The Distribution of File Transmission Duration in the Web[*]

Ronit Nossenson [†] Hagit Attiya [‡]

Abstract

It is well known that the distribution of files transmission duration in the Web is *heavy-tailed* [9]. This paper attempts to understand the reasons for this phenomenon by isolating the three major factors influencing the transmission duration: file size, network conditions and server load.

We present evidence that the transmission-duration distribution of the *same* file from the same server to the same client in the Web is Pareto and therefore heavy tailed. Furthermore, transmission delay for a specific client/server pair is not significantly affected by file size: all files transmitted from the same server to the same client have a very similar transmission duration distribution, regardless of their size.

We use simulations to estimate the impact of network conditions and server load on the transmission-duration distribution. When the server and the client are on the same network, the transmission duration distribution of each file is usually Pareto as well (for server files and client requests that are distributed in a realistic way).

By examining a wide-area network situation, we conclude that the network conditions do not have a major influence on the heavy-tailed behavior of transmission-duration distribution. In contrast, the server load is shown to have a significant impact on the high variability of the distribution.

1 Introduction

Understanding Web traffic's behavior is necessary in order to improve the quality of service experienced by Web users. In particular, it is important to study statistical characteristics of Web servers, as was done, for example, in [3, 4, 8, 9, 12, 19].

Describing a Web server as a stochastic queuing system requires knowledge about the distribution of service times. There are two major types of service in a Web server: *Static* service involves file transfer, while *dynamic* service includes additional work of answering client queries. In this paper, we restrict our attention to file transfer service experienced by the client (that is, the Round Trip Time), where the service time distribution is characterized by the files' transmission-duration distribution (in short, TDD).

Crovella et al. [9] present evidence that the distribution of files-transmission duration distribution is *heavy-tailed.* Roughly speaking, most of the samples in a heavy-tailed distribution are small but the probability for the appearance of a very large sample is non-negligible. A prominent example for heavy-tailed distribution is the *Pareto* distribution with *shape parameter* α and *location parameter* k, which has the following cumulative distribution function [11].

$$F(x) = P[X \leq x] = 1-(k/x)^{-\alpha}, \;\; \alpha,\, k > 0,\; x \geq k,$$

Heavy-tailed distributions are characterized by extremely high variability. Heavy-tailed distributions behave very different from exponential distributions, often used to analyze server systems, and their high variability has dramatic negative effects on the performance (see [5, 9, 19]).

---

[*] This research was supported by the *Israel Science Foundation* (grant number 105/01) and by Technion V.P.R. fund – I. Goldberg fund for electronics Research.

[†] Dept. of Computer Science, Technion, Haifa, 32000, Israel, E-mail: `ronitt@cs.technion.ac.il`

[‡] Dept. of Computer Science, Technion, Haifa, 32000, Israel, E-mail: `hagit@cs.technion.ac.il`

This paper studies the reasons for the heavy-tailed distribution of files transmission duration. Understanding the causes for the high variability of the TDD may help to design systems that reduce it. In addition, such understanding provides better tools for modeling the service time when analyzing server systems.

Three major factors determine the duration of file transfer: the file size, the server load, and the network conditions (server-client distance, and the load on the routes between them).

We start with a more refined analysis of a known Web trace, collected at Boston University [9], that reveals the TDD of the *same* file from the same server to the same client in the Web is Pareto and therefore heavy-tailed. We emphasize that this is not the same as analyzing as the collective distribution of the duration of transmitting *all* files. For example, if we transmit two files, 10 times each, then all 20 samples should be considered in order to estimate the TDD for all files, while only 10 relevant samples should be considered in order to estimate the TDD of a single file.

Since the TDD of the same file from the same server to the same client in the Web is heavy-tailed, every file request in the Web has a non-negligible probability for a very long completion time; this holds even for very small files. In addition, this empirical result means that every Web server, including servers that serve a single static page, suffers from bursty service time.

We also compare the TDD of all individual files transmitted from the same server to the same client in the BU trace. Interestingly, almost all the transmission durations of each of these files were distributed in a very similar way (that is, Pareto with *similar* shape parameter and *similar* location parameter). These distributions stayed similar for files with different order of sizes (from several bytes to more than 100KB), covering most text files.

This demonstrates that the distribution of file sizes, although heavy-tailed by itself, is not the major cause for the heavy-tailed behavior of the TDD. To confirm these findings and to isolate network conditions from server load, we performed controlled simulations on a local network and over wide-area connections. These simulations also provide a new, more up-to-date data set.

Our first simulation neutralized the effects of an outside network, by locating the server and the clients on the same local network. The size of the server files and their popularity are distributed in a realistic way (following [3]), varying from several bytes to more than 140MB, covering text and multi-media files.

The empirical result is that for almost all these files the transmission durations is Pareto. The shape parameters of their TDDs were very similar (that is, not correlated with the file sizes) and highly correlated with the server load. In contrast, the location parameters are highly correlated with the file sizes and not correlated with the server load.

We also run a simulation on a wide-area network, including a TCP server responding to requests made by clients at several remote locations. The file sizes vary from 1KB to 150KB, covering most text files. The simulation implies the following conclusions:

- When the influence of the file sizes is isolated, by comparing different files transmitted from the same server to the same client, we found that the TDDs of all files are Pareto, with similar shape parameters and location parameters that increase with file size.
- When the influence of the route is isolated, by comparing the same file transmitted over routes with different quality parameters, we found that the TDD of all files is Pareto, with similar shape parameters, but different location parameters.
- When the influence of the server load is isolated, by comparing the same file transmitted by the sever server, under different loads, to the same client, we found that both the shape parameters and the location parameters of the TDD are different.

Since the shape parameter is responsible for the heavy-tail of the Pareto distribution, we conclude that the server load is a major cause for the high variability of the TDD in our simulation. This important fact explains why even local service (that is, service provided by servers on the local network) suffers from negative effects of heavy-tailed distributions on the performance [14, 15].

Since the file sizes are highly correlated with the location parameter only, we conclude that the high variability of file sizes in Web servers is not a major cause for the heavy-tailed behavior of the TDD. Indeed, no correlation was found between file sizes and transmission-durations distribution, as specifically mentioned in [9] (Section 3.2 and Figure 9), Nevertheless, it is a common mistake to assume that the heavy-tailed file-sizes distribution is the major cause for the heavy-tailed TDD (for example, [10, 18, 16]).

2 Further Analysis of the BU Trace

2.1 Data Collection

We measure the TDD of files from the client point of view. Namely, the time interval from the moment the Proxy starts sending the client's request to the server until the moment the Proxy completes receiving the response from the server.

We start by further examination of the trace collected in Boston University, in April-May 1998 (see [2, 6, 7]). The trace data was collected in the Computer Science Department's undergraduate workstation lab [6], consisting of a cluster of 37 Sun SPARC station 2's running Solaris. The trace contains more than 72K requests, collected using lightweight non-caching HTTP/1.1 proxy servers running on each workstation.

This is the most recent public trace (known to us) that combines two essential properties for our experiments. First, the anonymity was achieved by hashing each column of the data set and not each line. This preserves the ability to compare members of a column for equality without leaking information regarding those members identities. For example, we can know that the same server sends file x and file y, without revealing the server's identity or the files. The second essential property is that the trace allows to calculate the duration of each request. Usually, traces supply for each request its start time but not its termination time.

Another important fact regarding the data is that we can not use all transactions in the data set. Since we need statistically reliable samples, we used only information on files which were transmitted at least 30 times. Although such files are not very common, their number and the way which their information was gathered supply us a statistically reliable sample.

2.2 The TDD of a file in the Web

In this section we analyze the TDD of a file transmitted from the same server to the same client. Our first empirical result presents evidence that the TDD of the same file from the same server to the same client is Pareto and therefore heavy tailed.

In order to find out the TDD of each file, we analyzed the BU trace data. We considered only files that have been transmitted more than 30 times. The TDD of each file turned out to be a Pareto distribution with shape parameter α that varies between 0.6 to 1.8.

Figure 1 shows an example of a very popular small file, requested more than 700 times. The figure plots the samples against the estimated distribution; The graph's X axis indicates the transmission duration, and the Y axis indicates the estimated $\bar{F}(x) = 1 - F(x)$. Note that $\bar{F}(x) = (k/x)^{-\alpha}$ for the Pareto distribution. (Additional examples can be found in [17].)

2.3 The TDD of files transmitted by the same server

In this section we analyze the TDD of files transmitted by the same server to the same client. Again, we analyze the TDD of each file separately and compare the results for all these files. As mentioned in Section 2.2, the TDD of each file is Pareto.

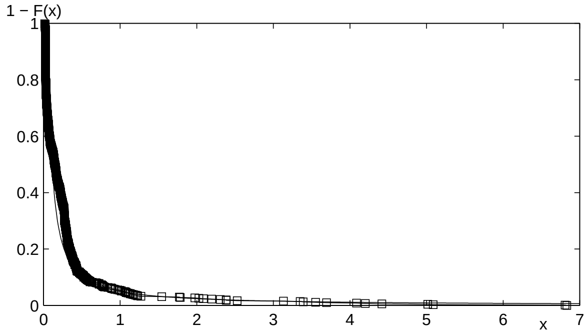

Figure 1: The distribution of file transmission duration in the BU trace: fitting Pareto distribution (with shape parameter 1.05 and location parameter 0.05) to empirical data (the file size is 68 bytes, the number of samples is 706).

Here we are testing two issues (i) the relation between the TDD location parameters of files transmitted by the same server to the same client, and (ii) the relation between the TDD shape parameters of these files. Since the location parameter represents the smallest sample, our intuition was that the location parameters would probably differ according to the files' size in contrast to the shape parameters that will probably be similar.

Surprisingly, this intuition is not fully correct. The BU trace contains evidence that for files transmitted from the same server to the same client, both the shape parameters and the location parameters are very similar. These distributions stayed similar even for files with different order of sizes (from several bytes to more than 100KB), that cover most text files.

In order to find the relation between the corresponding shape parameters and location parameters we processed the results of Section 2.2 further. The files were divided into groups according to the transmitting server. The most reliable sample consists of a group of 20 files, each file transferred more than 30 times. Although this set is not big, its size and the way which the information was gathered supply us a statistically reliable sample. Recall that to show that a specific property exists in a small statistically reliable sample set is a stronger result than to show that this property exists in a large set.

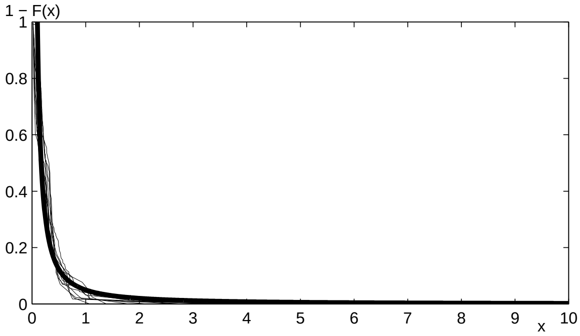

Figure 2: The transmission duration distribution of files transmitted by the a single server in the BU trace, plotted against Pareto distribution (with shape parameter 1.3 and location parameter 0.05).

Figure 2 shows for each file, its samples plotted against the specified distribution. The graph's X axis indicates the transmission duration, and the Y axis indicates the estimated $\bar{F}(x) = 1 - F(x)$.

Figures 3 and 4 summarize the results. The graphs are plotted with a logarithmic X axis, due to the high variance in the file sizes. As can be seen from Figure 3, 15 files have the same shape parameter $\alpha = 1.3$. Another interesting result which can be deduced from Figure 4 is that even the location parameter is not highly related to the file size. The location parameter of the files is in the range $[0.035, 0.7]$ and most of them (18) have location parameter in the range $[0.035, 0.1]$.

This experiment implies that for a client/server pair and ordinary text files (with size between several bytes up to around 100 KB), the TDD is not significantly affected by the file size. The consequence is that the TDD is highly correlated with the network condition and the server load.

3 Our Simulation Results

Our analysis of the BU trace, described in the previous section, indicates that the high variability of the distribution of file transmission duration is not caused by the high variability of file sizes. It remains to isolate the influence of net-

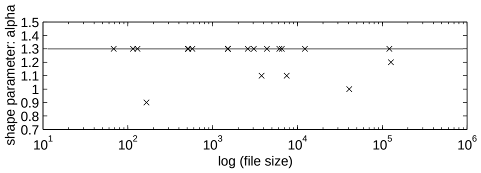

Figure 3: The shape parameters of transmission duration distributions of all files transmitted by a single server in the BU trace.

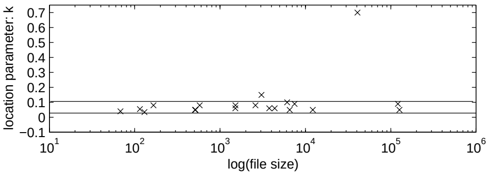

Figure 4: The location parameters of transmission duration distributions of all files transmitted by a single server in the BU trace.

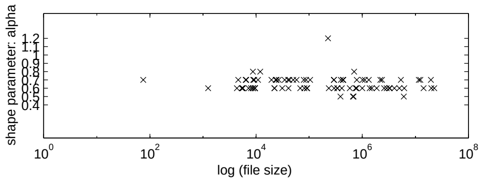

Figure 5: The shape parameters of transmission duration distributions of files transmitted by the simulation server on a local network (the smallest 84 files).

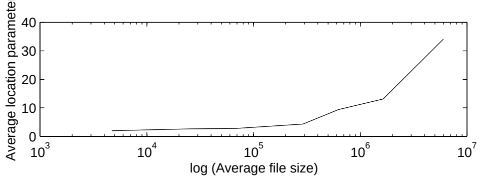

Figure 6: The location parameters of transmission duration distributions of files transmitted by the simulation server on a local network (the smallest 84 files).

work condition and server load on the TDD. At a university site, it is impossible to collect up-to-date, statistically reliable traces that contains sufficient information about server load and network conditions. Instead, we conducted several extensive simulations, which are described in this section.

3.1 The TDD in a Local Network

The first simulation was meant to corroborate the results of our analysis of the BU trace. The experiment, performed at the Technion's distributed systems laboratory, consists of one Apache server and three clients running the simulation code. The data was collected from the clients' log files.

The server has 100 different files, their sizes and popularity distributed according to [3]. The files sizes varies from several bytes to more than 140 MB. Each file was transmitted more than 30 times. The number of the transmissions and their generation method provides a statistically reliable sample.

The simulation statistics shows that for the smallest 84 files (with sizes in the range from 75 bytes to 22 MB), the TDD was Pareto and therefore heavy-tailed. The TDDs of the largest 16 files (with sizes of $26 - 140$ MB) although still clearly with heavy-tailed distribution, fit a littlest to the Pareto distribution.

Figures 5 and 6 summarize the results. The graphs are plotted with a logarithmic X axis, due to the high variance in the file sizes.

As can be seen in Figure 6, there was a high correlation between the location parameter and the file size. The location parameter varies in the range $[1.4, 90]$.

Considering all the simulation executions data, the shape parameter varies in the range $[0.4, 1.2]$. For most of the files (82) the shape parameter is in the range $[0.6, 0.7]$ (see Figure 5). There is no correlation between the shape parameter and the file size.

3.2 TDD in a Wide-Area Network

The local network simulation neutralizes the effects of network conditions. Its results indicate that file size is not the major factor influencing the heavy-tailed behavior of the TDD, as demonstrated by the shape parameter, leaving the server load as the main cause for the heavy-tail of the TDD.

This motivated a more elaborate study in a wide-area setting. Our simulation consists of a single TCP server, located at the Technion, and

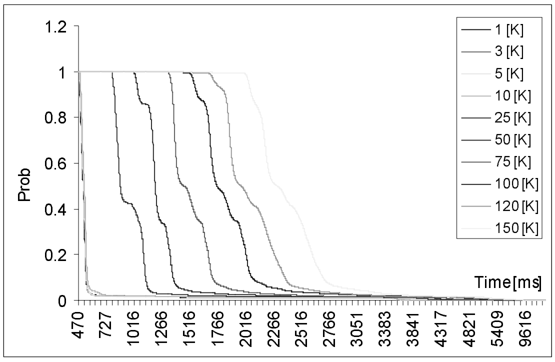

Figure 7: Impact of file size on the TDD of files transmitted between a specific server / client pair.

clients in several remote locations around the world, with very different route conditions.

To study the influence of the file sizes, we compare the TDD of different files transmitted from the same server to the same client. In the simulation, we generate a file request every two seconds and collect the transmission duration statistics. Such low frequency of requests' generation hardly influences the route and the server load.

The file sizes vary from 1KB to 150KB, covering most text files. To increase the reliability, we performed the simulation several times for several clients around the world.

Our simulation results appear in Figure 7. The TDDs of all files are Pareto, with similar shape parameters and location parameters that increase with the file size.

To study the influence of the route on the files TDD, we compare the TDD of the same file, transmitted from the same server to different clients. The server is located in Israel, while the clients are located in East and West USA, Portugal, Brazil, Australia and Taiwan. These routes are known from previous research to have different quality parameters [1].

The results appear in Figure 8. The TDD of all files is Pareto with similar shape parameters and different location parameters.

To study the influence of the server load on the files TDD, we compare the TDD of the same file, transmitted from the same server, working under different loads, to the same client. We examine five different load levels, from 'no-load' to 'high-load'. The server load is increased by adding more local clients. The simulation results appear in Figure 9. Both the shape parameters and the location parameters are different.

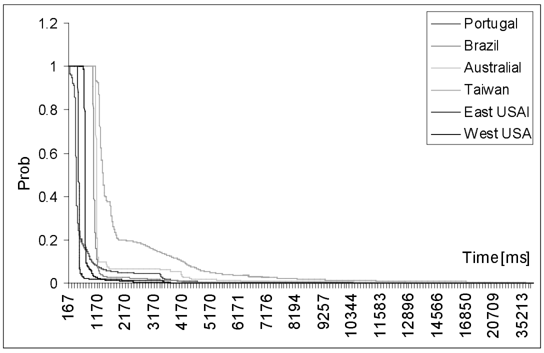

Figure 8: Impact of route selection on the TDD.

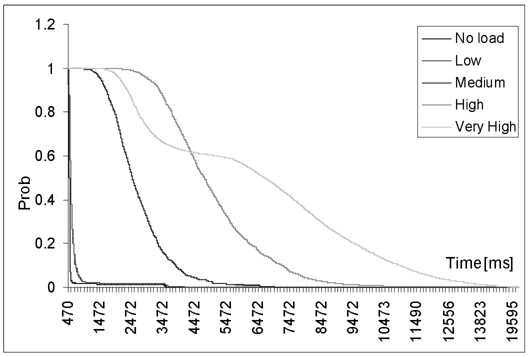

Figure 9: Impact of server load on the TDD.

These results validate our expectation that while file sizes and network conditions affect the absolute values of the transmission duration, the server load is, by far, the major cause for the high variance of its distribution. The next section attempts to explain these findings.

4 Interpreting the Empirical Results

Our results show extremely high variability in the time that a specific server needs to send the *same* file to a specific client.

One set of explanations for this phenomenon can be derived from the network design and in particular, from caching mechanisms. Either file caching hits or server-name resolving

(DNS) caching hits yield much faster transmissions. Intermediate caching over the network creates shortcuts and speeds up the file transmission. The difference in the transmission duration of file from an intermediate cache versus the transmission duration from the server might amount to several orders of magnitude. In addition, bottlenecks around the request route, due to overloaded routers or link failures, might cause severe delays.

The other possible explanation for this phenomenon can be derived from the server load. The server load might change dramatically in different hours and different days [13]. This might cause high variation (manifested as a heavy-tailed behavior) in the transmission duration of the same file transmitted in different times. Indeed, Barford and Crovella [4] show that high server load generates a significant gap between the transmission of connection setup packets and the first data packet flowing from the server.

The document request arrival process in Web servers shows *self-similar* behavior [13]: it is bursty over a wide range of time-scales. Using Process Sharing discipline, the server responds to all arrival requests, and rarely rejects requests. As a result, with the increased load, the server response (observed by the client) become slower. Thus, the result of a bursty server load is high variation in the response time, which might cause heavy-tailed behavior. The same phenomenon occurs when the server uses a First Come First Serve discipline and has a long queue: with the increased load, many requests arrive and enter the long queue, each client's request waits a long time in the queue and the client experiences a long response time.

Long request service times (that is, long transmission durations) might also cause a heavy server load. For example, assume that the server uses a First Come First Serve discipline and the server is busy for a long time with a long transmission. All other client's request waits a long time in the queue (even in a short queue) and the client experiences a long response time. The consequence is that long request service times of some requests might yield long response times at other requests. There is a non-negligible probability for the appearance of a requests with very large service times, which yield a heavy load server.

The server load might also explain why the transmission durations of almost all the files is Pareto distributed, even when network effects are neutralized. This happens in our local network simulation, where the transmission duration distribution of different files have similar shape parameters, which are highly correlated with the server load.

When considering files transmitted from the same server to the same client in the BU data, the small effect that file sizes do have on the TDD seems to vanish. For files with size between several bytes up to around 100 KB (that is, several IP-packets up to several hundred IP-packets), almost all the transmission durations have *similar* shape parameter and *similar* location parameter (for their Pareto distributions).

Findings of Cohen and Kaplan [8] reveal that DNS query times, TCP connection establishment times, and start-of-session delays at HTTP servers times, are the major causes of long waits, more than the actual transmission time. Note that, for every request made by the same client to the same server, the time spent on these operations is identically distributed regardless of the requested file size.

Acknowledgments: The authors wish to thank Mark Crovella for providing the BU data trace, and Roy Friedman, Danny Raz and Roman Vitenberg for useful discussions. Simulations at the Technion were performed by Idan Zach, Vadim Ryvchin, Nitzan Kashi and Yury Mogilevsky.

References

[1] H. Atzmon. Relacement policies for internet-wide caching of distributed objects. Master's thesis, Department of Computer Science, The Technion, 2002.

[2] P. Barford, A. Bestavros, A. Bradley, and M. E. Crovella. Changes in Web client access patterns: Characteristics and caching implications.

World Wide Web (special issue on Characterization and Performance Evaluation), 1999.

[3] P. Barford and M. Crovella. Generating representative Web workloads for network and server performance evaluation. In *Measurement and Modeling of Computer Systems SIGMETRICS '98 conference*, pages 151–160, June 1998.

[4] P. Barford and M. E. Crovella. Measuring web performance in the wide area. *Performance Evaluation Review*, Aug. 1999.

[5] O. J. Boxma and J. W. Cohen. The M/G/1 queue with heavy-tailed service time distribution. *IEEE Journal on Selected Areas in Communications*, 16(5):749–763, June 1998.

[6] A. Bradley. Bu computer science 1998 proxy trace. Technical Report BU-CS-99-011, Boston University Department of Computer Science, 1999.

[7] E. Carrera and R. Bianchini. Analytical and experimental evaluation of cluster-based network servers. Technical Report 718, Department of Computer Science, University of Rochester, 1999.

[8] E. Cohen and H. Kaplan. Prefetching the means for document transfer: a new approach for reducing web latency. In *19th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, volume 2, pages 854–863, 2000.

[9] M. Crovella, M. Taqqu, and A. Bestavros. *Heavy-Tailed Probability Distributions in the World Wide Web.* Appears in the book: A practical Guide To Heavy Tails: Statistical Techniques and Application, R. Adler, R. Feldman and M. S. Taqqu, editors, Birkhauser, Boston, 1998.

[10] A. B. Downey. The structural cause of file size distributions. In *SIGMETRICS/Performance*, pages 328–329, 2001.

[11] T. Hettmansperger and M. Keenan. *Tailweight, Statistical Inference, and Families of Distributions - A Brief Survey*, volume 1. in Statistical Distributions in Scientific Work, Patil, G.P., et al. (eds), Kluwer, Boston, 1980.

[12] B. A. Huberman, P. L. T. Pirolli, J. E. Pitkow, and R. M. Lukose. Strong regularities in World Wide Web surfing. *Science*, 280(5360):95–97, Apr. 1998.

[13] K. Kant and Y. Won. Server capacity planning for web traffic workload. *IEEE transaction on knowledge and data engineering*, 11(5):731–747, Oct. 1999.

[14] W. Leland, M. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, pages 1–15, Feb. 1994.

[15] W. E. Leland and D. V. Wilson. High time-resolution measurement and analysis of LAN traffic: Implications for LAN interconnection. *Proc. IEEE INFOCOM'91*, 3:1360–1366, 1991.

[16] A. V. Z. K. S. Molnar and G. Vattay. The propagation of long-range dependence in the internet. In *ACM SIGCOMM*, 2000.

[17] R. Nossenson and H. Attiya. The distribution of file transmission duration in the web. Technical Report CS-2002-09, Department of Computer Science, Technion, 2002.

[18] T. M. S. R. H. Rootzen and A. Stegeman. Is network traffic approximated by stable Levy motion or fractional Brownian motion? *Ann. Appl. Probab.*, to appear, 2002. available at www.orie.cornell.edu/trlist/trlist.html.

[19] D. Shuang. Empirical model of www document arrivals access link. In *IEEE International Communications Conference - ICC'96*, pages 17–23, June 1996.