A SYNTHETIC WORKLOAD MODEL FOR INTERNET MOSAIC TRAFFIC

Martin F. Arlitt and Carey L. Williamson
Department of Computer Science
University of Saskatchewan
Saskatoon, SK, Canada S7N 5A9
e-mail: carey@cs.usask.ca

Keywords: Network, Model Design, Parameter Identification, Telecommunications

ABSTRACT

Mosaic traffic (i.e., World-Wide Web) is the fastest growing component of the aggregate packet and byte traffic on the NSFNET backbone. Modeling the workload characteristics of these mosaic sessions is therefore deemed important in any simulation study of the Internet or future high speed networks.

The ATM-TN TeleSim project has designed and implemented a synthetic workload model for Internet mosaic traffic, to be used as input to a parallel simulator for high speed ATM networks. This paper describes the workload characterization and modeling process for this synthetic workload model, as well as the design, implementation, and validation of this traffic model.

1 INTRODUCTION

To many people, "the Internet" is synonymous with "the Information Superhighway". The Internet provides access to a vast array of information services and information resources. Information services include traditional services such as electronic mail, file transfer, and remote login, as well as new services such as resource discovery (Bowman *et al.* 1993; Obraczka *et al.* 1993). Information resources include news groups, software distributions, technical reports, library catalogs, product information, weather maps, satellite images, radio broadcasts, musical releases, video sequences, and many more. Almost everything that network users need is "out there" somewhere on the Internet; the hard part is finding what you want, when you want it. For this purpose, good resource discovery and navigation tools are needed.

One of the most popular resource discovery tools for navigating the Internet is `mosaic` (Andreessen 1993). Developed by the National Center for Supercomputing Applications (NCSA), mosaic provides an easy-to-use graphical user interface allowing users to browse the information resources of the Internet (e.g., technical reports, software, news groups, phone lists, images, and weather maps), using a wide array of document types (e.g., text, image, graphics, audio, video, hypermedia).

In particular, the mosaic application allows users to navigate sites that are part of the World Wide Web (Berners-Lee *et al.* 1993). The World Wide Web (WWW) defines a global naming convention for all documents that are part of the Web, a standard format for defining documents in the Web, and a standard protocol for exchanging documents between Web sites. The mosaic application interfaces to the WWW simply by using these standard document formats and protocols, as well as other standard Internet protocols.

The growth in popularity of the World Wide Web has been phenomenal. WWW traffic is by far the fastest growing component of the aggregate packet and byte traffic on the current NSFNET backbone (NSFNET 1995). In the last two years, WWW traffic has risen from 186th (December 1992) on the NSFNET activity list, to 7th in June 1994, to 2nd in December 1994. WWW accounted for 3.3 Terabytes of data on the NSFNET backbone in December 1994 (16% of the total bytes transmitted), second only to ftp (6.5 Terabytes).

A large part of this growth has been due to mosaic. Mosaic makes it very easy for users to access a variety of document types on the Internet, perhaps without even realizing the network resources (i.e., bandwidth) consumed by this traffic. In short, mosaic traffic will soon account for a substantial fraction of Internet backbone traffic.

Understanding the workload characteristics of these mosaic sessions is therefore deemed important in any simulation study of the Internet or future high speed networks. The key steps in this understanding process are empirical network traffic measurement, workload characterization, and workload modeling.

The purpose of this paper is to describe the workload characterization and modeling of Internet mosaic traffic that has been done as part of the ATM-TN TeleSim project. The ATM-TN TeleSim project has designed and implemented a synthetic workload model for mosaic traffic, to be used as input to a parallel simulator for high speed ATM networks. The traffic model is designed to represent the workload characteristics of wide area network traffic generated by Internet mosaic users.

The remainder of this paper is organized as follows. Section 2 presents background material on the NCSA mosaic application, including its terminology and protocols. Section 3 presents empirical network traffic measurement results for the mosaic application in a campus local area network environment. Section 4 describes the design of a synthetic workload model for Internet mosaic traffic, as well as a brief description of its implementation in SimKit. Finally, Section 5 concludes the paper.

2 MOSAIC OVERVIEW

Mosaic provides an easy-to-use graphical interface that allows users to browse the information resources of the Internet. In particular, the mosaic application allows users to navigate sites that are part of the World Wide Web (Berners-Lee *et al.* 1993).

Navigation in mosaic is done by going from one WWW "page" to another. Each of these pages (also called "home pages" or "Web pages") is a hypermedia document that can contain a combination of text, graphics, images, audio, and video, as well as *links* to other pages. Users move from one page to the next using simple "point and click" on the links (which are usually highlighted blue text in mosaic). Each page has a globally unique name in the WWW name space, called a Uniform Resource Locator (URL). Navigation can be done one step at a time (hierarchically) through a series of pages using links, or can be done directly using a URL. An example of a URL is `http://alf.usask.ca/personal/Arlitt/home.html`.

WWW pages are constructed using a language called the HyperText Markup Language (HTML). Since each page consists of one or more files, HTML defines which files should appear on the page, what format they are in, where they should appear, and which items should be highlighted as links to other pages. Links may lead to another file on the same site, or to another file at a different site.

When specific WWW pages are requested by users, the pages are transmitted to other sites using a standard protocol called the HyperText Transfer Protocol (HTTP). The HTTP protocol is an application layer protocol that defines the rules for exchanging HTML documents between WWW sites. HTTP uses TCP as its transport layer protocol (Tanenbaum 1988). TCP is a connection-oriented transport protocol that provides reliable byte stream delivery between two hosts on the Internet. A single HTML document can result in multiple TCP connections, one for each file that appears as part of the WWW page. The transfer of one file must complete before the transfer of the next file begins.

Mosaic provides full TCP/IP based communications support, including HTTP, Gopher, FTP (File Transfer Protocol), NNTP (Network News Transfer Protocol), WAIS (Wide Area Information Server), Archie, finger, and whois (Obraczka *et al.* 1993). It is also possible to mail documents from mosaic using SMTP (Simple Mail Transfer Protocol), and to telnet to various servers (Andreessen 1993).

Network traffic measurement of mosaic in our environment shows that the vast majority (well over 90%) of the network packet traffic generated by mosaic uses the HTTP protocol. In network traffic traces, mosaic traffic can be readily identified as a subset of all TCP traffic since all HTTP traffic uses port 80 at the TCP layer.

3 WORKLOAD CHARACTERIZATION

In order to characterize the network packet traffic workload generated by Internet mosaic traffic, empirical network traffic measurements were collected at the University of Saskatchewan. Measurements were collected using techniques similar to those reported in the literature (Claffy 1994; Paxson 1994; Paxson and Floyd 1994). This section describes the network measurement and workload characterization process.

For all network traffic measurements, the trace collection facility was `tcpdump` (Jacobson *et al.* 1989). Tcpdump is a software network monitor that provides user-level control of trace collection, including filtering on a host, protocol, or port basis. The traces were set up to record only TCP packets with the SYN (start) or FIN (finish) bits set, and only the headers of these packets were recorded. These traces provide enough information to deduce the start time, end time, duration, and bytes transferred in each direction by each TCP connection, as well as the source host, destination host, and port number used.

Two different sets of traces were used for workload characterization. The first set of traces recorded *all* TCP traffic destined from the campus Ethernet to the Internet, or from the Internet to the campus Ethernet. The traces were collected for a four day period in April 1994.

This paper presents data only for WWW traffic (i.e., HTTP traffic that uses TCP port 80), representing 5.5% of the total byte and packet network traffic in the four day trace. An analysis of the complete TCP traces is presented in (Zhu 1994).

A total of 7070 WWW-related TCP connections were recorded in the first set of traces. These connections accounted for 97 Megabytes of data transferred on the network, for a mean of 14,380 bytes per TCP connection. While the total volume of mosaic-generated traffic is fairly low in this trace, the measurement results are quite consistent with those reported for the NSFNET at that time (NSFNET 1995). Thus we believe that this trace provides a reasonable snapshot of mosaic traffic for further workload characterization.

The second set of traces recorded only mosaic-generated activity for a single workstation on the campus Ethernet. Several traces were collected of mosaic sessions initiated from this workstation, ranging in duration from several minutes to several hours. A total of 1985 TCP connections were recorded in these traces. These connections accounted for 30 Megabytes of data transferred on the network, for a mean of 15,646 bytes per TCP connection.

The two sets of traces were analyzed both separately and together. Only minor differences[1] were observed between the two traces, and thus only the combined analyses are presented in this paper. The combined traces have 9055 TCP connections with a mean size of 14,659 bytes. The largest number of data bytes transferred by a TCP connection was 2.2 Megabytes. Over 95% of the TCP connections transferred fewer than 42,000 bytes.

Figure 1 shows the distribution of the number of bytes per TCP connection for *requests* for WWW pages. The distribution is dominated by a peak at about 1130 bytes. Further analysis of network traces shows that request packets always contain the `GET` command, the name (URL) of the file or page begin requested, and a list of the document formats that are understood (i.e., acceptable) by the requesting machine. Figure 1 is thus showing that mosaic requests have a fixed overhead of approximately 1100 bytes for specifying document formats, and a small variable overhead due to the variable length of URL names.

Figure 2 shows the distribution of the number of bytes per TCP connection for the *responses* to requests for WWW pages. Note that the horizontal axis is logarithmic, showing the log (base 2) of the number of bytes in each file transferred by mosaic. Figure 2 shows that most of the files that appear on WWW pages are small. Approximately 50% of the files are in the 1-4 Kilobyte range, and over 90% of the files are less than 16 Kilobytes in size. Note, however, that there is a significant tail to the distribution, representing larger file sizes.

[1] For example, the traces of all campus TCP traffic include some WWW traffic generated by browsers other than mosaic, which may not support the full range of document types that mosaic does. Thus the mean request size and the mean file size are slightly lower in these traces than for the mosaic-only traces.

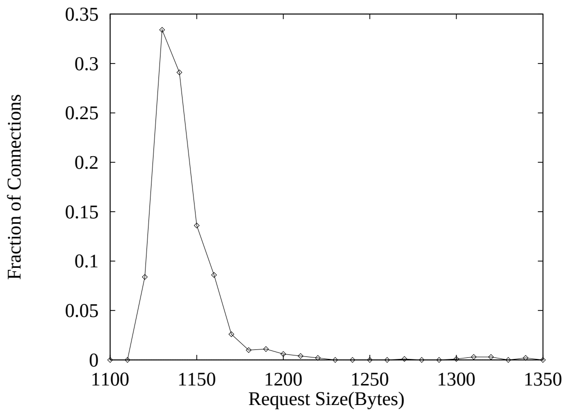

Figure 1: Distribution of Bytes per TCP Connection for WWW Requests

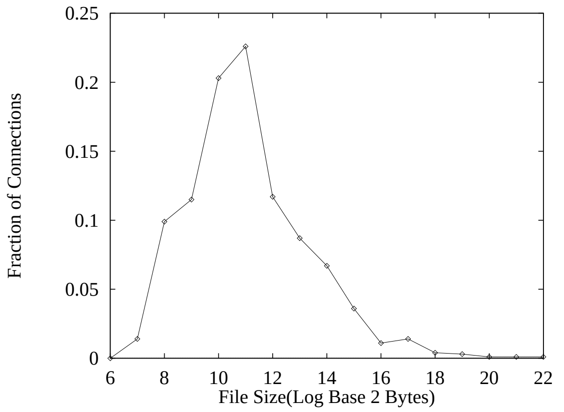

Figure 2: Distribution of Bytes per TCP Connection for WWW Responses

Figure 3 shows the distribution of the time duration of each TCP connection. The duration of a TCP connection is clearly correlated with the number of bytes sent by the connection. However, several other factors affect the connection duration as well, such as the geographical location of the destination host to which the connection is made, and the traffic load on the portion of the Internet traversed by the connection.

Figure 4 shows the interarrival times between connections (i.e., the time from the start of one connection to the start of the next connection). These are obviously related to the connection durations shown in Figure 3.

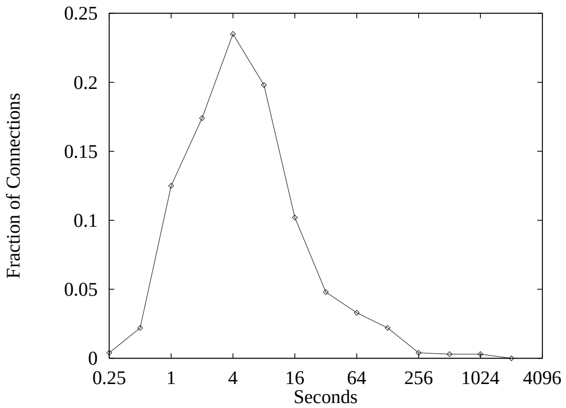

Figure 3: Distribution of TCP Connection Duration for WWW Traffic

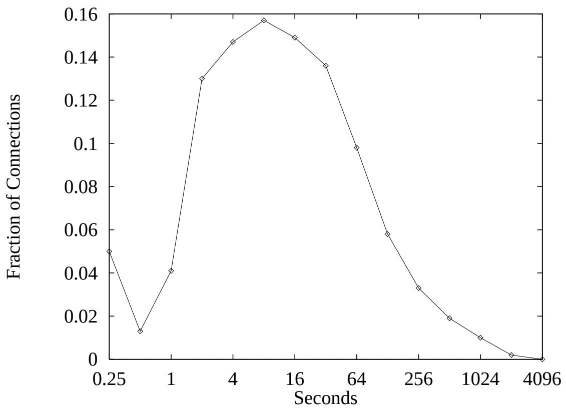

Figure 4: Distribution of TCP Connection Interarrival Time for WWW Traffic

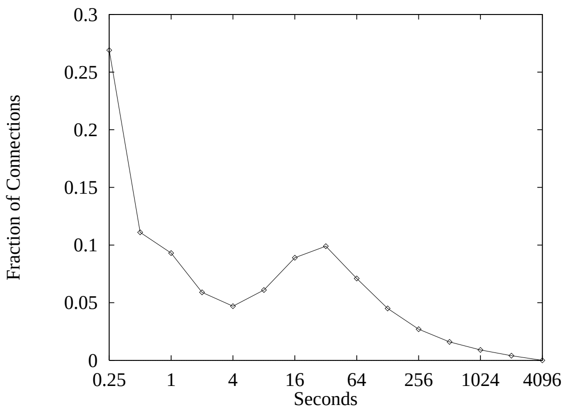

Figure 5: Distribution of TCP Connection Gap Time for WWW Traffic

However, the connection interarrival time also depends on the "think time" between connections. This think time could be machine generated (e.g., to access the next file that is part of the current Web page, as soon as the previous file is completely received) or human generated (e.g., the user reads the current Web page and then decides which link to traverse next).

Figure 5 shows more clearly the connection gap times (i.e., the time between the end of one connection and the start of the next one). There are many short (machine-generated) connection gap times when there are multiple files for one Web page, and there are many longer (human-generated) connection gap times when new Web pages are selected. The dividing point between the two regions appears to be around 4 seconds. Further analysis of the latter (human-generated) region has found these connection gap times to be exponentially distributed and independent.

The above data has been extremely useful in constructing our synthetic workload model for Internet mosaic traffic.

4 WORKLOAD MODELING

4.1 Model Overview

The Internet mosaic traffic model is designed to represent the wide area network traffic workload characteristics generated by Internet mosaic users.

The key part of the mosaic traffic model is modeling a mosaic *session* (i.e., a single human user seated at a single workstation, using the mosaic graphical user interface to explore and search for information all over the Internet). A single mosaic session can generate one or more *conversations* with different hosts on the Internet, where a conversation consists of a bidirectional exchange of packets with a single destination host. A conversation, in turn, may consist of one or more TCP *connections*, where each connection involves a bidirectional exchange of packets with a specific host. Each connection transfers the data bytes of an information unit, such as a text file, an image, graphics, etc. There may be several successive connections to the same destination host as part of a conversation. The next conversation may use the same destination host again, or a different one.

A typical mosaic session may last minutes or hours. Within the mosaic session, the model generates the number of conversations, the arrival time for each conversation, the number of connections for each conversation, the arrival time for each connection, the destination for the connection, and the number of bytes to be transferred in each direction for the connection.

4.2 Model Parameters

The mosaic session model has six parameters: (1) number of conversations per session; (2) conversation arrival model; (3) conversation destination model; (4) number of connections per conversation; (5) connection arrival model; and (6) number of bytes exchanged per connection.

These parameters are described in more detail as follows:

- **Number of conversations per session**.

  This parameter specifies the average number of conversations initiated during one mosaic session. Because of the wide variability in the mosaic sessions of individual users, it is extremely difficult to determine a proper setting for this parameter. For simplicity, our model assumes that the number of conversations is geometrically distributed, with a mean of 50 conversations per session.

- **Conversation arrivals**.

  Each conversation corresponds to a complete Web page. Once a conversation is complete, the arrival time of the next conversation must be determined. The next conversation may follow immediately after the previous conversation if the user is selecting successive links in a list of items, or is traversing links that have been previously explored. The next conversation may follow after a pause if the user takes some time to read the current document before selecting the next link, or if the user leaves the session idle, to return to it later.

  Empirical measurements have found both the conversation interarrival times (i.e., the time between the start of one conversation and the start of the next conversation) and the conversation gap times (i.e., the time between the end of one conversation and the start of the next conversation) to be exponentially distributed and independent. As a result, a simple Poisson arrival model is used for conversations. The mean conversation gap time in our model is 33 seconds.

- **Conversation destination**.

  When a conversation is to be started, the destination for the conversation must be determined. Several factors must be considered when making the choice of the next destination.

  First, when a conversation is established with a particular destination, there is a strong possibility that there will be more conversations to the same destination within the near future. This "locality of reference" phenomenon, called temporal locality or persistence in the literature, is common in many areas of computer systems, and can be modeled quite easily with Least-Recently-Used (LRU) stack methods. Analysis of the empirical trace data shows strong persistence in conversations, but with only the topmost stack element seeing significant reuse. Thus persistence is modeled with a single real-valued parameter DEST_SAME_PROB that specifies the probability of a conversation going to the same destination site as the last conversation. This parameter is set to 0.36.

  A second factor in determining the choice of the next destination is the distribution of network traffic workload across the Internet. Since mosaic makes it easy to connect to sites across the Internet, and individual users have their own preferences in what sites they visit, it seems reasonable to expect that mosaic will generate traffic widespread around the world. However, many studies of network traffic have identified the non-uniform distribution (i.e., concentration) of traffic on most networks. For example, since many users begin their mosaic sessions with the default NCSA home page, there will be many references to common hyperlinks. As another example, certain ftp sites on the Internet tend to be more popular than others. Standard techniques for modeling spatial and temporal locality have been used to incorporate these characteristics into the mosaic model.

- **Number of connections per conversation**.

  This parameter specifies the average number of TCP connections per conversation. In our model, the number of connections per conversation is geometrically distributed with a mean of 2.5.

- **Connection arrivals**.

  Each connection represents the transfer of one file, which may be part of a Web page (conversation). Within each conversation, the arrival time of each connection must be determined. The next connection may follow immediately after the previous connection (e.g., the graphics to appear on a Web page are transferred immediately after the transfer of the text of the page), or may follow after a pause (e.g., when the user selects a link to another page on the same server). Only the former are modeled with connection arrivals. The latter are modeled as conversation arrivals.

As with the conversation arrivals, empirical measurements have found that connection gap times (i.e., the time between the end of one connection and the start of the next connection) are exponentially distributed and independent. Thus a simple Poisson arrival model is used for connection arrivals. The mean connection gap time in our model is 0.5 seconds.

- **Bytes exchanged per connection**.

  Finally, the mosaic traffic model needs to select the amount of information to be exchanged in each direction. Empirical measurements show that traffic is bidirectional. The source sends relatively few bytes (typically 1136 bytes plus or minus a few bytes), and the destination sends relatively more bytes (anywhere between a few hundred bytes and a few megabytes). Source bytes are modeled using a Normal distribution. Destination bytes are modeled using an Erlang distribution.

4.3 Model Implementation

The mosaic model has been implemented in SimKit, a C++-based simulation interface to Shared Memory TimeWarp (SMTW), which provides an environment for efficient parallel simulation. At this point, only the sequential implementation of the model is complete. The model (1678 lines of SimKit code) uses 107 logical processes (LPs): 1 for the source, 100 for destinations, and 6 to model the segmentation and reassembly between mosaic pages, TCP/IP packets, and ATM cells. End-to-end retransmission of lost packets is also supported.

4.4 Model Validation

The model has been parameterized to generate synthetic workloads that closely match the empirical data collected at the University of Saskatchewan. Close agreement (often within 2%, and nearly always within 10%) has been obtained on all key characteristics of the model. Comparisons with mosaic workloads generated at other sites (e.g., Calgary, Waterloo) remains to be done. Further testing and validation of the model continues.

5 CONCLUSIONS

Since mosaic traffic (i.e., World-Wide Web) is the fastest growing component of the aggregate packet and byte traffic on the NSFNET backbone, the ATM-TN TeleSim project has designed and implemented a synthetic workload model for Internet mosaic traffic, to be used as input to a parallel simulator for high speed ATM networks. The model has six parameters, and has been parameterized to match mosaic workloads measured at the University of Saskatchewan. Modeling the workload characteristics of these mosaic sessions is deemed important in any simulation study of the Internet or future high speed networks.

REFERENCES

Andreessen, M. 1993. "NCSA Mosaic Technical Summary". National Center for Supercomputing Applications.

Berners-Lee, T.; R. Cailliau; A. Luotonen; H. Nielsen; and A. Secret. 1993. "The World Wide Web". *Communications of the ACM*, Vol. 37, No. 8 (Aug.), pp. 76-82.

Bowman, C.; P. Danzig; U. Manber; and M. Schwartz. 1993. "Scalable Internet Resource Discovery: Research Problems and Approaches". *Communications of the ACM*, Vol. 37, No. 8 (Aug.), pp. 98-107.

Claffy, K. 1994. "Internet Traffic Characterization". Ph.D. Thesis, Department of Computer Science, University of California, San Diego.

Jacobson, V.; C. Leres; and S. McCanne. 1989. `tcpdump` UNIX manual page. Lawrence Berkeley Laboratories, University of California, Berkeley (June).

NSFNET. 1995. Data available by anonymous ftp from nis.nsf.net:/nsfnet/statistics.

Obraczka, K.; P. Danzig; and S. Li. 1993. "Internet Resource Discovery Services". *IEEE Computer*, Vol. 26, No. 9 (Sept.), pp. 8-22.

Paxson, V. 1994. "Empirically Derived Analytic Models of Wide-Area TCP Connections". *IEEE/ACM Transactions on Networking*, Vol. 2, No. 4 (Aug.), pp. 316-336.

Paxson, V. and S. Floyd. 1994. "Wide-Area Traffic: The Failure of Poisson Modeling". In *Proceedings of the 1994 ACM SIGCOMM Conference* (London, UK, Aug. 31-Sept. 2), pp. 257-268.

Tanenbaum, A. 1988. *Computer Networks*, Second Edition. Prentice Hall, Englewood Cliffs, N.J.

Zhu, W. 1994. "Characterizing Wide Area Conversations on the Internet". M.Sc. Thesis, Department of Computational Science, University of Saskatchewan (July).

Martin Arlitt is an M.Sc. student in the Department of Computer Science at the University of Saskatchewan.

Carey Williamson is an Associate Professor in the Department of Computer Science at the University of Saskatchewan. His interests include high speed networks, multimedia networking, and ATM.