the process with the correlation structure intact. In order to do this, we use the original interarrival time data as input to a queue with exponential service time, 80% system utilization, and infinite buffer in the simulated *G/M/1* model and observe the queue statistics. On comparing the queue length statistics of Table 6 with Table 5, we observe that the correlation causes considerably higher buffer overflow than that predicted by the renewal models.

This behavior in BFW would have been a surprise had we not observed the oscillatory behavior and extremely slow decay of the correlation in the arrival count process (see Fig. 11(b)). It would have been difficult to perceive this from the autocorrelation behavior of the interarrival time process in BFW (see Fig. 10).

Mitchell [20], demonstrated a method to develop analytic Markovian traffic source models using matrix exponential distributions in which the correlation structure can be arbitrarily constructed leaving the marginals invariant (see Appendix A.2 for a brief discussion). Therefore, we create a non-renewal model for the connection interarrival time process with same marginals as that of the

Table 5
Model Verification Test - I

Models	Queue Length Statistics					
	DSL			BFW		
	Mean	Variance	c^2	Mean	Variance	c^2
Shuffled data	4.955	30.371	1.237	6.391	54.645	1.338
Weibull	5.148	34.502	1.302	5.673	41.530	1.290
TPT(M=4)	4.954	31.856	1.298	6.962	65.747	1.356
TPT(M=2)	4.907	31.212	1.296	6.650	59.690	1.350
Lognormal	4.877	30.788	1.294	7.140	69.471	1.363
Exponential	4.00	20.00	1.25	4.0	20.00	1.25
Pareto	-	-	-	0.292	0.235	2.751

TPT distribution in equation (4) and call it TPT-NR (Truncated Power Tail-Non Renewal distribution.) The most parsimonious model is constructed by using a single correlation parameter β, which is the geometric rate of decay of the autocorrelation. We try different values for lag coefficient of autocorrelation decay, and eventually find β=0.99 for DSL and β=0.50 for BFW based on the model's proximity to the performance of actual trace. We verify the model against the one achieved from trace-driven simulation in Table 6. The TPT-NR model does far better than the models in Table 5 in capturing the queue length statistics of the original data. This is obvious from the fact that correlation plays a significant role in queueing systems, which is captured in the TPT-NR model. Though TPT-NR does better than other models, it is still far from capturing the actual queueing behavior since we are approximating a complex correlation structure with only a single decay parameter.

5.3 Blocking

Finally, we observe the blocking probabilities using various models and compare them against the trace driven simulation model for the *G/M/1/C* queue while varying the buffer size between 5 and 90 with 80% utilization. We use Telpack [23] to solve the continuous finite QBD (Quasi-Birth-Death) system and calculate the blocking probabilities for arrivals that follow TPT and TPT-NR distributions. Results for the Weibull and lognormal distributions are obtained by simulating a finite queueing system.

Fig. 13 gives the plot for call blocking probabilities with varying buffer size for the various source models and the actual trace. Blocking remains high even for large buffer sizes in DSL as compared to BFW due to the existing correlation in the connection interarrival time process. Observe that the TPT-NR model provides blocking probabilities that closely represent the actual blocking as observed in the actual trace. These results seem to provide some credibility to the model that we have developed for the connection arrival process.

Table 6
Model Verification Test - II

Models	Queue Length Statistics					
	DSL			BFW		
	Mean	Variance	c^2	Mean	Variance	c^2
Original data	16.890	689.013	2.415	8.443	157.826	2.214
TPT-NR(M=4)	28.596	1523.623	1.863	9.156	119.344	1.424
TPT-NR(M=2)	9.212	125.281	1.476	8.205	93.788	1.393

6 Summary and Conclusion

In this paper, we presented the TCP/IP traffic behavior in DSL and BFW access schemes. We found some striking resemblances in the session lengths and byte counts in BFW and DSL access networks. We also observed that the TCP session duration and file size distributions are both heavy tailed. We found that *http* traffic exhibits heavy-tailed behavior, which is consistent with the results found by many others in their analysis. We also found strong heavy-tailed characteristics in the traffic generated by *kazaa*, a P2P based application. These applications dominate both types of broadband access that was studied.

We observed that the TCP connection interarrival time processes were both bursty yet significantly different in the two access schemes. We found significant correlation in the interarrival time process in DSL up to lag 100, while the correlation in BFW appeared insignificant. However, investigation into the autocorrelation structure of the TCP count arrival process in the BFW access network revealed oscillatory behavior and an extremely slow decay.

We also studied various theoretical models to see if a single model could closely represent the TCP connection interarrival time process in both the access schemes. We performed statistical goodness-of-fit tests and tested the models against trace driven simulations to observe their proximity to the actual queueing behavior.

We conclude that the assumption of traffic behavior being same across various broadband access schemes could lead to bad network resource allocation and design. The session behavior depends on the application protocol and hence stays consistent across different accesses, but the TCP connection arrival process depends on the underlying MAC protocol and the error control mechanisms, which differ for different access schemes.

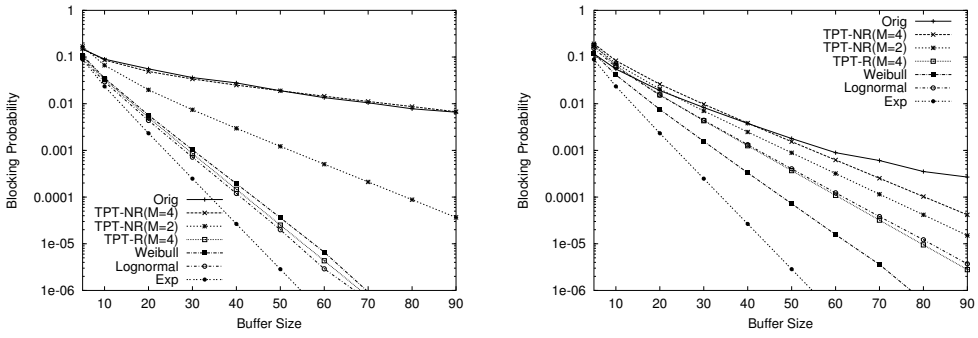

(a) DSL (b) BFW

Fig. 13. Blocking Probabilities for varying buffer size

Acknowledgment

The authors thank Sprint Corporation for making the representative data set available for this study; in particular, we thank Nick Gerber, and Sheldon Fisher for their help. We thank Robin Sommer for providing us the code for Flow-Reduce. We would also like to thank Dr. Deep Medhi for helpful discussions and insights.

References

[1] P. Abry and D. Veitch, "Wavelet analysis of long-range-dependent traffic," *IEEE Transactions on Information Theory*, vol. 44, no. 1, pp. 2–15, 1998.

[2] P. Bloomfield, *Fourier Analysis of Time Series: an Introduction.* New York: John Wiley, 1976.

[3] M. Cano, J. Malgosa-Sanahuja, F. Cerdan, and J. Garcia-Haro, "Internet measurements and data study over the regional network ciez@net," *IEEE Pacific Rim Conference on Communications, Computers and Signal Processing*, vol. 2, pp. 393–396, 2001.

[4] Cisco, "Netflow services and applications." `http://www.cisco.com/warp/public/cc/pd/iosw/ioft/neflct/tech/napps_wp.pdf`, 2002. White Paper.

[5] K. Claffy, H. Braun, and G. Polyzos, "A parameterizable methodology for Internet traffic profiling," *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 8, pp. 1481–1494, 1995.

[6] D. Comer, *Internetworking with TCP/IP: Principles, Protocols, and Architecture*, vol. 1. Prentice-Hall, 4 ed., 1996.

[7] M. Crovella and A. Bestavros, "Self-similarity in world wide web traffic: Evidence and possible causes," *IEEE /ACM Transactions on Networking*, vol. 5, no. 6, pp. 835–846, 1997.

[8] J. Farber, S. Bodamer, and J. Charzinsky, "Measurement and modelling of Internet traffic at access networks," *Proceedings of the EUNICE '98 - Open European Summer School on Network Management and Operation*, pp. 196–203, 1998.

[9] A. Feldmann, "Characteristics of TCP connection arrivals," technical report, AT&T Labs – Research, 1998. `http://www.research.att.com/~anja/feldmann/papers/ss_conn.ps`.

[10] A. Feldmann, A. Gilbert, W. Willinger, and T. Kurtz, "The changing nature of network traffic: Scaling phenomenon," *ACM Computer Communication Review*, vol. 28, no. 2, pp. 5–29, 1998.

[11] M. Fullmer, "flow-tools." http://www.splintered.net/sw/flow-tools/.

[12] "Gotomypc: Making life simpler for teleworkers and travelers." http://www.si.umich.edu/Classes/540/Placement/GoOvrvew.pdf.

[13] M. Greiner, M. Jobmann, and L. Lipsky, "The importance of power-tail distributions for modeling queueing systems," *Operations Research*, vol. 47, no. 2, pp. 313–326, 1999.

[14] K. Gummadi, R. Dunn, S. Saroiu, S. Gribble, H. Levy, and J. Zahorjan, "Measurement, modeling, and analysis of a peer-to-peer file-sharing workload," *ACM Symposium on Operating Systems Principles*, pp. 314–329, 2003.

[15] J. Hatem, L. Lipsky, and P. Fiorini, "Comparison of buffer usage utilizing multiple servers in network systems with power-tail distributions." INFORMS (1997). http://www.engr.uconn.edu/~lester/.

[16] J. Kilpi and I. Norros, "Call level traffic analysis of a large ISP," *ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*, pp. 6.1–6.9, 2000.

[17] L. Lipsky, *Queueing Theory: A Linear Algebraic Approach.* New York: MacMillan, 1992.

[18] L. Lipsky, P. Fiorini, W. Hsin, and A. Liefvoort, "Auto-correlation of lag-k for customers departing from semi-Markov processes," Technical Report TUM-19506, Technical University Munich, 1995. http://www.engr.uconn.edu/~lester/.

[19] Mesquite, "Csim: Process based simulator." http://www.mesquite.com/.

[20] K. Mitchell, "Constructing a correlated sequence of matrix exponentials with invariant first-order properties," *Operations Research Letters*, vol. 28, no. 1, pp. 27–34, 2001.

[21] K. Mitchell and A. Liefvoort, "Approximation models of G/G/1/N queueing networks with correlated arrivals," *Performance Evaluation*, vol. 51, no. 2-4, pp. 137–152, 2003.

[22] M. Neuts, *Algorithmic Probability: A Collection of Problems (Stochastic Modeling).* London: Chapman and Hall, 1995.

[23] N. C. Oguz, "Telpack version 2: Teletraffic analysis package." http://www.sice.umkc.edu/telpack/, 2002.

[24] P. Oliver and N. Benameur, "Flow level IP traffic characterization," *Seventh International Teletraffic Congress*, 2001.

[25] V. Paxson, "Empirically derived analytic models for wide-area TCP connections: Extended report," Technical Report LBL-34086, Lawrence Berkeley Laboratory, 1993.

[26] V. Paxson, "Empirically derived analytic models for wide-area TCP connections," *IEEE /ACM Transactions on Networking*, vol. 2, no. 4, pp. 316–336, 1994.

[27] V. Paxson and S. Floyd, "Wide-area traffic: The failure of Poisson modeling," *IEEE /ACM Transactions on Networking*, vol. 3, no. 3, pp. 226–244, 1995.

[28] M. Roughan and C. Kalmanek, "Pragmatic modeling of broadband access traffic," *Computer Communications*, vol. 26, no. 8, pp. 804–816, 2003.

[29] D. W. Scott, "On optimal data-based histograms," *Biometrika*, vol. 66, no. 3, pp. 605–610, 1979.

[30] A. Sinha, K. Mitchell, and D. Medhi, "Flow-level upstream traffic behavior in broadband access networks: DSL versus broadband fixed wireless," *IEEE Workshop on IP Operations and Management*, pp. 135–141, 2003.

[31] R. Sommer and A. Feldmann, "Netflow: Information loss or win?," technical report, Saarland University, Germany, 2002. `http://www.net.informatik.tu-muenchen.de/~robin/papers/`.

[32] D. Veitch, "Matlab code." `http://www.cubinlab.ee.mu.oz.au/~darryl/secondorder_code.html`.

[33] D. Veitch and P. Abry, "A statistical test for the time constancy of scaling exponents," *IEEE Transactions on Signal Processing*, vol. 49, no. 10, pp. 2325–2334, 2001.

[34] N. Vicari, S. Kohler, and J. Charzinsky, "The dependence of Internet user traffic characteristics on access speed," *IEEE Conference on Local Computer Networks*, pp. 670–677, 2000.

A Matrix Exponential Distribution and Correlation Match

Here, we briefly discuss the matrix exponential distribution and the method to create the analytic Markovian source model in which the correlation structure can be arbitrarily created leaving the marginals unchanged.

A.1 Matrix Exponential Distribution

A matrix exponential (ME) distribution [17] is defined as a probability distribution with representation $(\boldsymbol{p}, \boldsymbol{B}, \boldsymbol{e'})$, i.e.,

$$F(t) = 1 - \boldsymbol{p} \exp\left(-t\boldsymbol{B}\right) \boldsymbol{e'},\ t \geq 0, \tag{A.1}$$

where $\boldsymbol{p}$ is the starting operator for the process, $\boldsymbol{B}$ is the process rate operator, and $\boldsymbol{e'}$ is a summing operator. The n-th moment of the matrix exponential distribution is given by:

$$E[X^n] = n!\boldsymbol{p}(\boldsymbol{V})^n\boldsymbol{e'}, \tag{A.2}$$

where $\boldsymbol{V}$ is the inverse of $\boldsymbol{B}$.

The class of matrix exponential distributions is identical to the class of distributions that possess a rational Laplace-Stieltjes transform. A representation is not unique and the only limitations on $(\boldsymbol{p}, \boldsymbol{B}, \boldsymbol{e'})$ stem from the requirement that $F(t)$ must form a distribution function.

The true power of LAQT is the ability to choose a purely algorithmic representation. This creates a great deal of freedom in the algebraic manipulation of these processes as demonstrated in the next subsection describing the correlation matching process.

A.2 Constructing a Matrix Exponential Sequence Using Correlation Matching

The arrival and service processes we study are not renewal processes. We model these processes as a sequence of ME random variables $T_1, T_2, \cdots$ such that the joint probability *density* over any finite sequence of consecutive inter-event times is given by:

$$f_{T_1,\cdots,T_n}(t_1,\ldots,t_n) = \boldsymbol{\pi}(0)\exp(-\boldsymbol{B}t_1)\,\boldsymbol{L}\cdots\exp(-\boldsymbol{B}t_n)\,\boldsymbol{L}\boldsymbol{e'},$$

where $\boldsymbol{\pi}(t)$ is a vector representing the internal state of the process at time t. The matrix $\boldsymbol{L}$ is the (non-zero) event transition rate operator that generates an event and starts the next interval in the appropriate starting state. This process is interpreted as a stream of events (in this case, successive departures from the matrix exponential process representing handoffs or correlated arrivals) occurring at times $t_1, t_1+t_2, t_1+t_2+t_3, \ldots$, and inter-event times $t_1, t_2, \ldots$. Note that this process is similar to the Markov Arrival Process (MAP) described by Neuts [22, pp. 393–398], where $-\boldsymbol{B} = \boldsymbol{A}_0$ and $\boldsymbol{L} = \boldsymbol{A}_1$. The closed form representation of the state of the system at time t is

$$\boldsymbol{\pi}(t) = \boldsymbol{\pi}(0)\exp(\boldsymbol{Q}^*t), \text{for } t > 0, \tag{A.3}$$

where the generator matrix $\boldsymbol{Q}^*$ is

$$\boldsymbol{Q}^* = \boldsymbol{L} - \boldsymbol{B}. \tag{A.4}$$

See [18,20–22] for more detailed derivations. We assume that the process is in equilibrium, the steady state being represented by the vector $\boldsymbol{\pi}$. In equilibrium,

the starting vector $\boldsymbol{\pi}(0)$ is the conditional steady state at embedded arrival points and is denoted by the vector $\boldsymbol{p}$,

$$\boldsymbol{p} = \boldsymbol{\pi}(0) = \frac{\boldsymbol{\pi L}}{\boldsymbol{\pi L e'}}. \tag{A.5}$$

The moments of the process can be determined by using equation (A.2). Note that $\boldsymbol{L}$ only appears in the autocorrelation equation and not in the moment equation. A renewal process can be expressed in terms of uncorrelated sequence of matrix exponential by using the following:

$$\boldsymbol{L} = \boldsymbol{B e' p} \tag{A.6}$$

The covariance of the sequence of matrix exponentials is given as

$$\operatorname{cov}[X_0, X_k] = \boldsymbol{pV}\left(\boldsymbol{VL}\right)^k \boldsymbol{Ve'} - \left(\boldsymbol{pVe'}\right)^2. \tag{A.7}$$

See Lipsky [18] for a complete derivation of the above equations or [21,22] for a short review. If the process is covariance stationary, the autocorrelation can be obtained by dividing $\operatorname{cov}[X_0, X_k]$ by the variance.

Mitchell has been able to extend the work of Van de Liefvoort by constructing correlated point processes that retain the marginal distributions created by Van de Liefvoort's algorithm [21]. By examining equations (A.2) and (A.7), it can be seen that constructing a correlated point process with invariant marginals involves finding an $\boldsymbol{L}$ such that $\boldsymbol{p}$ and $\boldsymbol{B}$ remain invariant. This is due to the fact that $\boldsymbol{L}$ does not appear in (A.2). Define the matrix $\boldsymbol{Y}$ as

$$\boldsymbol{Y} = \boldsymbol{VL}, \tag{A.8}$$

where $\boldsymbol{V} = \boldsymbol{B}^{-1}$. In order to change $\boldsymbol{Y}$ and leave the marginals invariant, the following conditions must be met

$$\boldsymbol{pY} = \boldsymbol{p}, \tag{A.9}$$

$$\boldsymbol{Ye'} = \boldsymbol{e'}. \tag{A.10}$$

This ensures that $\boldsymbol{pY}^k \exp\left(-\boldsymbol{B}t\right) \boldsymbol{Be'}$ is a density for all $k = 0, 1, 2, \ldots$.

We are interested in developing parsimonious models, so we have a single parameter β, leading to geometrically decaying covariance. Define operator $\boldsymbol{L}_\beta$ by:

$$\boldsymbol{L}_\beta = (1-\beta)\boldsymbol{Q}^* + \boldsymbol{B}, 0 \leq \beta < 1. \tag{A.11}$$

By using equations (A.4) and (A.6), we can rewrite equation (A.11) as

$$\boldsymbol{L}_\beta = (1-\beta)(\boldsymbol{Be'p} - \boldsymbol{B}) + \boldsymbol{B}. \tag{A.12}$$

From equation (A.8) the above equation reduces to:

$$\boldsymbol{Y_\beta} = (1-\beta)(\boldsymbol{e'p}) + \beta\boldsymbol{I}, 0 \leq \beta < 1. \tag{A.13}$$

Thus, we can easily construct a second-order representation with three normalized moments and a single lag autocorrelation. In this form, if needed, an autocorrelated process can be constructed from an arbitrary number of moments and a sequence of lag-k autocorrelations.