

analysis of traces. This work did not present or test a Web traffic model.

In [10], Badford and Crovella built a Web model called "SURGE", based on the parameters: 1) Distribution of object size on the server, 2) Distribution of the size of requested objects, 3) Object popularity, 4) Number of in-line objects, 5) Temporal locality, and 6) User viewing time. The model showed self-similarity as evidenced by the variance-time plot. The browser "Mosaic" was modified to record user activities in the client. The instrumented browser was distributed to 37 clients in Boston University. This represents a relatively small cross-section of users, and because the browser on the client side is modified it is difficult to measure other user communities. Further, the configured HTTP version was 0.9, which was used before 1996. Because Mosaic is no longer evolving, this work can not be extended to current versions of HTTP.

Deng [8] measured Web traffic characteristics of individual subscribers and proposed a two-state ON/OFF model for the arrival process at the access link. The model did not simulate the detailed interaction of Web traffic and the accuracy of the model was not evaluated.

III. OVERVIEW

There are four different versions of HTTP currently available. In what might be called "pure" Version 1.0 [1] objects are downloaded back-to-back with each object requiring one TCP connection. In Version 1.0 with multiple connections, the browser opens multiple parallel connections to download objects for the earliest display of the page. The browser sets the limit on the number of multiple connections. Objects beyond these limits are downloaded after completing one of the outstanding connections. In Version 1.0 with "Keep-alive", multiple connections are possible, but a connection is not closed immediately on the chance that a new request for the connection will arrive before a time-out. Version 1.1 [2] permits persistent connections and requests are pipelined. The persistent connection is very similar to the Keep-alive connection, the exception being for a proxy.

A typical Web page consists of a Hypertext document with links to other objects that make up the whole page. An *object* is an entity stored on a server as a file. There are two kinds of objects, a main object and an in-line object. The file containing an HTML document is referred to as a *main object* and the objects linked from the Hypertext document are referred to as *in-line objects*.

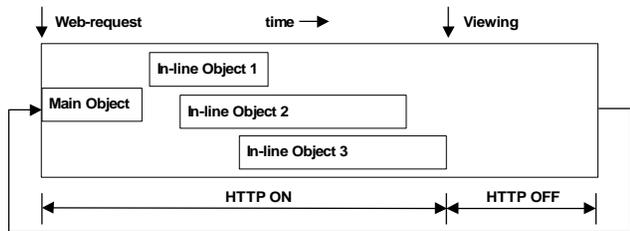


Figure 1: Overview of the basic model

The basic model of Web traffic is shown in Figure 1. A new Web-request is immediately generated after expiration of the viewing period. The model simulates an ON/OFF source

where the ON state represents the activity of a Web-request and the OFF state represents a silent period after all objects in a Web-request are retrieved. The duration of On state and Off state correspond to On-time and viewing time, respectively. Viewing time denotes any time that the browser is inactive. Viewing time includes, for instance, the situation where the browser is iconized while a user is working with another application. On-time is the time taken to fetch all objects in a Web-request. The ON state can be split into successive TCP connections used to deliver objects. The parallel connections for in-line objects are opened consecutively after the single connection for the main object.

IV. MEASUREMENT AND ANALYSIS

A. Traffic Measurement and Parsing

We would like to measure the traffic close to a browser source because we are modeling the traffic sent by, and to, the browser. At the same time, we want a large cross-section of traffic. We meet these objectives by recording a trace of traffic on the backbone network of the Georgia Tech campus (see Figure 2). The Georgia Tech campus network is composed of two B-class IP addresses (130.207 and 128.61) and a number of C-class IP addresses (about 170). We are able to record a large cross-section of Web interaction from across the campus. We further filter the traffic to gather only those sessions that originate on the campus and terminate elsewhere in the network. In this way we expect to obtain model parameters that are not unduly affected by the behavior of campus servers. Further, we exclude machine-generated traffic such as Web-crawler traffic. Traffic from major search engines is filtered out because it originates from outside of the campus. The traffic from the campus search engines is not recorded because it terminates inside the campus.

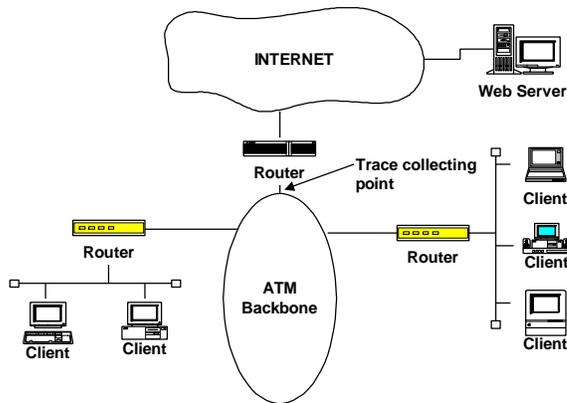


Figure 2: Perspective of the Georgia Tech campus network.

We use primarily a trace that was collected from 11 A.M. to 12 P.M. on Wednesday October 7 1998 running on a Sun Ultrasparc2 (180 MHz) workstation. More than 1900 clients participated in Web browsing sessions producing about 24,000 Web-requests. The details of the trace are listed in Table 1. We regularly record traces on the campus backbone and this particular trace is typical. We do not claim that our traffic source is representative of Web traffic in general. However, we are able to develop and evaluate a methodology that can be extended to other traffic sources.

Date	Number of Clients	HTTP Status Codes				
		1xx	2xx	3xx	4xx	5xx
10/7/98	1934	197	138948	39903	1594	2488

Number of Web-requests	Methods		
	GET	POST	HEAD
24014	448810	5811	320

Table 1: Summary of the trace

We use two tools to collect and parse data, Tcpcmdump [4] and Tcpshow [5]. The binary-mode option was set while Tcpcmdump was running. Tcpcmdump recorded TCP/IP headers as well as 300-byte TCP payloads. The 300 bytes of the TCP payload are large enough to capture the HTTP request header and the HTTP response header within the range of interest. For our study, the HTTP header is important for obtaining additional information about Web traffic. The HTTP header information is used in the parsing and analysis phase to separate Web-requests, to decide if the connection is Keep-alive connection or Close connection and to check if the object is cached.

Tcpshow interprets a binary-mode trace. From the binary-mode trace, two new traces are created in off-line processing. One contains the HTTP header information (HTTP trace). The other contains the TCP header information (TCP trace). The parameters are categorized into the HTTP layer and the TCP layer depending upon associated traces. HTTP-layer parameters are searched in the HTTP trace and TCP-layer parameters are searched in the TCP trace.

The parser script, written in *Practical Extraction and Report Language* (PERL), sorts out the HTTP trace by the client IP addresses and checks if the client is inside the campus. The sorted trace for a single client is separated into distinct Web-requests using the technique discussed in Section B. The start and end times of the Web-requests are then recorded and used to parse parameters in the TCP trace. Empirical distributions of HTTP-layer parameters are obtained. In the TCP trace, the parser script parses TCP-layer parameters based on the boundaries that were recorded when HTTP-layer parameters were parsed.

Once empirical distributions of the individual parameters are obtained, we compare each distribution with different standard probability distributions and select the best fit. The *Quantile-Quantile plot* (Q-Q plot) [11] is used to test the fit of the data to the model. If the model fits the data perfectly then the plotted points lie on a straight line. The best standard probability distribution is determined to be the one that minimizes the root-mean-square of the deviation from a straight line. We select the best distribution from among Weibull, Lognormal, Gamma, Chi-square, Pareto and Exponential (Geometric) distributions.

B. Web-request

A Web-request is a page or a set of pages resulting from a request by a user. By definition, (1) a Web-request is initiated by a human action and (2) the first object in a Web-request is an HTML document. Accurate identification of Web-requests is essential for our model to be accurate. We determine the Web-requests by setting up rules for a request. We apply

these rules to our parsing of the trace to identify Web-requests. Our rules are:

- For simultaneously requested multiple pages by a user from multiple browsers but from the same client, each request represents a Web-request by definition (1). In this case, the second or later Web-requests might include an object belonging to a page of previous Web-requests. If the Web page were a basic unit of the model, this inclusion would skew the model parameters. It does not matter if the basic unit is a Web-request because the boundary is determined by whether or not we have a user request.
- If a single request generates multiple pages (*e.g.* frame and Java-script), they belong to the same Web-request. However if subsequent pages are retrieved by a user from within the frame they would represent a complete Web-request by definition (1). We do not consider an automatically redirected page (HTTP status code 301 and 302) to be a Web-request.
- If a user clicks a hyperlink of a single object¹ such as an image (.jpg, .gif), a sound (.avi or .mp3) or a text document (.txt, .ps or .pdf), these single objects do not represent a Web-request by themselves (2). That's because the first object is other than an HTML document. Instead, they are included in the Web-request as in-line objects.

total	htm	/	cgi	asp	sml	stm	other
28833	12340	4515	4425	1231	59	58	6205

Table 2: List of HTML document extensions.

In order to determine a boundary of a Web-request, we take advantage of information in the HTTP header by inspecting the extension of the requested objects, the MIME type of the response or both. The extension indirectly implies the contents of an object. Objects are mostly named according to their type; for instance, most graphical images are named “.gif” or “.jpg” and the most popular extension for an HTML document is “.htm[l]”. This tendency is further enforced if the page is designed by Web publishing utilities². Table 2 shows the list of extensions of HTML documents found in our traces. Other in the last column of Table 2 is mainly due to query requests whose URIs are *WWW-URL-encoded*.

Extensions of objects do not always correlate with the contents of objects. To prevent inadvertently missing a Web-request by relying on the extension, the MIME type in the HTTP response header is also checked. The MIME type directly implies the type of an object. “text/html” is the reserved MIME type for an HTML document by *Internet Assigned Number Authority* (IANA).

In summary, a request becomes a Web-request:

- If it is used to request for an object whose extension contains either “.htm”, “.asp” or “.cgi”. A URI that finishes with “/” implies “index.html” in the directory, also becomes a Web-request.

¹ It is difficult to distinguish a hyperlinked single object from a regular in-line object because their trace records look the same. We could have used the HTTP header information, “referral”, to distinguish these objects. Because “referral” does not always guarantee to distinguish these two cases, we decided not to use it.

² Microsoft’s Frontpage 98, Adobe’s Pagemill and Netscape’s Composer. They use “.htm” as the default extension.

- If it results in a response of MIME type “text/html” and the HTTP status code 200 (OK).

Parameters		Mean	Median	S.D.	Best-fit
Request size		360.4	344	106.5	LN
Object size	Main	10710	6094	25032	LN
	In-line	7758	1931	126168	LN
Parsing time		0.13	0.06	0.187	G
Number of In-line objects		5.55	2	11.4	G
In-line Inter-Arrival time		0.86	0.17	2.15	G
Viewing (OFF) time		39.5	11.7	92.6	W
Number of Web-requests	Non-cached	12.6	5	21.6	LN
	Cached	1.7	1	1.7	GM

Table 3: Summary statistics for HTTP parameters (LN=Lognormal, G=Gamma, W=Weibull and GM=Geometric)

V. MODEL

From our traces and their analysis we have derived the parameters of our model. In the following we indicate a parameter by bold face.

A. HTTP Model

Statistics of HTTP parameters are shown in Table 3. **Number of in-line objects** is the number of in-line objects in a Web-request. The in-line objects that we count are only those which need to be downloaded. A requested object does not need to be downloaded if it is found in the cache with valid time-stamps. Hence **number of in-line objects** is always less than or equal to the total number of objects. The mean **number of in-line objects** is 5.55, almost three times larger than the mean observed in [7]. This may be explained by the following:

- The number of multimedia objects in a page is increasing with time, as pages become more complex.
- The hit ratio of the local cache has dropped. Because the number of Web servers has increased since the work [7] was done, the number of Web servers that a user accesses has also increased, so that a user’s access pattern becomes less predictable.
- The technique that was used to separate pages in [7] underestimated the parameter as discussed in Section II.

Viewing time is the inactive interval between Web-requests. The viewing time parameter measured in previous work [7], by definition, was always greater than one second. However our histogram of **viewing time** in Figure 4.b shows a large number of values less than one second. These samples may be attributed to situations where a user abandons the Web-request before the browser has finished fetching all of its objects and where Web-requests are requested from multiple browsers. The distribution of **viewing time** is fitted well by a heavy-tailed Weibull distribution (see Figure 4.a), which is consistent with previous work [7][8][9].

In-line inter-arrival time is the time between the opening of one in-line object and the next. It measures the starting time between subsequent in-line objects. The inter-arrival time, up to the maximum permitted number of simultaneous objects, is just a few tens of milli seconds. Further in-lines are sent only after outstanding objects complete, so that inter-arrival times

may be as large as a few seconds or more. The distribution of **in-line inter-arrival time** matches a Gamma distribution.

Parsing time is the time spent parsing the HTML code in order to determine the layout of a page after fetching the main object. This quantity depends on the client machine. It is well matched by a Gamma distribution.

Both distributions of **main-object size** (see Figure 5.a) and **in-line-object size** are well fitted by a Lognormal distribution. The mean of **main-object size** is larger than **in-line-object size** and the variance of **in-line-object size** is greater than **main-object size**. The HTTP object size is easily obtained from the *content-length* field in the HTTP response header. The histogram of **main-object size** is shown in Figure 5.b.

Request size is the size of the HTTP request header. It is best fitted by a Lognormal distribution.

B. Web Caching Model

In HTTP, cacheable objects are stored with a tag containing the expiration-time; after this time an object is no longer valid [1][2]. The expiration-time is estimated from the *last-modified* field in the HTTP response header. The longer the time since the object was modified, the longer the expiration-time. Before HTTP sends a request, it checks the local cache. If the requested object is found in a local cache, HTTP checks the expiration-time. If the object is valid, HTTP reads the object from the cache instead of retrieving it from a remote server. If not, HTTP sends a conditional request with the expiration-time tag in the *if-modified-since* field. The server compares the tag with the time the object was modified. The server responds with the HTTP status code 304 if it is still valid. If not, the server downloads the object.

We model two types of events. One is when the expiration-timer of the cached object is expired and the server proves the validation of the cache. The other is when the cached object is confirmed as expired by the server or when the object is not cached; in both cases an object is downloaded from the server. We do not model the case where the object in the cache is valid because such cases do not generate network traffic.

We can easily determine a cached object from the HTTP status code in the HTTP trace. It is a cached object if the HTTP status code is 304. In a Web-request, some objects can be cached and some objects are not. The Web-request is defined to be a cached Web-request if more than a half of the objects in the Web-request are cached. The reasoning is that two Web-requests share an object with low probability. Further, if the first object in a Web-request is cached, all following objects tend to be cached. Figure 6 shows the histogram of the fraction of non-cached objects in a Web-request. Two peaks at zero and one indicate whole objects are either cached or non-cached. Based upon the definition of a cached Web-request, we measure **number of consecutive cached Web-requests** and **number of consecutive non-cached Web-requests** from the trace.

Our Web caching model is a *two-state renewal process* where a Web-request is an event. A renewal process remains in one state for a number of self-returning events and moves to next states at the time of a renewal. This model has two states; one state represents a Web-request that is locally cached (Cached state) and the other represents a normal Web-

request (Non-cached state). We have extracted the distributions of the number of self-returning events of the two states; **number of consecutive cached Web-requests** and **number of consecutive non-cached Web-requests**. The model remains in the Cached (or Non-cached) state until the number of Web-requests generated in this state is the same as the number of Web-requests obtained from the distribution. Then, the model switches to the Non-cached (or Cached) state with probability one. **Number of consecutive cached Web-requests** is best fitted by a Geometric distribution and **number of consecutive non-cached Web-requests** is best fitted by a Lognormal distribution.

C. TCP model

In TCP, segments are transmitted in a burst until the TCP window is closed. The interval between two successive bursts is called a stall. We use a TCP model in which the number of segments transmitted and the time spent in a stall are chosen from empirically derived distributions. Our TCP model is described in more detail elsewhere [15]

D. Correlation

Parameters in our model are assumed to be uncorrelated with themselves as well as with one another so that the generation of one parameter is independent of the generation of other parameters. In order to test the accuracy of this assumption we calculated auto and cross correlation functions of our parameters. We selected only client sessions which had more than 40 samples of Web-requests. We calculated the correlation function up to a lag of 20 from these samples. We find that, generally, correlations are very low. An exception is the significant amount of auto-correlation in **request size** (see Figure 7) where the average auto-correlation at lag one is 0.59. Note that **request size** is far smaller than either **main-object size** and **in-line-object size**. Consequently we would expect the auto-correlation of **request size** not to affect significantly the accuracy of the model. Table 4 shows the average auto-correlations at lag one.

Request Size	0.59	Number of In-line Object	0.18
Viewing Time	0.23	In-line Inter-arrival Time	0.05
Parsing Time	0.07	In-line Object Size	0.13
Main-object Size	0.16		

Table 4: Auto-correlations at lag one

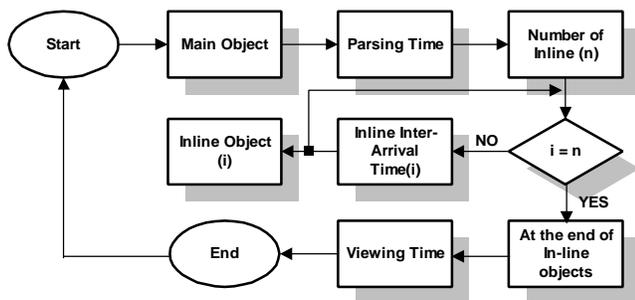


Figure 3: State transition diagram for Web traffic generation

VI. TRAFFIC GENERATION

The traffic model simulates an ON/OFF source. The state transition diagram shown in Figure 3 describes the traffic

generation process. At the beginning, the traffic corresponding to the main object is generated and is delayed for the period of **parsing time**. During this period a Web browser fetches the main object and parses **number of in-line objects** as well as the page layout. The model, however, generates **number of in-line objects** from the best-fit distribution and waits for the expiration of **parsing time**.

After the start of one in-line object there is a delay to the start of the next. The first in-line object starts after expiration of **parsing time**. The second in-line object does not wait until the first in-line object finishes but starts one **in-line inter-arrival time** after the start of the first. Subsequent in-line objects start until the number of in-lines started equals **number of in-line objects**. The number of outstanding connections in the model is not restricted to four or six which is the case of Netscape and an Internet Explorer. Instead we model the number of outstanding connections by the distributions that are collected from the trace. In the model, depending upon **in-line object size** and **in-line inter-arrival time**, the number of outstanding connections will vary. Frequently, **in-line inter-arrival time** is less than the duration of the connection, which is mainly determined by **in-line object size**. Hence, the model indirectly simulates the parallel downloading of in-line objects. After all objects are transmitted the model is silent for **viewing time**. After the expiration of **viewing time**, the model starts to generate a new Web-request.

The Web caching model influences the final model through **main object size** and **in-line-object size**. The main object is more frequently changed than the in-line object because modifying a Web page leads to modifying the main object. For example, famous news sites modify a text article and a few pictures associated with the article but leave a number of icons and commercial banners unchanged. Because the expiration-time of main objects is relatively short due to frequent changes, the main object is fetched most of time. The HTTP object size becomes zero except for the main object while in the cached state. While in the non-cached state, the sizes of both HTTP object types are generated from the distribution.

VII. VALIDATION OF THE MODEL

To validate the model we need measurements from the trace that are independent of any measurements used in constructing the model. Two such measurements are **on-time** and the variation of the required bandwidth in time.

A. On-time

On-time is not directly used in constructing our model. **On-time** is the function of a number of parameters; **number of in-line objects**, **in-line inter-arrival time**, **main-object size**, **in-line-object size** and **stall time**. These parameters interact with one another and combine to determine **on-time** in the model. Thus, it may be used to check the model by comparing with directly measured **on-time** from the trace. If there is any significant error in our measurement of parameters, the way we combine measurements or the completeness of our model we would expect differences between the model and the trace.

On-time from both the model and the trace match a Weibull distribution with the shape parameters 0.77 and 0.68,

respectively. The mean and standard deviation of the traces are 11.34 and 23.85 and those of the model are 10.49 and 20.33. The *cumulative density function* (CDF) comparison is shown in Figure 8.

B. Bandwidth Required

The variation of the required bandwidth with time is another independent check. To implement this test, we have recorded the sum of bytes in ten milli second granularity from the trace and the model. To obtain the value at the next level of the granularity we summed ten consecutive samples from the previous granularity and calculated the arithmetic mean of the samples. We measured four different granularities up to a ten second granularity and plotted 200 samples from both the trace and the model.

In this section, we are more interested in the largest granularity of 10 in order to compare the overall behavior of the trace and the model. The means of the required bandwidth of the model and the trace closely match as shown in Figure 9; the mean of the trace is 4656 kbytes and that of the model is 4699 kbytes.

C. Self-similarity

A process X is called self-similar if the aggregate process of X has the same auto-correlation function as X . The degree of the self-similarity is expressed using the Hurst parameter. The Hurst parameter is always less than 1.0 and greater than 0.5. The closer the Hurst parameter is to 1, the more self-similar the process. Further discussion about self-similarity can be found in [12][14].

We have tested the self-similarity of the model using two methods: variance-time plot and R/S (Rescaled Adjust) plot. We have also tested the burstiness of the traffic using four different time scales. For these three tests, we have used the same data described in Section B.

The two *variance-time plots* shown in Figures 10.a and 10.b exhibit the level of the self-similarity of the model and the trace. The *Hurst* parameters calculated are 0.805 for the trace and 0.78 for the model.

The *R/S plot* shows that the asymptotic slope is different from 0.5 and 1.0 (see Figures 11.a and 11.b). The estimated *Hurst* parameters are 0.8 for the trace and 0.77 for the model. The group of values near one on the x-axis fall below a slope less than half. That is mainly due to the small number of samples. The result of the rest test can be found in [15].

VIII. DISCUSSIONS

We collected data from Keep-alive and persistent connections. About 76 percent of connections in the trace are Keep-alive enabled connections. About 40 percent of these connections actually use Keep-alive and the rest are closed either after the expiration of the Keep-alive timer, because the user switched to another Web server, or because the server closed the connection due to the nature of a page. Most servers have a short period for the Keep-alive timer due to limited resources at the server.

IX. CONCLUSIONS

Web traffic has more structure to it than most types of Internet traffic. Further, the characteristics of the traffic change as browsers and servers evolve, as the behavior of users change, as the speed of the network increases and as the protocols change. This makes Web traffic modeling a challenge.

We have presented a model of Web traffic that attempts to capture the major aspects of the traffic. This has required us to extend the models used in previous work and to correct shortcomings of previous models. For example, the presence of a frame has required us to redefine the basic unit of the model.

From the extensive traces we have gathered of Web traffic we have extracted the distributions of the parameters used to represent the model. From the model, we have generated synthetic traffic and can now compare it with the trace data. We have shown that independent variables from the trace and from the model agree well. We believe that the model is accurate enough to predict the behavior of traffic as parameters change (such as the speed of a line or the size of Web-requests), however, this remains to be proved.

X. REFERENCES

- [1] T. Berners-Lee, R. Fielding, and H. Frustuk. *Hypertext Transfer Protocol – HTTP/1.0*. Internet RFC 1945, May 1996.
- [2] R. Fielding, J. Getty, J. Mogul, H. Frystyk, and T. Berners-Lee. *Hypertext Transfer Protocol – HTTP/1.1*. Internet RFC 2068, Jul 1997.
- [3] K. Thompson, G. J. Miller, and R. Wilder. Wide-Area Internet Traffic Patterns and Characteristics (Extended Version). *IEEE Network Magazine*, Nov 1997.
- [4] V. Jacobson, C. Leres, and S. McCanne. *Tcpdump manual and software*.
- [5] M. Ryan. *Tcpshow version 1.73* available at <http://http.cs.berkeley.edu/~daw/mike/>, Nov, 1996.
- [6] H. F. Neilson, *et al.* Network Performance Effects of HTTP/1.1, CSS1, and PNG. In *Proceedings of SIGCOMM '97*, Cannes, French Riviera, France, Sep 1997.
- [7] B. Mah. An empirical model of HTTP network traffic. In *Proceedings of INFOCOM '97*, Kobe, Japan, Apr 1997.
- [8] S. Deng. Empirical model of WWW document arrivals at access link. In *Proceedings of ICC '96*, Jun 1996.
- [9] M. Crovella and A. Bestavros. Self-Similarity in World Wide Web traffic, Evidence and Possible cause. In *Proceeding of ACM SIGMETRICS '96*, Philadelphia, PA, Apr, 1996.
- [10] P. Badford and M. Crovella. Generating representative Web workloads for network and server performance evaluation. In *Proceedings of ACM SIGMETRICS '98*.
- [11] R. B. D'Agostino and M. A. Stephens. *Goodness-of-Fit Techniques*, Marcel Dekker, Inc., 1986.
- [12] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Transaction on Networking*, pages 1-15, 1994.
- [13] C. A. Cunha, A. Bestavros, and M. Crovella. Characteristics of WWW client-based trace. *Technical Report TR-95-010*, Boston University, Dept. of Computer science, Apr 1995.
- [14] J. Beran. Statistics for long-memory processes. Monographs on Statistics and Applied Probability. Chapman and Hall, New York, NY, 1994.
- [15] Hyoung-Kee Choi and John O. Limb. A behavioral model of Web traffic (Extended Version) available at <http://users.ece.gatech.edu/~hkchoi/model.pdf>.

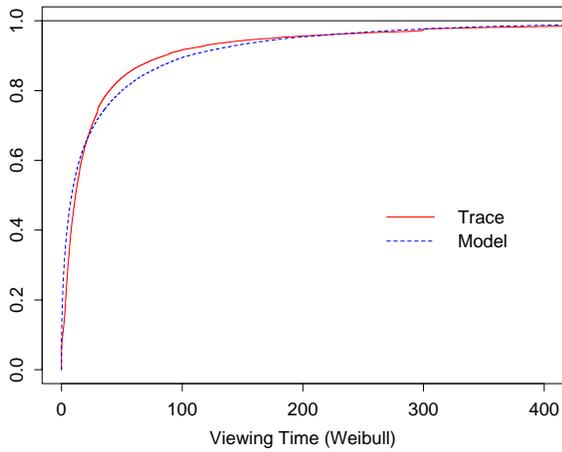


Figure 4.a: CDF comparison of viewing time with a Weibull distribution

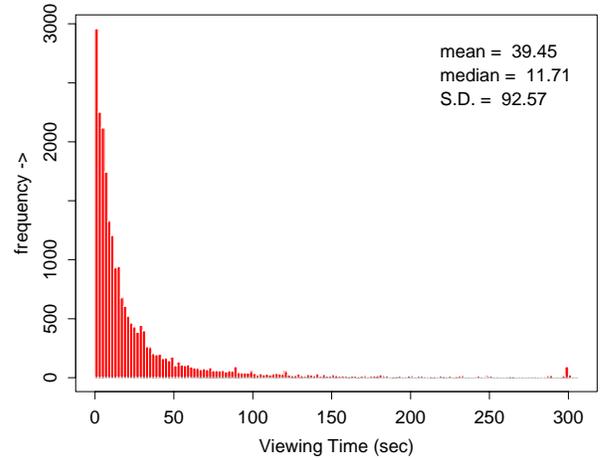


Figure 4.b: Histogram of viewing time. The peak near 300 sec are due to periodically re-freshed pages.

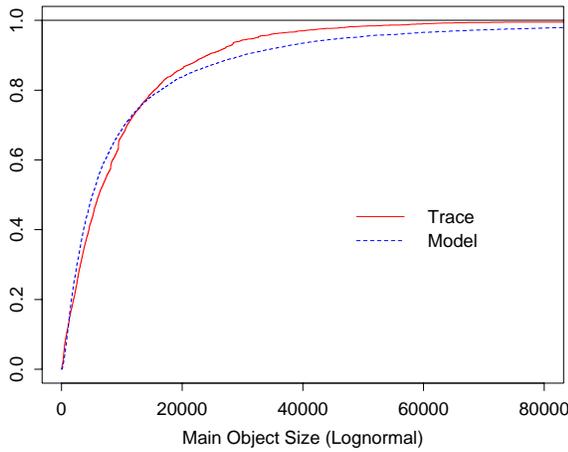


Figure 5.a: CDF comparison of main-object size with a Lognormal distribution

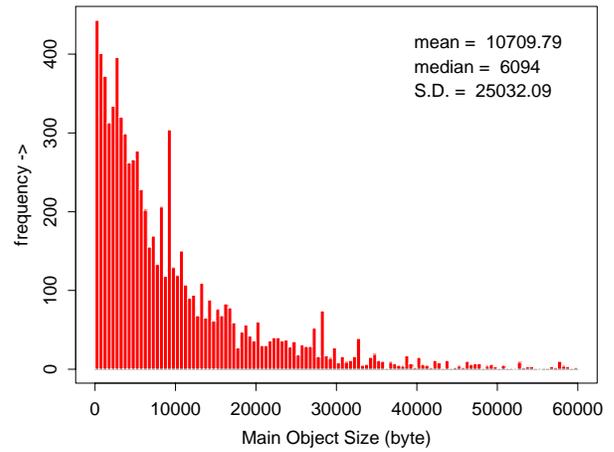


Figure 5.b: Histogram of main-object size

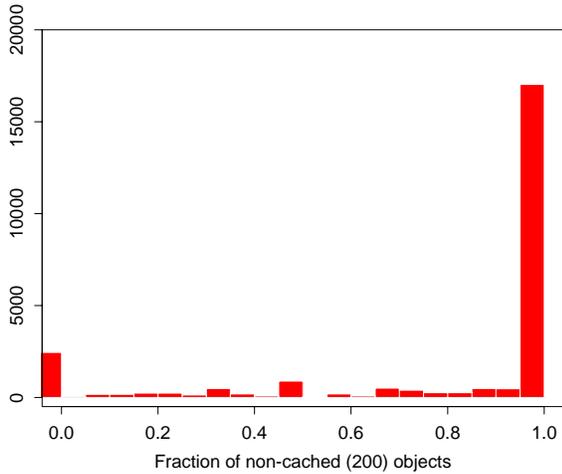


Figure 6: Histogram of the fraction of non-cached objects. A sample is calculated by dividing the number of non-cached objects by number of in-line objects in a Web-request.

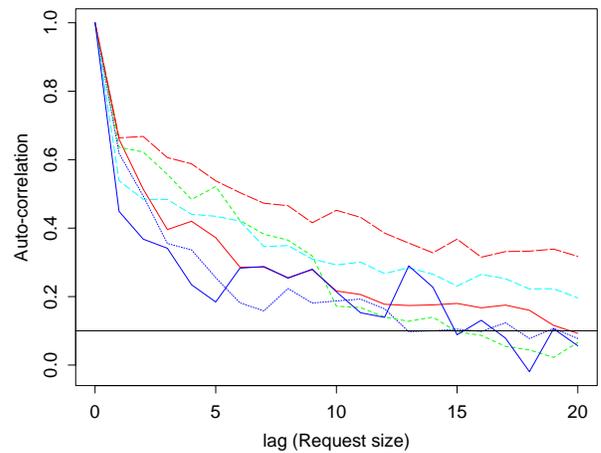


Figure 7: Auto-correlation function of request size. Six different samples are plotted. The mean of auto-correlation at lag one is 0.59.

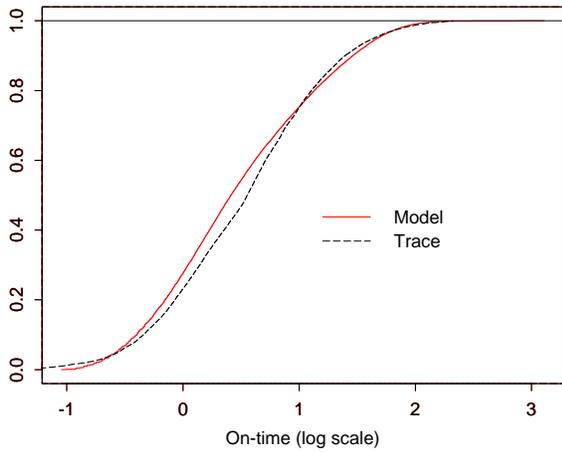


Figure 8: CDF comparison of On-times - On-time of Trace and On-time of Model. X-axis is log-scaled

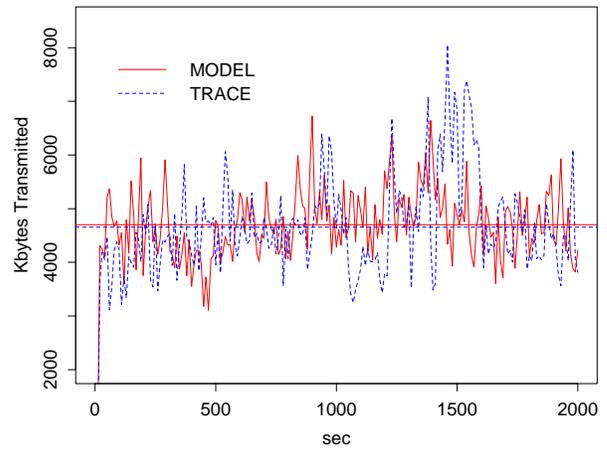


Figure 9: The variation of the demanded bandwidth in time. Two parallel lines indicate the mean of samples.

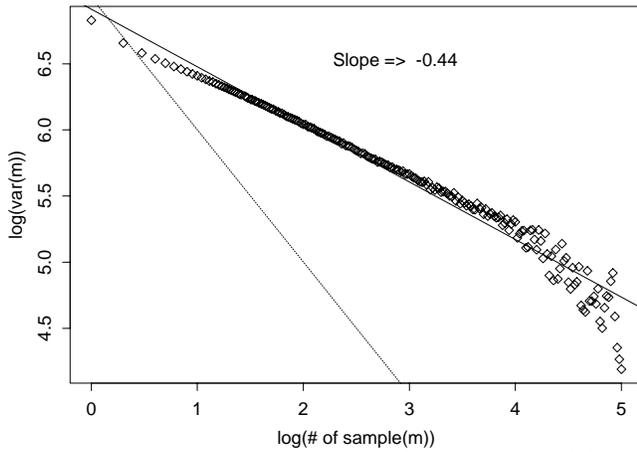


Figure 10.a: Variance-time plot of the model. The Hurst parameter is 0.78. The slope of lower line is -1.

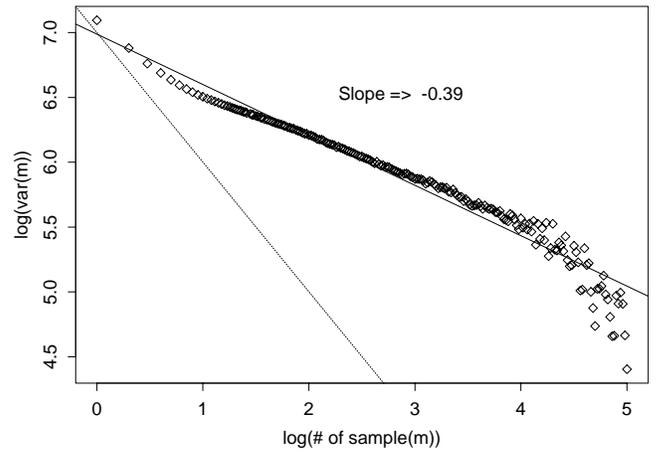


Figure 10.b: Variance-time plot of the trace. The Hurst parameter is 0.805.

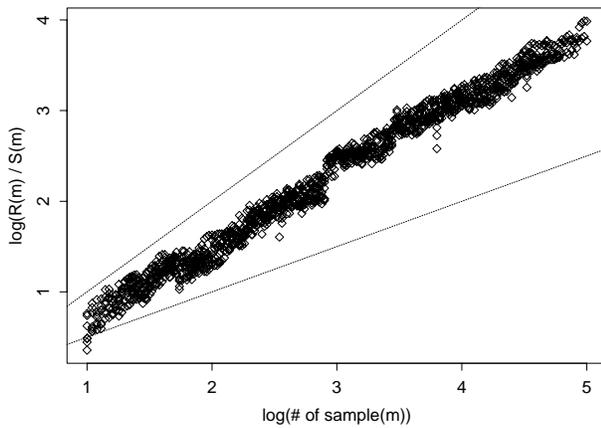


Figure 11.a: R/S plot of the model. The Hurst parameter is 0.77. The slopes of two straight lines are 1 and 0.5, respectively.

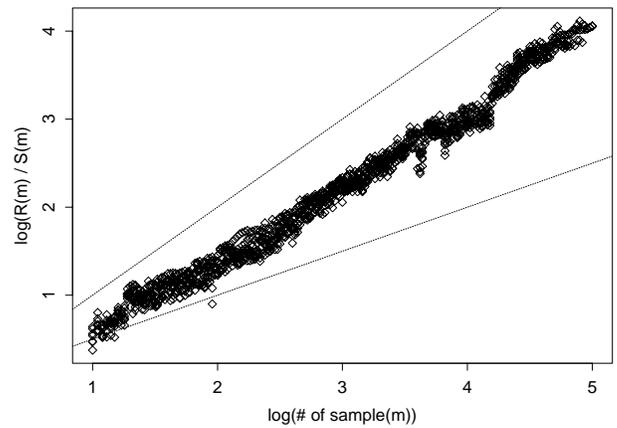


Figure 11.b R/S plot of the trace. The Hurst parameter is 0.8.