Characterizing Wide Area Conversations on the Internet

A Thesis Submitted to the College of

Graduate Studies and Research

in Partial Fulfillment of the Requirements

for the Degree of

Master of Science

in the

Department of Computational Science

University of Saskatchewan

Saskatoon, Saskatchewan

by

Wanning Mandy Zhu

July 1994

PERMISSION TO USE

Characterizing Wide Area Conversations on the Internet

Candidate: Wanning Mandy Zhu

Supervisor: Rick Bunt

Master of Science Thesis

Presented to the College of Graduate Studies and Research

University of Saskatchewan

July 1994

ABSTRACT

Physically, the *Internet* is a collection of packet switching networks interconnected by gateways along with the TCP/IP (Transmission Control Protocol/Internet Protocol) protocol suite that allows them to function logically as a single large virtual network. Logically, the Internet is a vast library of information and services being exchanged among a worldwide user community. It provides electronic mail, file transfer capabilities, remote login, network news, and other services on top of both reliable, full duplex stream delivery and unreliable, connectionless packet delivery services.

Internet traffic measurement allows a statistical analysis of the international information flow. Such analysis is the first step to provide a snapshot of real workloads on wide area networks.

This thesis presents results from measurements of wide area network TCP conversations between the Campus Ethernet at the University of Saskatchewan and the Internet in April 1994. A total of 103,016 TCP conversations were recorded over a four-day period. The SMTP, Gopher, FTP-Data, WWW, Domain, Telnet, Finger, FTP-Control, and NNTP protocols accounted for the majority of the conversations and the majority of the bytes transferred. The interconnectivity between the Campus Ethernet and the rest of the world was quite rich with 53 countries being reached over the four-day period. Characteristics of these conversations were studied in terms of number of bytes transferred, conversation duration, bandwidth, interarrival time, and conversation setup latency. These results are reported in the thesis. Also reported is an analysis of conversations involving the Computational Science subnet.

It is hoped that this measurement and analysis provide a better understanding of usage of the Internet by people at the University of Saskatchewan and will contribute to an accumulating body of knowledge on Internet usage more generally.

ACKNOWLEDGEMENTS

Special thanks is given to my supervisor, Dr. Richard Bunt, for suggesting this interesting research topic and giving excellent guidance throughout my research work. Without his support I may never have got all this together. I would also like to thank other members of my committee, Dr. Carey Williamson and Dr. Winfried Grassmann, and my external examiner, Dr. Marshall Gilliland, for their helpful suggestions.

Many thanks to Dr. Dave Emigh for his help in obtaining traces in the Department of Computing Services and in providing information on the Campus Ethernet.

I would like to thank the faculty, staff, and the fellow graduate students in the Department of Computational Science, especially Yuru, Kerhong, and Nan, who have kept me company and made my life in Saskatoon more joyful. A big thanks to all the netfriends who have pointed me to the resources on the Internet.

I dedicate this work to my parents, YueJuan Qin and HueGuan Zhu. The values bestowed in me by these two wonderful people have made it possible for me to see this day. A special thanks goes to my sister and brother-in-law, Michelle and Chuck Yeung, for your financial support and encouragement. I am so lucky to have your love and care.

Most importantly, I thank my dear friend, Evon YuanFang Liu (although I prefer your nickname Ar Bing :-), for your continual help and your patience. Whenever my morale needed a boost, you could be counted on to do a good job. We share so many good old memories together. Your appreciation has guided me to be whom I am.

All of these efforts are very much appreciated.

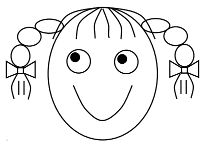

Contents

List of Figures

List of Tables

List of Protocols and Glossary

Chapter 1

Introduction

A Wide Area Network (WAN) is any physical network that spans large geographic distances [Comer 91], with the Internet being the most successful example of a WAN. The Internet is a collection of networks connecting through gateways using the TCP/IP protocol suite that function as a single cooperative virtual network. It forms a three-level hierarchy consisting of a cross-country backbone (CA*net for Canada, NSFnet for the United States, etc.), a set of "mid-level" or "regional" networks such as Sask#net (Saskatchewan area) and ONet (Ontario area) that each span a small geographic area, and a set of "campus" or "access" networks at a number of universities and research institutions and a number of military networks. Mid-level networks attach to the backbone and campus networks attach to the mid-level networks. All country backbones are connected to the NSFnet backbone. The Internet provides universal connectivity and three levels of network services: unreliable, connectionless packet delivery; reliable, full duplex stream delivery; and application-level services such as file transfer, remote login, electronic mail, and network news that are built on top of the previous two. The Internet reaches many universities, government research labs, corporate sites, and military installations throughout the world.

Recent growth of the Internet has been dramatic. As of January 1994, 2,217,000 hosts and 30,000 domains throughout the world were registered in the Internet domain system as compared to 2,056,000 Internet hosts and 28,000 Internet domains in October 1993 [Merit 94]. There were 81 countries having access to the Internet or general electronic mail services as of June 1994. The Internet grows not only in size, but also in services. Merit Network Information Center Services [Merit 94] also presents a history of usage by service in terms of percentage of packets.

With the computer networking infrastructure expanding rapidly, as exemplified by the growth of the Internet user base, robust and predictable networking environments to support these networking clients are now in demand. The proper design of future international networks, such as choosing appropriate routing and queueing strategies, and designing protocol interfaces, requires detailed performance studies, a key part of which is an accurate workload model. Measurement of network traffic is the first step providing an accurate description of real workloads and is the way as well to collect information that would be helpful in evaluating future network designs, capacity planning, and traffic predictions. Periodic usage updates are useful and required by users.

A number of recent studies have characterized IP, TCP, and application traffic on the Internet. Measurement studies of IP packet statistics have documented the growth in traffic volume [Paxson 93a, Claffy 93, Rutkowski 94] and the growth of Internet hosts and domains [Lottor 92, Lottor 94, Merit 94, Lynch 93]. Some studies focused on end-to-end behaviour on the Internet [Mills 83, Mukherjee 92, Sanghi 93, Claffy 93], while others focused on the study of wide area traffic [Heimlich 90, Asaba 92, Wakeman 93, Paxson 91, Caceres 91, Schmidt 92, Danzig 92b, Maffeis 93].

1.1 Thesis Objectives

Since the Internet continues to grow rapidly in size as well as in services, users and service providers at the University of Saskatchewan might ask questions like:

- How do people at the University of Saskatchewan use the Internet? What kinds of services does the local area network on campus provide to its users?
- How much wide area TCP traffic is distributed to the Campus Ethernet by each department subnet?

- How many countries does this local area network reach?
- How many networks on the Internet does this local area network communicate with?
- Which hosts at the University of Saskatchewan use the Internet the most frequently? Which networks on the Internet are used the most frequently by hosts at the University of Saskatchewan?
- How fast can wide area conversations be established?
- What are the characteristics of TCP conversations in terms of data transferred, duration, bandwidth, and interarrival time? How many simultaneous conversations are there in a given time period?
- Is traffic flow unidirectional or bidirectional?
- And, finally, how can these data be represented in a meaningful way?

All these questions have motivated this research. The objective of this research was to characterize wide area conversations (as defined in Section 2.8) based upon traces collected at the University of Saskatchewan.

The methodology used was as follows. TCP and IP packets were collected using a software tool called *TCPDUMP* [Jacobson 89]. Since TCP packets constitute a large percentage of the wide area network traffic, TCP conversations formed the basis of this research.

The purpose in presenting the measurement results is to share this information with the Internet community in the hope that it will result in better understanding of Internet behaviour.

1.2 Thesis Outline

The remainder of this document is organized as follows. Chapter 2 provides background information on the Internet. Chapter 3 surveys some studies related to this research. Chapter 4 gives experimental methodology, including the measurement environment and tool used, data collection and data reduction. Chapter 5 presents the measurement results in the following four categories: IP traffic breakdown by transport level protocols; TCP traffic breakdown by department subnets, protocols, and domains; TCP conversation setup latency; and per-conversation characteristics of TCP conversations, such as conversation size, duration, bandwidth, interarrival time, number of simultaneous conversations, and direction of data flow. The conclusions of the thesis are presented in Chapter 6.

Chapter 2

Background

2.1 Layering Model

Computer networks are designed in a highly structured way. To reduce design complexity, most networks are organized as a series of *layers* or *levels* such that each layer builds upon (i.e., uses functions and/or services provided by) its predecessor. The rules and conventions used in each layer are known as the layer *protocol*. Between each pair of adjacent layers there is an *interface* that defines primitive operations and services that the lower layer offers to the upper one.

The International Standards Organization (ISO) has proposed a standard reference model, the ISO *Open Systems Interconnection* (OSI) model, which is shown in Figure 2.1 [Zimmermann 80]. There are seven layers in this system. The physical layer is concerned with transmitting raw data bits over a communication channel. The data link layer takes a raw transmission facility and transforms it into a line that appears to the network layer to be free of transmission errors. The network layer is concerned with sending packets from a specific source node to a specific destination node. Routing concerns are usually dealt with at this layer. The reliability of the communication is usually left to the transport layer, although some transport services will send isolated packets with no guarantee of delivery. The session layer allows users on different machines to establish sessions. The presentation layer deals with the representation of the user's data, such as the syntax and semantics of the information transmitted. The definition of the application layer is left to the individual users.

7	application layer
6	presentation layer
5	session layer
4	transport layer
3	network layer
2	data link layer
1	physical layer

Figure 2.1: ISO Open Systems Interconnection Model

Layer protocols are usually grouped into sets corresponding to specific layering models, and those sets are called *protocol suites*.

Internetworking extends the system-to-system delivery service across a set of interconnected subnetworks that may have different properties. Several approaches have been used in internetworking, such as repeaters, bridges, routers, and gateways. A *repeater* just forwards bits from one network to another at the physical layer, making the two networks look logically like one network. This approach has minimal processing burden and no protocol conversion duties. A *bridge* can be used to connect two networks at the data link layer. This approach can copy data frames selectively and make necessary changes while doing so. A *router* implements internetworking functions at the network layer. This approach must deal with a wide range of procedures and protocols. A *gateway* is used to connect different subnetworks that have different transport protocols. This approach adds delay to the relay operation since it needs to perform protocol conversions through the transport layer if necessary [Tanenbaum 89].

Each of these approaches assumes that end-systems, or user machines, share common layers above the ones that handle the internetworking operation.

2.2 The Internet Layering Model

Physically, the *Internet* is a collection of packet switching networks interconnected by gateways along with the TCP/IP (Transmission Control Protocol/Internet Protocol) protocol suite that allow them to function logically as a single large virtual network.

The Internet started with the *ARPAnet*, which was created by the Defense Advanced Research Projects Agency (DARPA) of the United States Department of Defense (DoD). The Internet architecture model [Cerf 83] proposed by the DoD differs from the ISO model at the higher layers as shown in Figure 2.2.

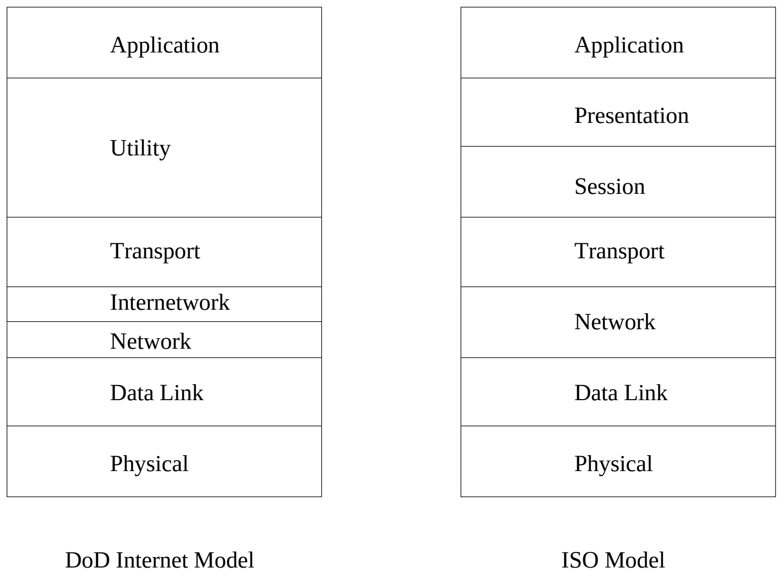

Figure 2.2: DoD Internet and ISO Models

The Internet model adds an internetwork layer between the network layer and the transport layer. At the internetwork layer, all network services, such as global Internet addressing, Internet routing, and error handling, are unified and viewed by hosts as an Internet datagram service. A special "type of service" field in each Internet datagram (defined in Section 2.4) can be used to select appropriate lower level network services. The principal protocols at this level are the *Internet Protocol* (IP) and the *Internet Control Message Protocol* (ICMP), which are used to coordinate host/Internet interactions including routing advice from gateways to hosts and warning messages related to congestion or unrecoverable failures.

There are two primary protocols at the transport layer. They are the *Transmission Control Protocol* (TCP) and the *User Datagram Protocol* (UDP).

The session and presentation layers are combined into a single utility layer in the Internet model. Protocols become much more application-specific at this layer, and the popular ones include the File Transfer Protocol (FTP), a remote terminal protocol Telnet, the Simple Mail Transfer Protocol (SMTP), the Domain Name Server Protocol (Domain), and the information service protocols (Gopher, World-Wide Web).

Actual programs that use the lower level protocols to accomplish specific applications are found in the application layer.

Figure 2.3 shows part of the DoD Internet protocol hierarchy. The details of some of the protocols in the TCP/IP protocol suite will be introduced in the following sections.

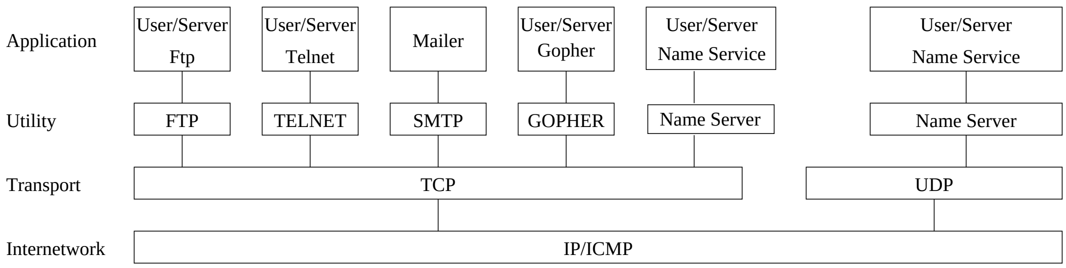

Figure 2.3: Part of the DoD Internet Protocol Hierarchy ([Cerf 83])

As the Internet spans every continent across the globe, the *CA*net*, the *NSFnet*, and thousands of other networks have become part of it. NSFnet, the National Science Foundation Network, is a high speed, centrally managed national (US) backbone network. It reaches out to Canada, Mexico, Europe, and the Pacific Rim. The CA*net backbone network connects to the NSFnet which, in turn, connects to the backbone networks of other countries. There are other major wide area networks that are not based on the TCP/IP protocol suite, such as *BITnet* and *UUCP*. However, it is possible to communicate between them and the Internet via electronic mail because *mail gateways* can act as "translators" between the different network protocols by connecting to two or more electronic mail systems and transferring mail messages among them.

The architecture of the wide area internetwork being considered throughout this study is shown in Figure 2.4.

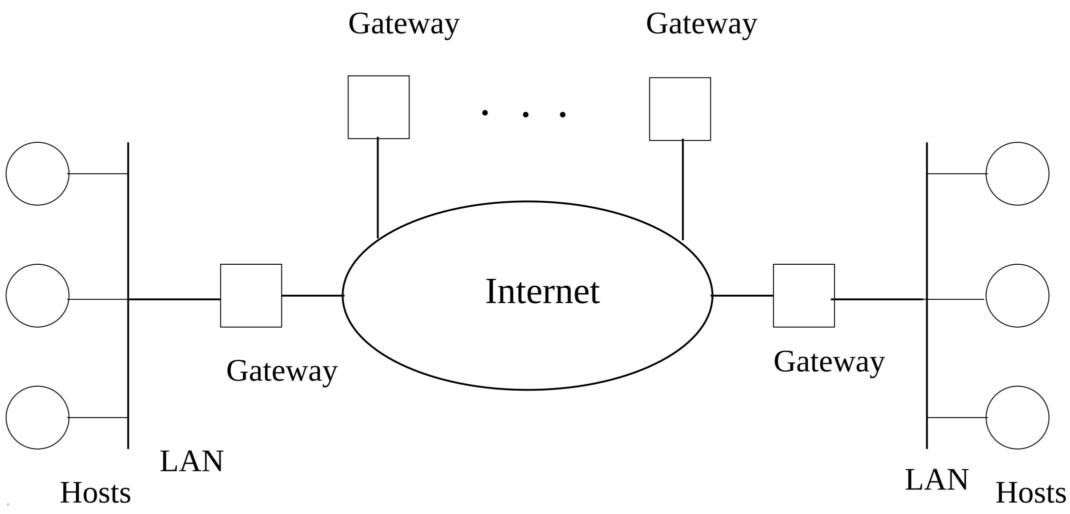

Figure 2.4: A Wide Area Internetwork

Host computers are connected to a local area network (*LAN*). Through *Internet gateways*, local area networks are connected to the wide area portion of the Internet. Internet gateways communicate with other Internet gateways through the Internet. The Internet is itself composed of a collection of switching nodes and communication links (not shown in the figure). In general, data flows from a source host, over a LAN, through an input Internet gateway, over the Internet, through an output Internet gateway, over other LANs and gateways, and finally to a destination host.

Most of the information written about the Internet, including its architecture, protocols, and history, can be found in a series of reports known as *Request For Comments* (RFCs). The Network Information Center (NIC) which is funded by the Defense Communication Agency (DCA) handles many administrative details for the Internet in addition to distributing documentation. There are several ways to obtain copies of RFCs. One can use anonymous FTP (Section 2.7.2) to server *nic.ddn.mil*, or send e-mail to *info-server@sh.cs.net* or *service@nic.ddn.mil* to get an electronic copy of an RFC, or one can obtain a paper copy of an RFC by calling 1-800-235-3155 (within US).

2.3 Internet Addressing

Each host on the Internet has a unique address called its *Internet address* or *IP address*. To ensure that the network portion of an IP address is unique, all Internet addresses are assigned by the central authority, the NIC. Each host on the Internet is assigned a unique 32-bit Internet address that is used in all communication with that host. Conceptually, each address is a pair (*netid, hostid*), where *netid* identifies a network, and *hostid* identifies a host on that network. In practice, each IP address must have one of the first three forms shown in Figure 2.5. The three primary forms, Classes A, B, and C, can be distinguished by the first two bits. Class A addresses are used for the handful of networks that have more than 2^{16} hosts. Class B addresses are used for intermediate size networks that have between 2^8 and 2^{16} hosts. Class C networks have fewer than 2^8 hosts.

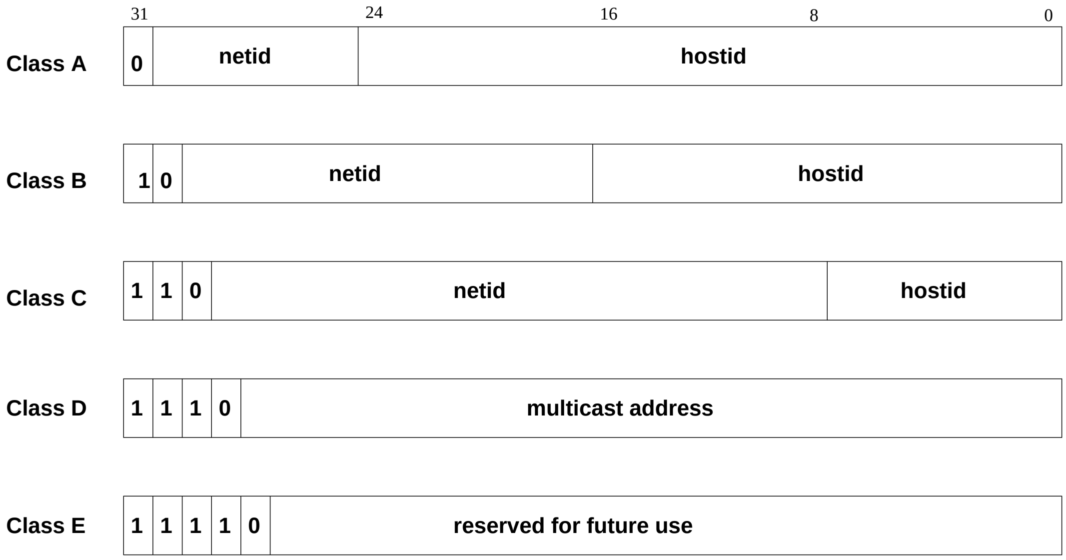

Figure 2.5: The Five Forms of Internet (IP) Addresses [Comer 91]

In order to make them more readable, IP addresses are written as four base 10 numbers separated by periods, each of which represents 8 bits of the 32-bit IP address. For example the 32-bit Internet address 10000000 11101001 10000010 00000001 is written as 128.233.130.1.

Although such addresses are convenient and efficient for machines to communicate with each other, users prefer to assign machines pronounceable, easily remembered names. In order to support a large and rapidly expanding set of names, the Internet uses a hierarchical name structure. The mechanism that implements a machine name hierarchy for the Internet is called the *Domain Name System* (DNS). It is used to map human-readable machine names into Internet addresses. Each of the levels in the hierarchical name structure is referred to as a *domain*. The hierarchical naming scheme used by DNS is known as *domain names*. Conceptually, the highest level domains permit two completely different naming hierarchies: geographic and organizational. The geographic scheme divides the universe of machines by country. A 2-letter identifier is assigned for a country when this country wants to register its machines in the Domain Name System. For example, AU is for Australia, CA is for Canada, HK is for Hong Kong, FR is for France, and US is for United States. For the organizational scheme, the highest level domains allow organizations to be grouped by organizational type. For example, it is possible for a university to register itself under the EDU domain and a company under the COM domain. The highest level domains in the hierarchy (tree) are listed in Table 2.1. All hosts that are to be attached to a network or subnet of the Internet must be registered with one of these top level domains.

Table 2.1: Highest Level Internet Domains

Scheme	Domain name	Description
Geographical	Country Code	2 letters per country (CA, AU, HK, etc.)
Organizational	ARPA	Old style Arpanet
	COM	Commercial organizations
	EDU	Educational institutions
	GOV	Government institutions
	INT	International organizations
	MIL	US Military
	NATO	Nato field
	NET	Network
	ORG	Non-Profit Organization

The general scheme and naming convention are shown in Figure 2.6. A domain name consists of a sequence of subnames separated by a delimiter character, the period, and is written with the local subname first and the top domain last. This scheme provides a unique name for each host on the Internet. For example, the domain name for the Department of Computational Science at the University of Saskatchewan is *cs.usask.ca*.

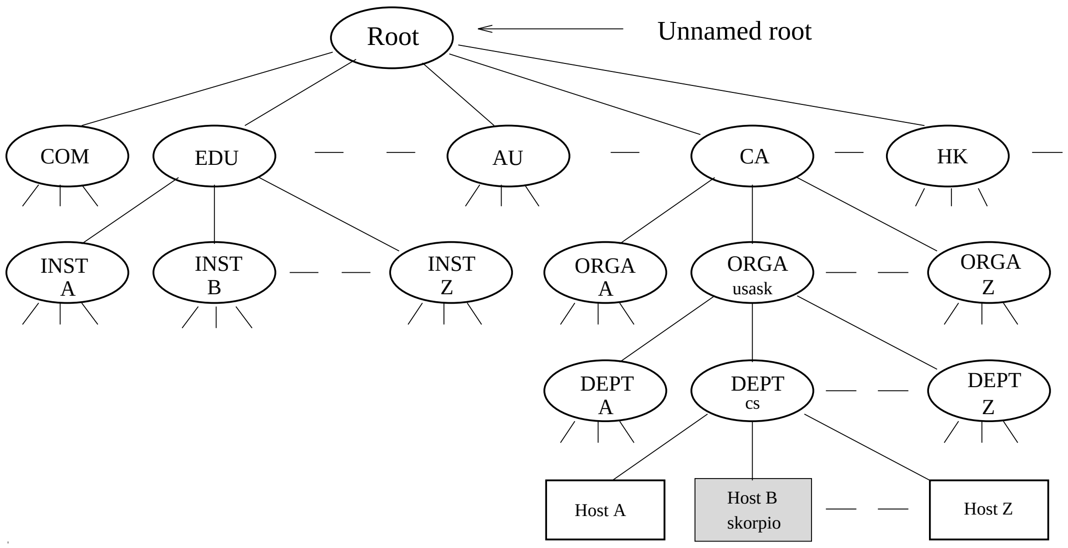

Figure 2.6: Domain Name Hierarchy

The Domain Name System allows multiple naming hierarchies to be embedded in one system. A given name may map to more than one item in the domain system. For example, domain name *cs.usask.ca* can map to any of the host machines in the Department of Computational Science, such as *skorpio.usask.ca* (shown as the shaded box in Figure 2.6).

2.4 The Internet Protocol (IP)

The Internet Protocol (IP) defines an unreliable, connectionless delivery mechanism for network packets [Comer 91, Postel 81a]. It is responsible for defining the basic unit of data transfer called an *Internet datagram*, finding a route for the datagram and getting it to the other end, and lastly, for providing a set of rules that characterize the unreliable packet delivery.

A datagram is divided into header and data sections. The format of the IP datagram header is shown in Figure 2.7.

32 bits

Version	IHL	Type of service	Total length	
Identification			Flags	Fragment offest
Time to live		Protocol	Header checksum	
Source address				
Destination address				
Options				

Figure 2.7: The Format of IP Datagram Header

The *version* field tells which version of the protocol the datagram belongs to. The *IHL* is the header length measured in 32-bit words. The *type of service* field specifies the kind of service needed, such as its reliability and speed requirement. The *total length* gives the length of the datagram measured in bytes including both header and data. Since it may be necessary to fragment the datagram into smaller units as it goes through the Internet, the *identification* field is used to let the destination host know which datagram a newly arriving fragment belongs to. All fragments of the same datagram contain the same identification value. The *flags* and the *fragment offset* are used to keep track of the pieces when a datagram has to be split up. The *time to live* field specifies how long, in seconds, the datagram is allowed to remain in the Internet system in order to prevent the datagram from traveling around the Internet forever. It is decremented each time the datagram passes through a gateway on the Internet. When the value reaches zero, the datagram is destroyed. The *protocol* field identifies the higher-level protocol that the datagram belongs to, such as TCP, UDP, etc. The *header checksum* is used as an error detection mechanism to check that the header has not been changed during transit. The most significant fields in the header are the *source address* and the *destination address* that indicate the IP address of the source and destination hosts. The destination host address is necessary for any gateways in between to know where the datagram needs to be sent to. The *options* field is used for security, source routing, error reporting, debugging, time stamping, and other optional information.

2.5 Transmission Control Protocol (TCP)

TCP is a connection-oriented, end-to-end reliable protocol designed to fit into a layered hierarchy of protocols that support multi-network applications [Postel 81b]. It is responsible for breaking up a user's message into datagrams, reassembling them at the other end, resenting anything that gets lost, and putting things back in the right order upon arrival.

The unit of transfer between the TCP software on two machines is called a *segment* [Comer 91]. It contains a header with a minimum of 20 bytes followed by data. The format of the TCP header is shown in Figure 2.8.

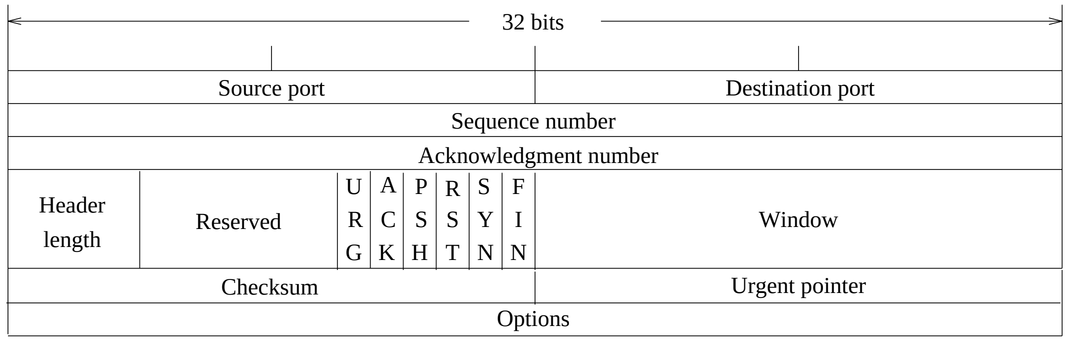

Figure 2.8: The Format of TCP Header

The *source port* and *destination port* fields identify the end points of the connection. An end point is defined as a pair of integers *(host, port)*, where *host* is the IP address for a host and *port* is a TCP port on that host. A port in the TCP protocol is an identifier that indicates a particular application to the network software. For example, the end point *(128.233.128.5, 25)* specifies TCP port *25* (SMTP) on the machine with IP address *128.233.128.5*. The *sequence number* field tells the position in the byte stream of the data from the sender in the segment. This is used so that the receiver can make sure that it gets the datagrams in the right order, and that it has not missed any. The *acknowledgment number* field contains the value of the next sequence number the sender of the segment is expecting to receive. The *header length* field tells how many 32-bit words are contained in the TCP header. The *reserved* field is reserved for future use. Following the *reserved* field, there are six control bits. They are *URG* (urgent pointer field), *ACK* (acknowledgment field), *PSH* (push function), *RST* (reset the connection), *SYN* (synchronize sequence numbers), and *FIN* (no more data from sender). The *window* field specifies the number of bytes that the sender of this segment is willing to accept. The *checksum* field provides reliability and the *urgent pointer* field allows one end to tell the other to skip ahead in its processing to a particular byte. The *options* field is used for miscellaneous things.

Each machine contains a set of TCP ports in order to to name the end points of logical connections. Each protocol port is identified by a positive integer. The local operating system provides an interface mechanism that processes use to specify a port or access it. TCP combines static and dynamic port binding, using a set of *well-known port assignments* for commonly invoked programs but leaving most port numbers available for the operating system to allocate as programs need them [Comer 91]. The well-known ports are controlled and assigned by the Internet Assigned Numbers Authority (IANA). Recently, the range of assigned well-known ports managed by the IANA has been expanded to 0-1023. Ports in the range 1024-65535 are registered but cannot be controlled by the IANA [Reynolds 92].

Table 2.2 shows some frequently used assigned TCP port numbers and their respective applications.

Table 2.2: Examples of Assigned TCP Port Numbers.

Decimal	Protocol keyword	Description
20	FTP-Data	File Transfer Protocol (data)
21	FTP	File Transfer Protocol (control)
23	Telnet	Terminal Connection
25	SMTP	Simple Mail Transfer Protocol
49	Login	Login Host Protocol
53	Domain	Domain Name Server
70	Gopher	Internet Gopher
79	Finger	Finger
80	WWW	World-Wide Web
113	Ident	Identification Protocol
119	NNTP	Network News Transfer Protocol

2.6 User Datagram Protocol (UDP)

UDP is a connectionless datagram delivery service that does not guarantee delivery in the transport level [Postel 80]. It provides a procedure for application programs to send messages known as *user datagrams*, to other programs with a minimum of protocol mechanism. Each user datagram consists of two parts, the UDP header and the UDP data area. Figure 2.9 shows the header format of the UDP datagram.

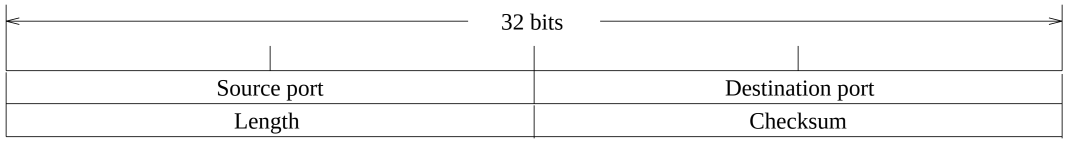

Figure 2.9: The Format of UDP Header

The *source port* and *destination port* fields are the UDP protocol port numbers. The path of communication between an application and UDP is through UDP ports. Each application program must obtain a protocol port and an associated port number from the operating system before it can send a UDP datagram.

UDP uses the same port assignment scheme as TCP. For example, the well-known port number 53 indicates a domain name server application. The *source port* is optional. When used, it specifies the port to which replies should be sent; if not used, a value of zero is inserted. The *length* field is the length of the UDP datagram in bytes, including the UDP header and the user data. The *checksum* field is optional. A value of zero in the *checksum* field means that the checksum has not been computed.

2.7 Application Level Protocols

There are many application level protocols in the TCP/IP protocol suite that permit users to use the network to carry out useful communication tasks. The number of these protocols is growing continually. Some of the applications, such as Telnet, SMTP, and FTP, have existed since the beginning of the Internet technology. Others, such as Gopher and WWW, are relatively new. The following is a brief description of some widely used applications.

2.7.1 Simple Mail Transfer Protocol (SMTP)

SMTP is used to transfer electronic mail (e-mail) reliably and efficiently [Postel 82]. It uses a TCP connection that is opened to a remote host on the SMTP port 25. When a user presents a mail request, the sender-SMTP establishes a connection to the receiver-SMTP. Once the connection is established, the sender-SMTP sends a MAIL command, indicating the sender of the mail. If the receiver-SMTP can accept mail it responds with an OK reply. The sender-SMTP then sends a RCPT command identifying the intended recipient of the mail. If the receiver-SMTP can accept mail for that recipient, it responds with an OK reply; if not, it responds with a reply rejecting that recipient but not the whole mail transaction. The sender-SMTP and receiver-SMTP may negotiate several recipients. When the recipients have been negotiated the sender-SMTP sends the mail data, terminating with a special sequence. If the receiver-SMTP successfully processes the mail data it responds with an OK reply and the connection is closed.

2.7.2 File Transfer Protocol (FTP)

FTP is used for transferring files between hosts [Postel 85]. The primary function of FTP is to transfer files among hosts efficiently and reliably and allow the convenient use of remote file storage capabilities. User FTP sessions are subdivided into FTP-Control and FTP-Data connections. The FTP-Control connection deals with the transactions related to interactive user commands, such as the commands that initiate data transfer, while FTP-Data connection deals with data transfer. Each FTP process can initiate one FTP-Control connection and zero or more FTP-Data connections. When a host wants to ftp to a remote server, the client's FTP-Control process connects to the server's FTP-Control process using a TCP connection. The port of the server used for this connection is 21. User name and password are needed by the FTP server for verification. After that, the user can execute FTP commands to list available files (*dir* or *ls*), *get* a file, or *put* a file to the remote FTP server. The data transfer of FTP is done using a separate FTP-Data connection. Each data file transfer uses a new data transfer connection and the connection is terminated when the data file transfer is done. The FTP server always sends the FTP-Data connection request to port 20 of the client. In general, the FTP-Control connection remains alive as long as the user keeps the FTP process going. Usually, FTP is used to represent FTP-Control.

An anonymous FTP permits any of the Internet users to retrieve files from any anonymous FTP server without a personal login name. This service is commonly used for publishing documents, software, and other information on the Internet.

2.7.3 A Remote Terminal Protocol (Telnet)

Telnet is a remote terminal protocol [Postel 83, Comer 91]. It allows a user to open a TCP connection to a login server on the Telnet port 23. After the connection is established, the originator sends keystrokes from the user's terminal and the responder replies with the output from the remote machine. Telnet offers three basic services. First, it defines a *network virtual terminal* that provides a standard interface to remote systems. Second, it includes a mechanism that allows the client and server to negotiate options, and it provides a set of standard options. Third, it treats both ends of the connection symmetrically.

2.7.4 BSD Rlogin Protocol (Login)

The rlogin facility is a remote login protocol for Unix machines that provides a remote-echoed, locally flow-controlled virtual terminal with proper flushing of output [Kantor 91]. Rlogin uses a TCP connection that is opened to a remote host on the Login port 513. It is widely used between Unix hosts for two reasons. First, it provides transport of more Unix terminal environment semantics than the Telnet protocol does. Second, it can be configured on many Unix hosts not to require user entry of passwords when connections originate from trusted hosts.

2.7.5 Finger User Information Protocol

Finger is a simple protocol that provides an interface to a remote user information program (RUIP) [Zimmerman 90]. It is used to present information on a remote site or a local site such as someone's login and personal names, whether or not the user you are asking about is currently logged in, etc. Finger uses a TCP connection that is opened to a remote host on the Finger port 79. It consists of the originator sending a query with the name of the remote user and electronic address to a remote host and the RUIP in the remote host processing the query, returning an answer, and terminating the connection.

2.7.6 Gopher

The Internet Gopher protocol is designed for distributed document search and retrieval [Anklesaria 93]. Originally created as a fast, simple, distributed, campus-wide information system, Gopher has become increasingly popular on the Internet. While documents (and services) reside on many servers, Gopher client software presents users with a hierarchy of items and directories much like a file system. The Gopher protocol consists of a client connecting to a server and sending to port 70 on the server a selector (a line of text, which may be empty) via a TCP connection. The server responds with a block of text terminated with a period on a line by itself, and closes the connection. If the chosen resource is a file, the Gopher server will *ftp* it for the user. If the resource is a "Login" resource, it creates a *Telnet* session. If the resource is a collection indexed by Archie (an electronic directory service for the Internet) or wide area information server (WAIS) [Danzig 93], the Gopher server uses Archie or WAIS to find out what is relevant. Each service request from the client uses one Gopher connection. The Gopher server terminates the connection after sending the response.

When looking for resources with Gopher, two general areas are of interest: the list of other Gopher servers available, and the list of information services provided by the particular Gopher.

2.7.7 Network News Transfer Protocol (NNTP)

NNTP [Kantor 86] specifies a protocol for the distribution, inquiry, retrieval, and posting of news articles among the Internet hosts using TCP connections. New articles are stored in a central database and separated by newsgroups. Newsgroups are shown in a catalogue by topics. Subscribers can select only those items they want to read. Indexing and expiration of aged messages are also provided. When a client host wants to have news services, it establishes a TCP connection to port 119 at the news server. After this, the server sends a series of newsgroups to the host. When the client responds with selecting a newsgroup and one news article from this newsgroup, the server transferred the article to the client. Hosts can post their news, save the news articles, etc. The TCP connection is terminated when the client finishes using the news server.

Usually each newsgroup will maintain a list of *Frequently Asked Questions* (FAQs) to help new users learn what the newsgroup is about.

2.7.8 World-Wide Web (WWW)

The World-Wide Web (WWW) is a way of viewing all the online information available on the Internet as a seamless, browsable continuum [Berners-Lee 93]. It merges the techniques of information discovery, wide area networking, and hypertext. Hypertext is a method of presenting information where selected words in the text can be "expanded" at any time to provide other information about the word. That is, these words are links to other documents, which may be text, files, pictures, or anything else [Krol 92]. Using hypertext, users navigate through an information world from existing databases and information systems. WWW clients provide comprehensive point and click network access tools, while WWW servers provide an efficient and friendly method to fetch data to users.

Besides its native HyperText Transfer Protocol (HTTP) [Danzig 93], WWW clients understand FTP, Gopher, NNTP, and WAIS. HTTP is a new document retrieval and full-text search protocol. HTTP objects are identified by the HTTP protocol type, the corresponding name of server, and the path name to the file where the contents of objects reside.

To access the WWW server, a browser program is needed to fetch documents that are set up by information providers. The browser program in the client establishes a TCP connection to port 80 at the WWW server, sends the information request, and reads the resulting response. This is done repeatedly. The connection is closed by the server upon completion of each response. When servers other than WWW servers are requested by clients, TCP connections other than WWW connections are established between the client and the selected server. For example, when the client selects a FTP server, a TCP connection from the client will be established to port 21 at the FTP server. All requests and responses between the client and the FTP server are done using an FTP session. The connection is closed when the client finishes using this FTP server.

Currently, a public-domain browsing tool called *Mosaic*, from the National Center for Supercomputing Applications (NCSA) in Urbana-Champaign, is the most popular client interface for accessing WWW servers.

2.7.9 Domain Name Server Protocol (Domain)

The Domain Protocol is used for exchanging hostname information among the distributed name servers on the Internet. When a client invokes an application program and supplies the name of a host with which the application will communicate, the application program must find the IP address of this host. In order to fetch the IP address, the client passes the domain name to a local name server, which maintains caches of hostnames and their corresponding IP address. If the IP address of this domain is found in the cache of the local name server, the IP address is returned to the client. If not, the local name server presents itself as a client, formats a message, finds the best servers to ask, sends the message to port 53 [Mockapetris 87] at these name servers, and waits for a reply. If a timeout occurs, a failure message will be sent back to the client.

Almost every wide area TCP connection, whether it be SMTP, Telnet, or FTP, begins with two or three DNS requests. This might not generate wide area traffic because some mappings could be cached locally [Danzig 92b].

Name servers manage two kinds of data. One kind of data is called *zones*. Each zone is the complete name service database for a particular subtree of the domain architecture. A name server periodically checks to make sure that its zones are up to date. If not, it obtains a list of updated zones from master files stored locally or in another name server. When the Domain protocol uses TCP connections for transferring zone information between hosts, the messages will be large. The other kind of data is cached data that was acquired by a local name resolver. This data improves the performance of the name retrieval process when non-local data is accessed repeatedly. Cached data is eventually discarded by a timeout mechanism.

2.7.10 Identification Protocol (Ident)

The Identification Protocol [Johns 93] is used to determine the identity of a user of a particular TCP connection. It is not intended as an authorization or access control protocol. It is used only to provide some additional auditing information with respect to TCP connections.

When an Ident client receives a TCP connection request, it sends another TCP connection to the host that originated the previous TCP connection. Now the previous originator is an Ident server. The Ident server listens for TCP connections on its port 113. After an Ident connection is established, the server reads the query from the client. Usually the query is asking about the user information of the previous TCP connection. If an answer exists, it is sent from the Ident server to the Ident client. The server may then either shut the connection down or continue to answer multiple queries.

2.8 A Wide Area Conversation

Since application programs running on hosts are the ultimate sources of network traffic and TCP packets make up a large amount of all Internet network traffic, the characteristics of TCP applications are the focus of this research.

In general, applications fall into two broad categories: *bulk transfer* and *interactive*, each of which may offer a different workload to the network and demand a different level of performance from the network. A *bulk transfer* application is one where a computer sends a batch of data without direct human intervention. An example is an exchange of e-mail between two hosts. An *interactive* application is one where human actions are directly involved in generating individual packets. An example is a remote terminal session where a human generates the keystrokes that go over the network as individual packets [Caceres 93].

Central to TCP applications is the notion of a *conversation*. A conversation refers to the stream of packets flowing between an application program on one Internet host and an application program on another Internet host. To establish a TCP conversation, the originator picks an initial *sequence number* and sends it to the responder in a TCP packet with the SYN (*synchronize sequence numbers*) bit set in the header. If the responder accepts the conversation, it replies to the originator with another TCP packet with a separate sequence number and the SYN bit set, acknowledging the first sequence number given in the packet from the originator. Figure 2.10 shows the creation procedure for a TCP conversation (a "three-way handshake").

The conversation is established after the originator sends the final acknowledgment to the responder, and the two hosts can now exchange packets. Each new packet contains an acknowledgment number equal to the sequence number of the last packet received plus the number of bytes in that packet. Therefore, all sequence numbers are given as the sum of the initial sequence number plus the number of bytes transmitted so far.

When either host wants to close the conversation, it sends a packet with the FIN (*Finish*) bit set to its partner and the partner then replies with another FIN

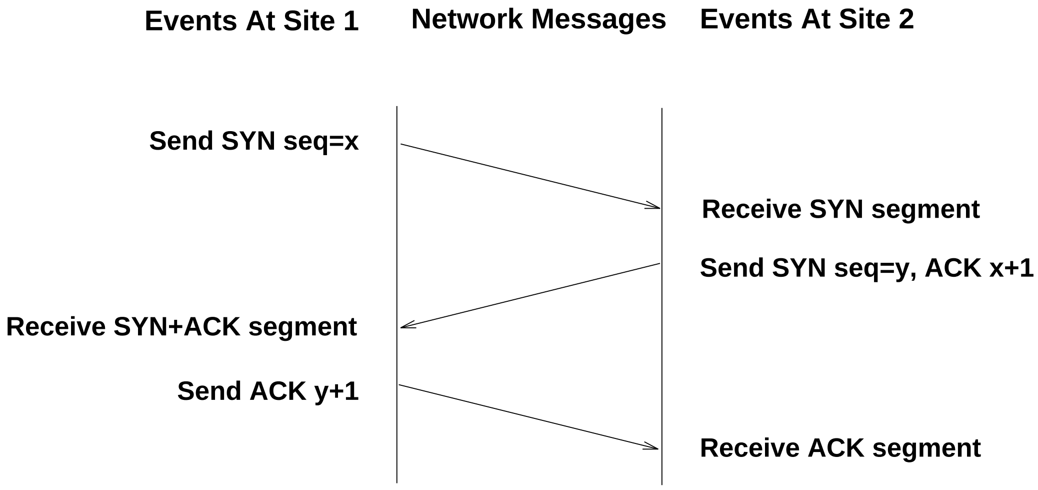

Figure 2.10: The Creation Procedure of a TCP Conversation [Comer 91]

packet. All packets received will be acknowledged and the acknowledgments are sent to the partner host. At this point, the conversation is terminated. A more detailed description of the termination procedure of a TCP conversation is shown in Figure 2.11.

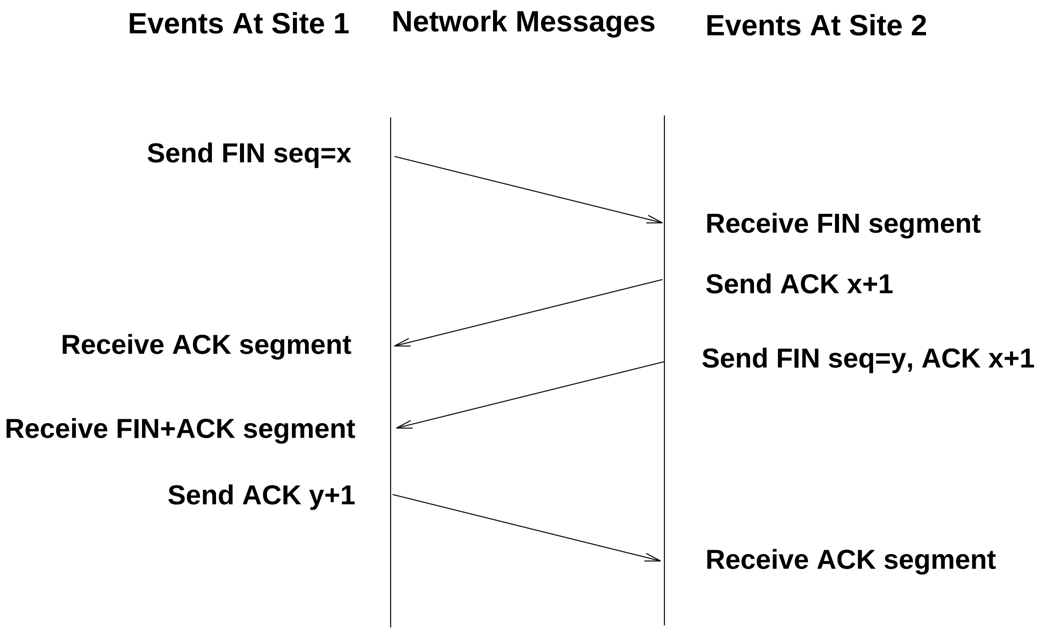

Figure 2.11: The Termination Procedure of a TCP Conversation [Comer 91]

Conversations with the whole creation and termination processes described above are called *complete conversations*. Sometimes abnormal conditions arise that force an application program or the network software to break a conversation. To reset a connection, one side sends a packet with the RST (*reset*) bit set to its partner. The other side responds with a RST packet to abort the conversation.

Conversation size, conversation duration, conversation bandwidth, conversation interarrival time, and the number of concurrent conversations are several important parameters to characterize wide area conversations. It is sufficient to extract them by the information implied in SYN and FIN packet headers.

The *conversation size* is the number of data bytes transmitted in both directions by each conversation. The *conversation duration* is the lifetime of a conversation. It is calculated by taking the difference of the timestamp values between the first TCP SYN packet and the last TCP FIN packet of that conversation. The *conversation bandwidth* is a rate of information exchange computed by dividing the conversation size by the conversation duration. The *interarrival time of conversations* is defined as the time interval between the start of one conversation (from any hosts) and the start of the next conversation (from any hosts) and it is calculated by taking the difference of the timestamp values of the starting time between adjacent conversations. The *number of concurrent conversations* is the number of conversations occurring over a particular time interval, which is chosen to be 15 minutes in this study. It is calculated by incrementing a counter, which is initialized to be 0 at the beginning of each time interval, when a conversation begins in this time interval or when a conversation begins before this time interval but is still active in this time interval.

Part of the conversation creation process is the selection by the originator of the TCP ports that will be used by the conversation. Because TCP identifies a connection by a pair of end points (application programs), a given TCP port number can be shared by multiple connections on the same host. A host can have more than one conversation at the same time as long as the pairs of port numbers of the conversations are unique.

2.9 Summary

The Internet is a world-wide "network of networks" that are connected to each other using the TCP/IP protocol suite. It provides file transfer, remote login, electronic mail, network news, and other services on top of both reliable, full duplex stream delivery and unreliable, connectionless packet delivery services. In order to study the behaviour of the Internet, measurements of several aspects have been undertaken by previous researchers. The next chapter surveys what has been done in this area.

Chapter 3

Previous Work

For many years, networking communities have enjoyed the power of electronic communications and data transfer through the Internet. As the Internet grows, in both geography and usage, drawing maps of and measuring the Internet have become popular activities. Many studies of Internet behaviour have appeared in the literature in recent years. In this chapter, statistics relating to the Internet are reviewed to provide a preliminary picture of Internet growth, end-to-end behaviour of paths in the Internet, and characteristics of Internet traffic.

3.1 Internet Growth

Network growth statistics show that the Internet traffic, size, and the count of reachable domains have grown very fast with time.

Lottor [Lottor 92] provides statistics on the growth of the Internet by examining the number of Internet hosts and domains over a 10-year period (from August 1981 to January 1992). During those years, the number of Internet hosts has grown exponentially (a mean growth rate of 126%/year), although the rate of increase has decreased (a mean growth rate of 109%/year from January 1989 to January 1992). The number of Internet domains has also increased, with an annual growth rate of 161%/year from July 1988 to July 1991. A total of 17,000 domains were found in January 1992.

More recent data [Lottor 94] shows the growth rate of Internet hosts dropped to 82.54%/year from October 1991 to October 1993. There were 2,217,000 hosts and 30,000 domains on the Internet in January 1994.

NSFnet statistics are reported at monthly intervals. Merit [Merit 94] shows there are 81 countries having access to the NSFnet or general electronic mail services as of June 1994 and the number of networks connected to the NSFnet has increased from 173 in 1988 to 14,964 in 1994 (a mean growth rate of 111%/year). The traffic on the NSFnet (measured in packets) has also been growing exponentially, from 8.5×10^7 packets/month in January 1988 to 5.79×10^{10} packets/month in February 1994. The byte traffic has increased from 1.27×10^{10} bytes/month in March 1991 to 1.1×10^{13} bytes/month in February 1994.

The Internet Activities Board (IAB) has estimated that the Internet has a monthly growth rate in the number of computer hosts of 10 - 15%, making it by far the most rapidly growing electronic network in the world [Lynch 93]. At the time of Lynch and Rose's writing, traffic growth on the major backbones of the Internet exceeded 25% per month.

After the NSFnet was established, it evolved from a 56 Kbps network in the mid-1980s, to a 1.5 Mbps (T1) network in July 1988, to a 45 Mbps (T3) network in May 1992. The long-term traffic characteristics of the T1 NSFnet backbone were studied by Claffy *et al.* [Claffy 93]. Their statistics show that the traffic, both in packets and in bytes, and the number of networks increased steadily from 1988 to 1990, but began to decrease in late 1991 because of the migration of traffic to the T3 backbone. Their statistics suggest that the increase in traffic volume in bytes seemed to be quadratic, while the increase in number of networks served appeared to be linear.

NSFnet statistics have been published by the Internet Society (ISOC) [Rutkowski 94]. A new traffic record was reached during the month of March 1994 when a total of 11.226 Terabytes (10^{12} bytes) were transmitted through the NSF backbone, nearly a 2 Terabytes increase from February 1994. The use of browsing services continues to grow at remarkable rates. Gopher traffic grew by 17.6% and WWW (HTTP) traffic grew by 32.9% to a new high of 0.5 Terabytes per month.

Indeed, WWW traffic grew by a total of 0.7% of the total NSFnet traffic. It is now the fifth ranking service after FTP, Telnet, network news, and SMTP mail. The statistics also present a summary breakdown of the Internet networks: 7% of the networks belong to government, 8% to education, 10% to defense, 31% to commercial, and 44% to research organizations.

On a smaller scale, growth trends of wide area TCP connections in a medium-sized research laboratory were studied by Paxson [Paxson 93a] over a period of more than two years. His results showed that the site traffic grew exponentially and at a rate greater than that at which the site added computing resources.

Schwartz *et al.* [Schwartz 93] presented measurements of Internet growth based on tests of what services could be reached at over 13,000 sites, in four different measurement cycles over the course of a year. Their analysis uncovers the count of reachable domains over time is currently experiencing exponential growth. Telnet, FTP, and SMTP remain the most commonly available services, although there was a slight decrease in the final measurement cycle.

Since the Internet continues to grow bigger and bigger, understanding of its behaviour in general is needed. Measurements of Internet behaviour and Internet traffic have been taken by many experimenters. The following sections present the results of some measurement studies of end-to-end behaviour and traffic characteristics.

3.2 End-to-end Behaviour of the Internet

The end-to-end behaviour of a path on a network is characterized using measures such as round-trip time, packet loss, throughput, packet reordering, and duplicates.

Mills [Mills 83] analyzed ICMP [Deering 91] Echo/Reply messages with embedded timestamps and reported a strong correlation between Internet delays and packet size, with the longest packets showing delays two to three times that of the shortest packets. There were two sets of measurement experiments. One of these was based at the local area network at the Linkabit Corporation offices in McLean, Virginia, while the other was based at the local area network at the Norwegian Telecommunications

Administration (NTA) near Oslo, Norway.

Mukherjee [Mukherjee 92] presented a detailed study of Internet load using measured statistics of end-to-end delays, packet losses, and out-of-order acknowledgments. Several segments of the Internet, including a regional network in the New Jersey/Pennsylvania area (the Jon Von Neumann Center Network), a segment of the NSFnet backbone, and a cross-country network consisting of regional and backbone segments, were studied. His results show that the distribution of delay is a constant plus a Gamma-distributed random variable for all three network paths studied. The shape and scale parameters for this distribution change with load. The correlation between packet loss and delay varies considerably across networks. Low loads show lower correlations between delay and loss while higher loads show higher correlations. The number of out-of-sequence acknowledgments is positively correlated with the different delay statistics for the cross-country segment. Analogous to the packet-loss case, the conditional expectation of fraction of out-of-sequence acknowledgments also increases as a function of delay.

Sanghi *et al.* [Sanghi 93] conducted a series of experiments to study end-to-end behaviour of the Internet. The experiments were conducted using three different hosts located at the University of Maryland, Stanford University, and MIT. The measurements indicate that the IP level service provided in the networks yields high losses, duplicates, and reordering of packets. There were 2.1% to 10.1% packet losses in different experiments. In several experiments, the number of duplicates was approximately 1%. Reordering of packets was observed in all experiments. In addition, the round-trip transit delay varied significantly. The maximum round-trip time was 2,339 msec and the minimum round-trip time was 19 msec. The maximum average round-trip time over all the experiments was 114 msec.

Claffy *et al.* [Claffy 93] also studied delay statistics on the NSFnet backbone by recording the round-trip time between all pairs of backbone nodes in May 1992. Delay statistics showed that typical end-to-end delays on the NSFnet backbone do not exceed 100 msec and typical link delays do not exceed 45 msec.

3.3 Traffic Characteristics on the Internet

Internet traffic studies usually fall into three categories: backbone level, TCP conversation level, and packet level.

Heimlich [Heimlich 90] studied the packet arrival process on the NSFnet and verified that the packet train model [Jain 86] can be used to characterize the combined traffic as well as specific application and transport protocols.

Asaba *et al.* [Asaba 92] evaluated the current utilization of the international link between Japan and Hawaii and evaluated the scope of international information exchange grounded in the Internet application protocols. All TCP traffic between Japan and other nations passing through an Ethernet at Shounan Fujisawa Campus, Keio University, was considered in this analysis. Traces were captured for one week in January 1992. The TCP traffic occupied an average of approximately 9.5% of the link over the seven days, a relatively light load, but the maximum inbound traffic rate was 133 Kbps as compared to the outbound rate of 37 Kbps. Many connection requests originated in Japan from noon to midnight Japanese time, and significantly fewer connection requests occurred during the morning. FTP-Data traffic accounted for 77% of the total inbound traffic volume, and 71% of the total outbound traffic volume. 37 domains including 31 countries were observed in the trace.

Wakeman *et al.* [Wakeman 93] analysed the international IP traffic that passed along the UK-US academic network link, as well as some traffic between UK academic networks and sites in Europe and the US. The study was carried out in March and August of 1991 at the University of London Computing Center (ULCC). At the time of this study, UK-US academic networks (JANET and NSFnet) were interconnected by a 384 Kbps leased line. The statistics presented for IP traffic include the number of packets per protocol, number of packets by geographic destination, change in traffic patterns over the day, DNS (UDP) analysis, and statistics of TCP conversations. Packet train statistics of wide area conversations were calculated as well.

Paxson [Paxson 91] described measurements of wide area network TCP conversations between a local area network at the Lawrence Berkeley Laboratory and the rest of the world on the Internet, for the months of November 1990 and March 1991. Overall characteristics of these conversations, such as the distributions of amount of data transferred, network bandwidth used, conversation lifetimes, and conversation interarrival times were studied. Temporal traffic variation and long term variation were investigated as well. Traffic to 48 of the 50 states in the United States and 23 foreign countries was discovered. A more recent study [Paxson 93b] provided a diverse look at wide area traffic. 2.5 million TCP connections that occurred during 14 wide area traffic traces gathered at five "stub" networks and two internetwork gateways were analyzed. Telnet, NNTP, SMTP, and FTP connections were studied in detail. Some of the major findings were random variables associated with wide area network connections can be described as well by analytic models as by empirical models; bulk transfer traffic is best modeled using log-normal distributions; except for NNTP, connection interarrivals are well modeled using a non-homogeneous Poisson process with fixed hourly rates.

Caceres *et al.* [Caceres 91, Caceres 92, Danzig 92a] presented conversation-level analysis of wide area TCP traces collected on three local area networks: two campus networks (the University of Southern California (USC) and the University of California at Berkeley (UCB)) and one industrial research site (Bellcore). Traces from UCB and USC were collected over a period of one day and traces from Bellcore were collected over a period of three days. The statistics show that bulk sources transfer large amounts of data per conversation and they send large packets in only one direction, while interactive sources send small packets in one direction and receive echoes of comparable size in the opposite direction. The statistics also suggest that internetwork traffic can be modeled by either a Poisson interarrival process or a packet train model.

Schmidt and Campbell [Schmidt 92] reported statistics on TCP conversations and IP traffic from data captured at the CSOnet and NSFnet gateways. The University of Illinois serves as one of a number of Internet backbone sites throughout the United States and CSOnet is a departmental local area network in the Computing Services Office at the University of Illinois at Urbana-Champaign. Samples were taken in periods of 24 hours during weekdays in December 1991 and January 1992. On the CSOnet, 60% of the packets were TCP and 25% were UDP. The NSFnet data contained 85% TCP and 15% UDP packets. The statistics presented include the number of concurrent conversations, conversation start up time, and the relationship between connection lifetime and bytes transmitted, etc. The mean number of concurrent connections on the NSFnet was 209.3. The mean start up time was 6.59 msec on the CSOnet and 10.73 msec on the NSFnet, respectively. They also concluded that the longer conversation interactive sessions did not show a great deal of traffic.

Both Caceres *et al.* [Caceres 91] and Claffy *et al.* [Claffy 93] reported that locality of reference exists between network pairs in wide area internetworks.

Some experimenters have studied application traffic in detail. Danzig *et al.* [Danzig 92b] studied wide area name server traffic of the Internet's Domain Name System (DNS). This study was based on two 24-hour packet traces collected from DNS root name server *a.isi.edu* and three other name servers that replicated various domains. They found that the portion of wide area network traffic due to DNS would decrease as defective name servers and resolvers are replaced. Defective implementations of name servers for DNS should be actively checked to reduce the number of packets to the Internet. Negative caching, which refers to the caching of bad names or the absence of a resource record in order to answer future queries quickly, is unnecessary on properly running name servers on the Internet although it can serve as a firewall between a wide area network and malicious or badly broken programs.

Maffeis [Maffeis 93] studied FTP traffic from an investigation of access to public files which took place over three months. The results show that most file transfers relate to only a small fraction of files in an archive and that a considerable part of the operations to public files are updates of files.

3.4 Summary

This chapter has reviewed some measurement studies on the Internet. For this research, the Internet traffic study will focus on the conversation level and will be carried out on the University of Saskatchewan Campus Ethernet. The next chapter presents details of the measurement methodology. Results of the analysis are presented in Chapter 5.

Chapter 4

The Measurement Methodology

The goal of this research is to characterize wide area conversations between the hosts at the University of Saskatchewan and the rest of the world on the Internet. Since TCP traffic makes up a large amount of the wide area traffic, TCP conversations are the focus in this research. This chapter presents the measurement methodology for this research.

Section 4.1 describes the measurement environment and the measurement tool used in the research. Section 4.2 describes the processes of data collection and data reduction.

4.1 Measurement Environment and Tool

A study of traffic flow between a local area network and a wide area network was carried out, using the 10 Mbits/sec Campus Ethernet at the University of Saskatchewan as the local area network and the Internet as the wide area network. The wide area traffic between the Department of Computational Science Ethernet (a subnet of the Campus Ethernet, which is referred to as the "CS Ethernet"), a subset of the campus wide area traffic, was also studied in detail. Figure 4.1 shows the nature of the connection between these networks. The Campus Ethernet is a collection of heterogeneous networks, one of which is the CS Ethernet. These networks are connected to the campus network which, in turn, is connected to Sask#net. Sask#net is the

Saskatchewan regional network that connects all of the research and education sites in Saskatchewan. It, together with other regional or provincial networks, connects to CA*net, which is the Canadian backbone network, and to the NSFnet which, in turn, connects to the backbone networks of other countries.

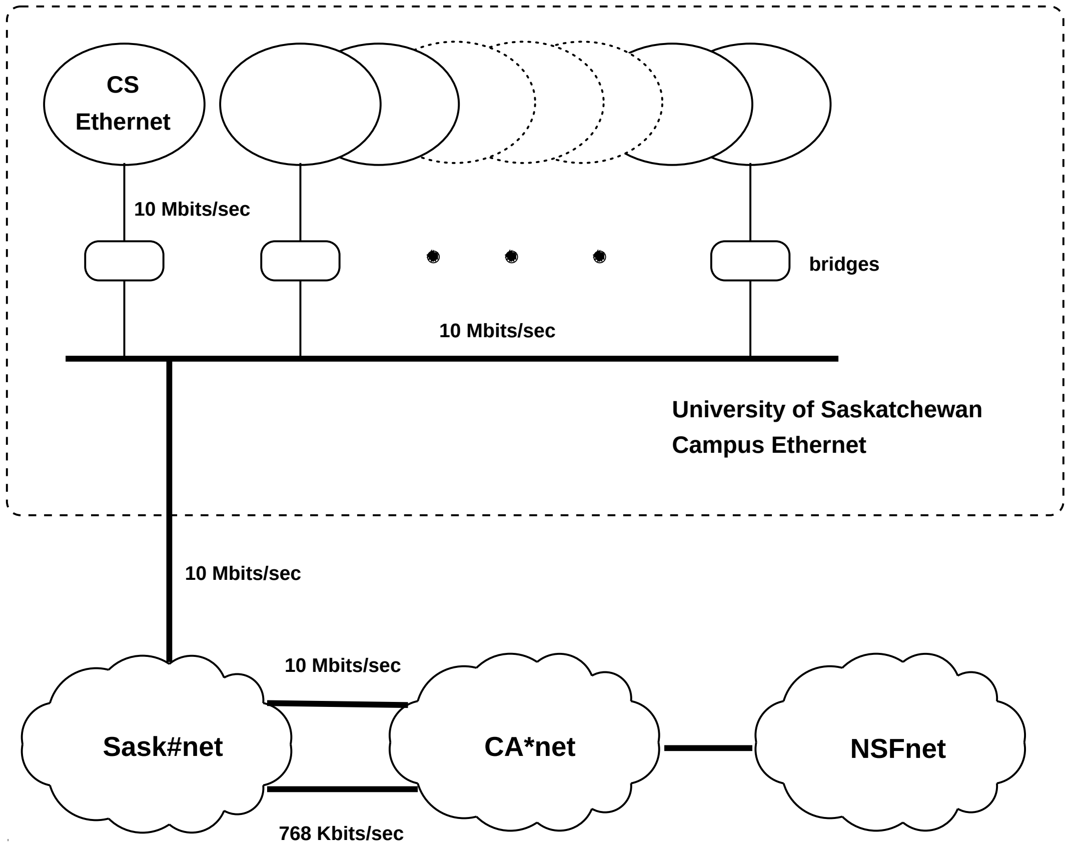

Figure 4.1: The Network Measurement Configuration

The link between the Campus Ethernet and Sask#net runs at 10 Mbits/sec. There are two kinds of links between Sask#net and CA*net. One is a 10 Mbits/sec link between the Saskatchewan regional router and the CA*net Edmonton (Alberta) backbone router. The other is a 768 Kbits/sec link between the Saskatchewan regional router and the CA*net Winnipeg (Manitoba) backbone router. In April 1994, traffic between Saskatchewan and the United States (NSFnet) was taking the path Saskatchewan – Manitoba – Quebec – United States because of temporary bandwidth problems on the Alberta – British Colombia – United States link. The link between Manitoba and Quebec runs at T1 speed (1.5 Mbits/sec) and the link between Quebec and the United States runs at 3.0 Mbits/sec.

For this analysis, data (TCP or IP packets) was gathered on a SUN Sparcstation 1+ workstation, which was installed and configured on the Campus Ethernet using the *TCPDUMP* program written by Jacobson *et al.* [Jacobson 89]. Only the wide area traffic between the Campus Ethernet and the Internet was captured. Wide area traffic between the Internet and the CS Ethernet was extracted from the Campus trace using the list of the IP numbers assigned to the Department of Computational Science.

TCPDUMP is a network monitoring and packet acquisition software tool widely used on UNIX systems to print out the headers of selected packets on a network interface that match a user-supplied boolean expression. Because TCPDUMP runs as a user-level process, the kernel must provide some facility that gives the user-level program access to raw, unprocessed network traffic. In this experiment, TCPDUMP works with the Berkeley Packet Filter (BPF) [McCanne 93], a new kernel architecture for packet capture. BPF has two main components: the network tap and the packet filter. The network tap collects copies of packets from the network device drivers and delivers them to the TCPDUMP program. The filter decides if a packet should be accepted and, if so, how much of it to copy to the TCPDUMP program. BPF shows good performance since it copies only the selected packets to memory and it can capture packets that are generated by the local host on which the BPF resides as well.

A TCPDUMP command line is as follows:

tcpdump [options] expression

TCPDUMP provides 19 possible options for users to describe their data collection requirements. The *expression* field selects which packets will be dumped. TCPDUMP prints out the headers of packets for which the boolean *expression* is 'true'. The *expression* consists of one or more primitives that usually consist of an *id* (name or number) preceded by one or more qualifiers. There are three different kinds of qualifier. *Type* qualifiers say what kind of thing the *id* name or number refers to.

Possible types are host, net, and port. *Dir* qualifiers specify a particular transfer direction to and/or from *id*. Possible directions are src, dst, src or dst, and src and dst. *Proto* qualifiers restrict the match to a particular protocol. Possible *proto*s are ip, tcp, udp, etc. For example, *'ip and host skorpio'* captures all the IP packets whose source host or destination host is *skorpio*. There are some other allowable primitives, such as *expr relop expr* where *relop* is one of $>$, $<$, $\geq$, $\leq$, $=$, !=, and *expr* is an arithmetic expression composed of integer constants, the normal binary operators (+, -, *, /, &, |), a length operator (the length of the packet), and special packet data accessors. To access data inside the packet, the following syntax is used: *proto [expr : size]*. *Proto* indicates the protocol layer for the index operation. *Expr* gives the byte offset which is relative to the indicated protocol layer. *Size* is optional and indicates the number of bytes in the field of interest. For example, the 13th byte of a TCP datagram (see Figure 2.8) contains all of the TCP flags, the 0th, the 1st, and the 2nd bits of which represent the FIN, SYN, and RST flags, respectively. For this experiment, only those TCP packets with one of the FIN, SYN, or RST flag set in their headers are of interest. The expression 'tcp[13] & 7 != 0' limits the collection to TCP packets with any of the FIN, SYN, or RST flag set.

4.2 Data Collection and Data Reduction

The commands that were used are as follows:

tcpdump -S -n -w trace-file 'ip and not src and dst 128.233 and not host 0.0.0.0' 2> error

tcpdump -S -n -w trace-file 'tcp[13] & 7 != 0 and not src and dst 128.233 and not host 0.0.0.0' 2> error

Option -S is for printing absolute TCP sequence numbers rather than relative ones. Option -w is to write the raw packets to *trace-file*. The *trace-file* is a binary file, which can only be read using command *tcpdump -r*. Option -n means that Internet addresses are not translated to host names while data capturing. Both -w and -n are for speeding up TCPDUMP processing since it can take some time for TCPDUMP to print the ASCII version of each packet and translate Internet addresses to host names. *Tcp[13] & 7 != 0* means only the TCP packets with any of the FIN, SYN, or RST flag set will be captured. *128.233* is the *netid* of the local area network at the University of Saskatchewan. *Not src and dst 128.233* means only those IP or TCP packets that the source host and the destination host are not both on the network 128.233. *Not host 0.0.0.0* is for preventing some local area broadcasting packets from being captured. File *error* saves the messages of stop status of TCPDUMP including how many packets were received by the packet filter and how many packets have been dropped by the kernel since the user process running the TCPDUMP is not fast enough to process the packets that are accepted by the packet filter and are buffered inside the kernel.

Starting and stopping trace capturing are controlled by a series of UNIX commands. To capture a trace starting from 0 AM (midnight) on November 28 and lasting until 0 AM (midnight) on November 29, the following commands are used.

echo "sh begin" | at 00:00 Nov 28

echo "sh end" | at 00:00 Nov 29

Begin is a Unix shell script that contains the TCPDUMP command used.

Begin:

tcpdump -S -n -w trace-file 'tcp[13] & 7 != 0 and not src and dst 128.233 and not host 0.0.0.0' 2> error

or

tcpdump -S -n -w trace-file 'ip and not src and dst 128.233 and not host 0.0.0.0' 2> error

End is another Unix shell script that contains the following UNIX commands to kill the TCPDUMP process.

End:

pid=`ps -ax | grep trace-file | nawk 'print $1'`

kill $pid

The translation of Internet addresses to host names was done after data capture. Option -r is used to translate the binary file to ASCII file. The command is as follows:

tcpdump -S -r trace-file > data-file

A preliminary trace that contained all IP packets (subsequently referred to as the IP trace) was collected for 24 hours on April 22, 1994, in order to get an overview of the breakdown of wide area network traffic by protocols. The preliminary trace consists of 6 subtraces, 4 hours per subtrace. The preliminary trace was about 594 Mbytes in size, with a total of 6,556,683 packets collected. Only 3,133 packets (0.05% of the total) were dropped by the user process running the TCPDUMP. Table 4.1 shows the number of packets in each subtrace of this preliminary trace.

Table 4.1: Packet Counts of Subtraces from the IP Trace

total	subtrace1	subtrace2	subtrace3	subtrace4	subtrace5	subtrace6
6,556,683	1,005,472	870,149	1,083,634	1,318,907	999,612	1,278,909

The analysis of the preliminary trace shows that TCP packets make up the majority (96.01%) of all wide area IP packets. For this reason it makes sense to concentrate the detailed analysis on TCP traffic. A second trace (subsequently referred to as the Campus TCP trace) was collected for the more detailed study. This trace consists of four consecutive days of TCP traffic.

The focus of this investigation is on the TCP *conversation* as defined in Section 2.8. With that in mind, it is not necessary to keep many of the TCP packets observed, but only those that mark the beginning or end of a conversation. As described in Section 2.8, such packets have either the SYN, FIN, or RST bits set in them. In other words, SYN, FIN, and RST packets are sufficient to characterize TCP conversations and thus are all that need to be recorded for this analysis. Since such packets made up only 2.47% of the traffic observed in the preliminary IP trace (a total of 161,809 packets), this also provides a significant reduction in storage requirement.

The Campus TCP trace contained all such TCP packets and was collected from April 25 to 28 (4 consecutive weekdays from Monday to Thursday). This trace consists of 4 continuous subtraces, each of which is 24 hours long. The Campus TCP trace is about 64 Mbytes in size, with a total of 617,806 packets collected (about 154,452 packets per day). No packet was dropped by the user process running the TCPDUMP.

Each record in this trace has the following format.

timestamp src.port > dst.port: flag data-seqno ack window urgent options

Src and *dst* are the IP addresses of the source host and the destination host. The *port* fields contain TCP port numbers that identify the application programs at the ends of the connection. *Flag* is some combination of S (SYN), F (FIN), P (PUSH), or R (RST), such as S, SP, F, R, etc. *Data-seqno* describes the sequence number of the first and the last data byte in this packet. *Ack* is the sequence number of the next data byte expected to come from the other direction on this connection. Other fields are not important to this study. The following is an example of such a trace record:

10:17:49.283413 skorpio.4525 > marsh.cs.curtin.edu.au.ftp: S
1282752000:1282752000(0) win 4096

This trace record is interpreted as follows: at 49.283413 seconds after 10:17 AM, port 4525 of host *skorpio* requests to establish an *FTP* conversation to host *marsh.cs.curtin.edu.au*. A random number (1282752000) is chosen to be the sequence number. No data is sent in this packet. The window size (the number of bytes of receiver buffer space that is available on the other direction of this connection) used is 4096 bytes. This is not important here.

A complete conversation (as defined in Section 2.8) usually contains many such records. As long as there is one packet with the SYN flag set and one packet with the FIN flag set sent in each direction of one conversation, it is assumed that a complete conversation is found. A conversation that has more than 4 such packets in both directions is also considered to be a complete conversation. This is because hosts re-send hand-shake packets when a timeout occurs before the ACK packet is received.

The raw data is reduced so that a complete conversation is described in one record. The reduced data is saved in a file, each record of which consists of 10 fields as follows:

hour minute second src dst duration application src-dst-bytes dst-src-bytes direction

(*Hour, minute, second*) denotes the start time of the conversation. *Src* and *dst* are the IP addresses of the source host and the destination host. *Duration* is the conversation duration. *Application* is the application protocol used by this conversation, such as FTP, SMTP, etc. *Src-dst-bytes* is the number of bytes transmitted from the source host to the destination host and *dst-src-bytes* is the number of bytes transmitted from the destination host to the source host. *Direction* is a flag that identifies the conversation initiator. When *direction* is 1, the conversation is initiated by a host at the University of Saskatchewan. Otherwise, the conversation is initiated by a host on the Internet.

To illustrate the data reduction process, the following example shows the record(s) for one complete conversation before and after data reduction.

Records before reduction:

08:40:06.017467 Sun.COM.1686 > skorpio.smtp: S
1142016000:1142016000(0) win 4096

08:40:06.019736 skorpio.smtp > Sun.COM.1686: S
1273920000:1273920000(0) ack 1142016001 win 4096

08:40:09.419736 Sun.COM.1686 > skorpio.smtp: F
1142019475:1142019475(0) ack 1273920331 win 4096

08:40:09.799100 skorpio.smtp > Sun.COM.1686: F
1273920331:1273920331(0) ack 1142019476 win 4096

After data reduction, the following record is built:

8 40 6.017467 Sun.COM skorpio 3.781633 smtp 3474 330 0

This record indicates that at 6.017467 seconds after 08:40 AM on the day that the data was collected, an *SMTP* conversation began from a host at *Sun.COM* to host *skorpio* at the University of Saskatchewan. The conversation duration was 3.781633 seconds (from 08:40:06.017467 to 08:40:09.799100). During the conversation 3474 bytes (1142019475-1142016000-1) were sent from the originator to the responder and 331 bytes (1273920331-1273920000-1) were sent from the responder to the originator. The conversation size was 3805 bytes (3474+331). This conversation was initiated by the host in *Sun.COM*. As can be seen, data has been compressed without losing any meaningful information.

Incomplete conversations are discarded in the data reduction process. Incomplete conversations would be either conversations that began before April 25 or ended after April 28, or not complete because the kernel was not fast enough and caused packet loss during the trace capturing (4.44% of the total packets), or conversations that were terminated because of abnormal conditions (16.89% of the total conversations). After the data reduction, 103,016 complete TCP conversations were stored in the final file (named *CAMPUS-TCP-1*): 24,844 for the first day, 27,437 for the second day, 24,595 for the third day, and 26,140 for the fourth day. The average number of complete conversations per day was 25,754 (standard deviation is about 1135).

The characteristics of these conversations are presented in the next chapter.

Chapter 5

The Measurement Results

This chapter presents a detailed analysis of wide area traffic on the Campus Ethernet. The characteristics of interest are described under five categories in Section 5.1 to Section 5.5. Section 5.1 presents an overview of the composition of the wide area IP traffic on this local area network using the IP trace. Later, detailed studies of wide area TCP conversations on the Campus Ethernet are described in Section 5.2 to Section 5.5 using file *CAMPUS-TCP-1*, which is the reduced data from the Campus TCP trace. Section 5.2 presents the breakdown of TCP traffic by subnets at the University of Saskatchewan and by types of application protocols. Section 5.3 provides information as to how wide the wide area conversations involving the Campus Ethernet can reach. Section 5.4 deals with the setup speed of wide area conversations on the Campus Ethernet. Section 5.5 analyzes the wide area conversations in several subcategories and deals with particular application protocols as well. Section 5.6 gives conclusions reached about the characteristics of wide area traffic on the Campus Ethernet.

Wide area traffic on an important subnet of the Campus Ethernet, the Computational Science Ethernet (the CS Ethernet), was also studied in detail in this research. Complete results of the wide area traffic on the CS Ethernet are presented in Appendix A. Section 5.7 gives a brief comparison of conversation characteristics of the Campus data and the CS data.

5.1 Transport Level Protocols

Recall that the preliminary trace captured contained all IP packets in a 24 hour period. The first step in the analysis is to break the traffic down based on the protocol type field in the IP packet header to determine the distribution of traffic by protocols. Table 5.1 shows the number of packets captured in this trace, categorized by protocols. TCP packets made up 96.01% of the wide area network traffic and UDP packets made up only 2.27%. Other protocols, such as ICMP, NTP (Network Time Protocol [Mills 92]), etc., made up 1.72% of the wide area network traffic. The results of detailed analysis of TCP applications (from file *CAMPUS-TCP-1* that was derived from the 4 day trace) are given in the following sections.

Table 5.1: Packet Counts from the IP Trace

Protocol		Number	Percent
TCP	(with S, F, or R flags set)	161,809	2.47%
	others	6,133,283	93.54%
	total	6,295,092	96.01%
UDP		148,793	2.27%
Others		112,798	1.72%
Total		6,556,683	100.00%

5.2 TCP Traffic Breakdown

People might ask how much wide area traffic is contributed by the various units in the University of Saskatchewan. Table 5.2 shows this breakdown. Each row of this table gives the name of the unit, the total number of wide area conversations involving this unit, the percentage of such conversations among the total wide area conversations, the total amount of data transmitted by this unit, and the percentage of the data among the total data transmitted by the Campus Ethernet. The units are listed in order of number of conversations.

Table 5.2: Breakdown of Complete TCP Conversations by Units

Campus Units	Number of Conv.	Percent	Total MBytes	Percent
Computing Services	70206	68.15	993.145	53.88
Computational Science	8888	8.63	333.973	18.12
Other Services	7077	6.87	71.567	3.88
Engineering	6262	6.08	159.503	8.65
Linac	2045	1.99	80.575	4.37
Commerce	1488	1.44	18.972	1.03
Physics	1398	1.36	38.402	2.08
Library	1194	1.16	13.412	0.73
Math	1147	1.11	15.568	0.84
Psychology	772	0.75	21.858	1.19
Geological Sciences	442	0.43	11.269	0.61
Crop Science	359	0.35	4.001	0.22
Mac/Ethernet Gateways	352	0.34	2.648	0.14
Chemistry	323	0.31	16.163	0.88
Computer Store	258	0.25	28.464	1.54
Health Sciences	213	0.21	4.199	0.23
Agriculture	213	0.21	4.590	0.25
Education	82	0.08	0.421	0.02
Campus Dialup Facility	78	0.08	0.926	0.05
Animal Science	63	0.06	10.984	0.60
Recruitment	48	0.05	6.675	0.36
Extension Division	41	0.04	1.557	0.08
Physical Education	28	0.03	3.453	0.19
Law	13	0.01	0.522	0.03
Research Services	10	0.01	0.264	0.01
Humanities	7	0.01	0.125	0.01
Administration	6	0.01	0.024	0.00
Development Office	3	0.00	0.003	0.00
Total	103016	100.0	1843.263	100.0

The Department of Computing Services has subdomains 128.233.1.x to 128.233.9.x. It contributed 68.15% of all the wide area conversations and accounted for 53.88% of all the wide area data to the Campus Ethernet over this period. The Department of Computational Science has subdomains 128.233.128.x to 128.233.136.x. It contributed 8.63% of all the wide area conversations and accounted for 18.12% of all the wide area data to the Campus Ethernet. The group listed as "other services" includes subdomains such as 192.139.75.x which belong to the University and are used for routing, etc. This group contributed 6.87% of all the wide area conversations and accounted for 3.88% of all the wide area data to the Campus Ethernet.

The CS Ethernet is a heterogeneous network of 178 interconnected hosts in the Department of Computational Science and represents an important subset of the campus user community. These users tend to have more Internet experience and thus might be considered more sophisticated for the purposes of this analysis. The usage patterns associated with this subnet were isolated for separate study. The subset of wide area TCP conversations was extracted from file *CAMPUS-TCP-1* and saved in a file called *CS-TCP-1*. A total of 8,888 complete TCP conversations make up this file with 1,985 on the first day, 2,226 on the second day, 2,313 on the third day, and 2,364 on the fourth day. The complete results for this subset of conversations are presented in Appendix A. The results presented in this chapter include these conversations as part of the campus total.

It is important to study multiplexing of different traffic types at the entrance to a wide area network. Since the application defines the type of traffic, network applications are identified by breaking down the TCP conversations according to the source or destination port number. A total of 419 application protocols were identified in file *CAMPUS-TCP-1*. Statistics on the 9 most common protocols (comprising more than 94% of the total) are given in Table 5.3. Each row gives a protocol followed by the total number of complete conversations of this application, the percentage of such conversations among all wide area conversations, the total size of these conversations, the average data rate, the average conversation size, and the average conversation duration. The total size of the conversations is the number of data bytes transferred by all conversations in the same protocol and is calculated by adding all individual conversation sizes. The average data rate is the ratio of this total size divided by the total duration of all these conversations (that is the sum of all individual conversation durations in the same protocol). The average conversation size is the ratio of total size divided by the number of all conversations in the same protocol. The average conversation duration is the ratio of total conversation duration divided by the number of individual conversations in the same protocol.

Table 5.3: Breakdown of Complete TCP Conversations by Protocol

Protocol	Number	Percent	Total MB	Data Rate (Bytes/sec)	Con. Size (KB/con)	Duration (Secs/con)
SMTP	43508	42.23	234.978	233.414	5.401	23.138
Gopher	18338	17.80	174.769	385.238	9.530	24.739
FTP-Data	14606	14.18	785.345	2594.704	53.769	20.722
WWW	5829	5.66	83.711	1092.362	14.361	13.147
Domain	5457	5.30	0.751	97.124	0.138	1.417
Telnet	4275	4.15	93.992	39.136	21.986	561.801
Finger	2495	2.42	1.825	88.676	0.732	8.250
FTP	2117	2.06	4.162	7.214	1.966	272.522
NNTP	1036	1.01	277.355	997.450	267.717	268.401
Others	5355	5.19	186.377	71.905	34.804	484.033
All	103016	100.00	1843.264	238.868	17.893	74.908

Table 5.3 shows that SMTP conversations represent the most popular protocol among the wide area conversations, making up 42.23% of the total number. FTP-Data, NNTP, SMTP, and Gopher conversations transmitted a large amount of data, comprising about 80% of all the data bytes transmitted by the wide area conversations during the four days. FTP-Data conversations themselves transmitted 42.61% of all the data bytes transmitted. These conversations had a high average data rate of 2.6 Kbytes/sec and the highest average conversation size of roughly 54 Kbytes.

Telnet, FTP, and NNTP are interactive applications. They have very long average conversation durations, which include network transmission time and user thinking and reaction time. SMTP, Gopher, FTP-Data, WWW, Domain, and Finger are bulk transfer applications. These conversations are characterized by short average durations, which include only network transmission time.

Although Gopher and WWW are relatively new applications on the Internet, Gopher conversations made up 17.8% and WWW conversations made up 5.66% of all the conversations observed.

The average wide area data rate of the Campus Ethernet is defined to be the ratio of the total size of all conversations divided by four days (345,600 seconds). It was 5.33 Kbytes/sec during the four days. Compared to the capacity of the 10 Mbits/sec link that is used to connect the Campus Ethernet to the Sask#net, this average data rate of complete wide area TCP conversations was relatively low during the study was carried out.

Table 5.4 shows the volumes of data that were transmitted between the Campus Ethernet and the Internet by various TCP protocols and the ratio between the outgoing and incoming volumes. A total of 1,220 Mbytes of data were sent from the Internet to the Campus Ethernet over the four days. This was almost double the volume of data that was sent from the Campus Ethernet to the Internet over the same period. Incoming SMTP data was almost 3 times the volume of outgoing SMTP data and outgoing Gopher data was more than 3 times the volume of incoming Gopher data. Clearly there is a non-symmetric relationship in volume of data transferred.

Table 5.4: Volumes of Data from the Campus Ethernet to the Internet and Vice Versa

Protocol	UofS to Internet (Mbytes, total)	Internet to UofS (Mbytes, total)	Ratio (Out:In)
SMTP	61.920	173.058	1:2.79
Gopher	133.886	40.883	3.27:1
FTP-Data	118.701	666.644	1:5.62
WWW	14.369	69.341	1:4.83
Domain	0.220	0.531	1:2.41
Telnet	60.741	33.251	1.83:1
Finger	0.754	1.072	1:1.42
FTP	1.342	2.819	1:2.10
NNTP	164.142	113.212	1.45:1
Others	67.038	119.339	1:1.78
All	623.113	1220.150	1:1.96

The breakdown by protocols of TCP conversations initiated by hosts on the Internet and by hosts at the University of Saskatchewan is shown in Table 5.5 and Table 5.6, respectively. Only the top 5 protocols are listed.

By comparing Table 5.3, Table 5.5, and Table 5.6, it is seen that 65% of the observed conversations were initiated by hosts on the Internet while 35% were initiated by hosts at the University of Saskatchewan. In other words, the number of conversations initiated by hosts on the Internet was almost double the number of conversations initiated by hosts at the University of Saskatchewan. Three protocols, SMTP, Gopher, and FTP-Data account for most of this difference.

Table 5.5: TCP Conversations Initiated by Hosts on the Internet

Protocol	Number	Percent	Total MB	Data Rate (Bytes/sec)	Con. Size (KB/con)	Duration (Secs/con)
SMTP	33674	50.29	167.166	197.334	4.964	25.157
Gopher	13046	19.48	134.354	419.368	10.298	24.557
FTP-Data	11530	17.22	656.826	2928.815	56.967	19.450
Telnet	2998	4.48	60.990	42.363	20.343	480.218
WWW	1635	2.44	12.399	318.163	7.583	23.834
Others	4082	6.10	194.323	283.468	47.605	167.938
All	66965	100.00	1226.057	344.791	18.309	53.101

Table 5.6: TCP Conversations Initiated by Hosts on the Campus Ethernet

Protocol	Number	Percent	Total MB	Data Rate (Bytes/sec)	Con. Size (KB/con)	Duration (Secs/con)
SMTP	9834	27.28	67.812	424.941	6.896	16.227
Domain	5457	15.14	0.751	97.124	0.138	1.417
Gopher	5292	14.68	40.415	303.206	7.637	25.188
WWW	4194	11.63	71.312	1893.398	17.003	8.980
FTP-Data	3076	8.53	128.519	1639.085	41.781	25.491
Others	8198	22.74	308.398	82.370	37.619	456.704
All	36051	100.00	617.207	148.341	17.120	115.412

5.3 How Wide do Wide Area TCP Conversations Reach?

An examination of the Internet domains involved in the conversations gives an overview of the different types of hosts with which the hosts at the University of Saskatchewan communicate. *TCPDUMP* can convert some numeric IP addresses to corresponding domain names by using the name servers. The **whois** database in FTP site *nic.ddn.mil* and a program called *hf* [Leres 93], which can convert raw Internet addresses to domain names, were used to map the remaining numeric IP addresses to the corresponding domain names. The information on domain names is used to determine the source and destination nations of information exchanged within each conversation. Even with these aids there were still some IP addresses which could not be mapped to a particular domain name. Conversations involving these addresses were excluded from the results reported in this section. These unresolved connections comprised 2.51% (2,581 conversations out of a total of 103,016) of all complete conversations. The deleted data include 457 unresolved hosts with 27 kinds of wide area conversations.

The reduced data contain 100,435 complete conversations, and was saved in a file named *CAMPUS-TCP-2*. The subset of these conversations contributed by the CS Ethernet was saved in file *CS-TCP-2*, in which there are 8,860 conversations.

Table 5.7 lists the counts of conversations, different subdomains, and different protocols that were involved in a given highest level domain. A total of 61 highest Internet domains were found in the reduced data. 53 countries (including Canada) out of the 81 countries that have access to the Internet or general electronic mail services were captured. 27.02% of all the conversations were with hosts within Canada, 44.23% were with hosts in the EDU domain, and 9.49% were with hosts in the COM domain.

Subdomains usually correspond to a single entity or local area network. For example, all hosts at the University of Saskatchewan are in the subdomain *usask.ca*. A total of 1,821 subdomains (local area networks) on the Internet were found in file *CAMPUS-TCP-2*. 387 different hosts at the University of Saskatchewan engaged in conversations with these 1,821 Internet local area networks. Table 5.8 lists the top 10 hosts at the University of Saskatchewan and Table 5.9 lists the top 10 local area networks on the Internet, with the counts of conversations and different Internet LANs involved with this host or different hosts at the University of Saskatchewan involved with this local area network.

Table 5.8 shows that host *skycat* participated in 18.90% of the wide area conversations and host *jester* participated in 12.77%. These are general purpose campus machines. *Skycat* is a VAXstation 4000-90 running VMS and *jester* is a DECstation 5000-125 running Ultrix, both in the Department of Computing Services. These machines are available for use by any faculty and staff member, as well as by graduate students whose advisors request an account for them. Undergraduate students may get class accounts if their professor requests it. *Skdad* (the file server of the CS Ethernet) participated in 2.77% of the wide area conversations.

Table 5.7: Number of Conversations as a Function of Domains

Domain	Description	No. of Conversations	Percent of Total	No. of Subdomains	No. of Protocols
AR	Argentina	1	0.00	1	1
AT	Austria	157	0.16	3	9
AU	Australia	985	0.98	6	64
BE	Belgium	188	0.19	2	6
BR	Brazil	64	0.06	14	4
CA	Canada	27140	27.02	134	70
CH	Switzerland	1010	1.01	15	7
CL	Chile	67	0.07	6	7
CN	China	11	0.01	1	1
CR	Costa Rica	2	0.00	1	2
CZ	Czech Republic	27	0.03	8	5
DE	Germany	714	0.71	60	15
DK	Denmark	125	0.12	10	6
EC	Ecuador	8	0.01	1	2
EE	Estonia	7	0.01	1	2
EG	Egypt	5	0.00	1	2
ES	Spain	43	0.04	13	7
FI	Finland	519	0.52	21	9
FR	France	330	0.33	41	13
GB	Great Britain	5	0.00	1	1
GR	Greece	9	0.01	4	3
HK	Hong Kong	93	0.09	7	7
HR	Croatia	181	0.18	6	11
HU	Hungary	82	0.08	7	9
IE	Ireland	14	0.01	4	4
IL	Israel	189	0.19	4	8
IN	India	51	0.05	2	5
IS	Iceland	4	0.00	3	3
IT	Italy	412	0.41	27	12
JP	Japan	574	0.57	7	14
KR	Korea (South)	27	0.03	3	6

abcdef

Table 5.8: Ten Most Popular Hosts on the Campus Ethernet

Hosts	Number of Conversations	Percent of Total	Number of Internet LANs
Skycat	18981	18.90	677
Jester	12824	12.77	815
Duke	11040	10.99	354
Access	7378	7.35	545
Herald	5148	5.13	468
Sask	4251	4.23	243
Skynet	3630	3.61	317
Dvinci	3276	3.26	244
Doc.siast.sk.ca	2962	2.95	101
Skdad	2781	2.77	271

Conversations with a small number of local area networks (LANs) dominated as can be seen in Figure 5.1. Only 12% of external LANs were responsible for roughly 80% of the conversations observed over the four-day period. In fact, only 10 LANs (out of a total of 1,821) were responsible for 20.88% of all the conversations. This "top ten" list is given in Table 5.9. The local area network at TR*Labs* headed the list, participating in 3.73% of the wide area conversations with the Campus Ethernet. The conversations with this network involved 35 different hosts at the University.

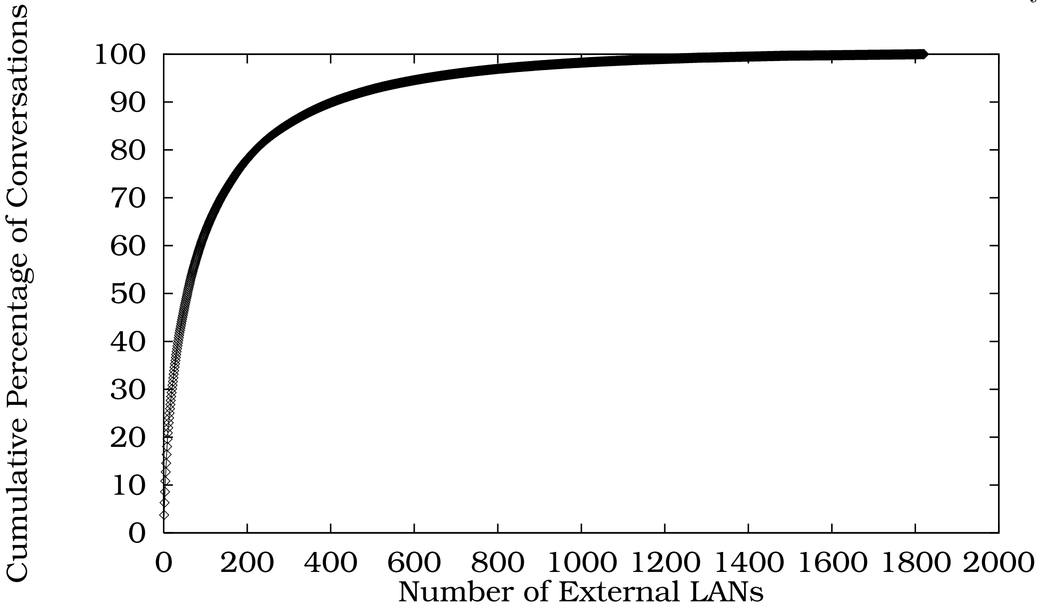

Figure 5.1: Cumulative Percentage of Number of Conversations with External LANs

There are many popular FTP, Gopher, NNTP, and WWW servers on the Internet. Table 5.10, Table 5.11, Table 5.12, and Table 5.13 show the 5 most popular FTP, Gopher, WWW, and NNTP servers that were used by hosts at the University of Saskatchewan, respectively. Each row of these tables consists of the name of the server, the total number of FTP, Gopher, WWW, or NNTP conversations with the server, and the percentage of these conversations among the total FTP, Gopher, WWW, or NNTP conversations, respectively.

Table 5.9: Ten Most Popular Local Area Networks on the Internet

LANs	Number of Conversations	Percent of Total	Number of UofS hosts
Trlabs.ca	3742	3.73	35
Sk.ca	2609	2.60	18
Umn.edu	2274	2.26	74
Uregina.ca	2230	2.22	40
Ualberta.ca	1919	1.91	56
Ac.uk	1857	1.85	58
Buffalo.edu	1829	1.82	21
Uiuc.edu	1634	1.63	83
Uga.edu	1573	1.57	20
Mit.edu	1303	1.30	40

Table 5.10: Five Most Popular FTP Servers on the Internet

Hosts	Num. of Con.	Percent of Total
Winftp.cica.indiana.edu	76	3.76
Wuarchive.wustl.edu	56	2.77
Oak.oakland.edu	46	2.28
Gowinnt.microsoft.com	38	1.88
Ftp.usr.com	30	1.49

In total, 361 FTP, 467 Gopher, 227 WWW, and 8 NNTP servers were observed in file *CAMPUS-TCP-2*. Again, no group of local area networks dominated. The Microsoft windows FTP server at the Center for Innovative Computer Applications (CICA), Indiana University (*winftp.cica.indiana.edu*), is an anonymous FTP clearinghouse for MicroSoft Windows applications, tips, utilities, drivers, bitmaps, etc. The FTP server at Washington University in Saint Louis (*wuarchive.wustl.edu*) has a large

Table 5.11: Five Most Popular Gopher Servers on the Internet

Hosts	Num. of Con.	Percent of Total
Wired.com	208	1.19
Beech.cic.net	130	0.74
Wiretap.spies.com	129	0.74
Ashpool.micro.umn.edu	126	0.72
Owl.nstn.ns.ca	121	0.69

Table 5.12: Five Most Popular WWW Servers on the Internet

Hosts	Num. of Con.	Percent of Total
Sils.umich.edu	317	5.58
Chico.ncsa.uiuc.edu	306	5.39
Vanbc.wimsey.com	300	5.28
Harpo.ncsa.uiuc.edu	195	3.43
Thrall.cm.cf.ac.uk	137	2.41

collection of publicly available software for PCs, Macs, Amigas, and UNIX systems. There is also a large collection of images at this site. It also contains some Internet documentation, such as RFCs, statistics on NFS traffic, monthly usage statistics for wuarchive in terms of FTP traffic, etc. The Gopher server *wired.com* provides complete information about the computer magazine *WIRED*. Detailed information about the CICNet is available at the Gopher server *beech.cic.net*. The WWW server *harpo.ncsa.uiuc.edu* presents all the new information provided by the National Center for Supercomputing Applications (NCSA). The WWW server *vanbc.wimsey.com* is Canada's premier WWW site and British Columbia's leading Internet service provider. It is located in Vancouver, and provides information through FTP, Gopher, WAIS, and WWW.

Table 5.13: Five Most Popular NNTP Servers on the Internet

Hosts	Num. of Con.	Percent of Total
Softy.softwords.bc.ca	252	24.32
Canopus.cc.umanitoba.ca	58	5.60
Quartz.ucs.ualberta.ca	45	4.34
Sue.cc.uregina.ca	16	1.54
News.nofc.forestry.ca	12	1.16

5.4 Conversation Setup Latency

Conversation setup latency can be defined to be the latency or delay that a conversation experiences before it can transmit user data [Schmidt 92] and can be calculated by measuring the difference between the time the first TCP SYN packet arrives from the originator and the time a TCP SYN packet arrives from the responder to the monitor. Since the difference between the time it takes for a packet to travel from a host on the Campus Ethernet to the monitor (which was on the Campus Ethernet) and the time it takes for a packet to travel from a host on other sites on the Internet to the monitor is significant, setup latency of conversations that were initiated by hosts on the Campus Ethernet and setup latency of conversations that were initiated by hosts on the Internet will be analyzed separately.

An analysis of the file *CAMPUS-TCP-1* revealed that the average setup latency for the conversations that were initiated by hosts at the University of Saskatchewan was 0.85 seconds, with a maximum of 114.43 seconds and a minimum of 0.02 seconds, and the average setup latency for the conversations initiated by hosts on the Internet was 0.1 seconds, with a maximum of 291.54 seconds and a minimum of 0.15 msec. Figure 5.2 and Figure 5.3 show the cumulative percentage of conversations that were initiated by hosts on the Campus Ethernet and by hosts on the Internet, respectively, whose setup time is in a particular range.

From the graphs it can be seen that about 86% of the TCP conversations initiated at the University of Saskatchewan had a setup latency shorter than 0.5 seconds and about 90% of the TCP conversations initiated on the Internet had a setup latency shorter than 10 msec. 0.29% of the TCP conversations initiated at the University of Saskatchewan had a setup latency longer than 10 seconds.

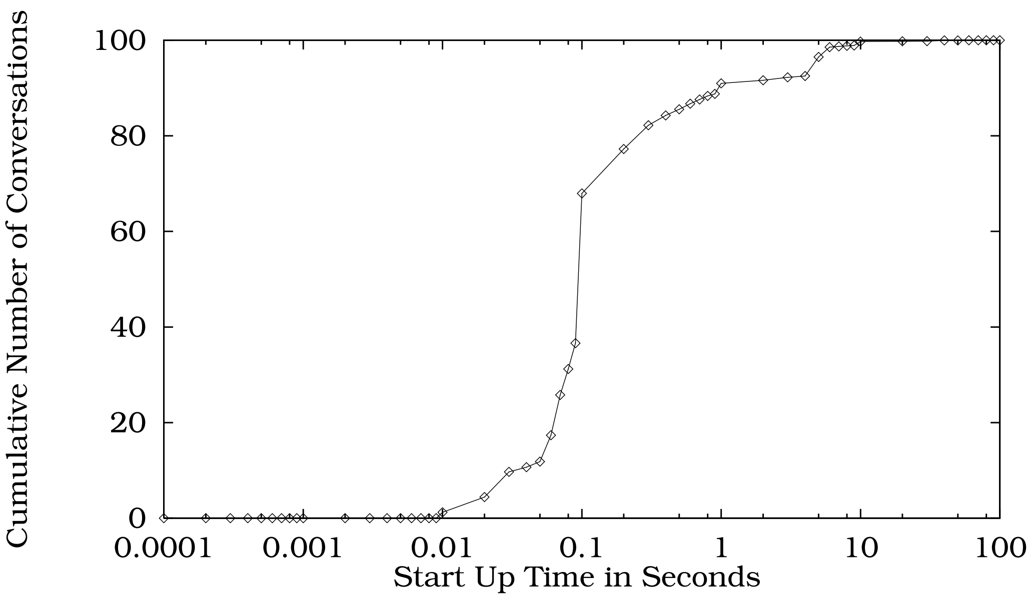

Figure 5.2: Setup Latency of Conversations (UofS-initiated), All Protocols

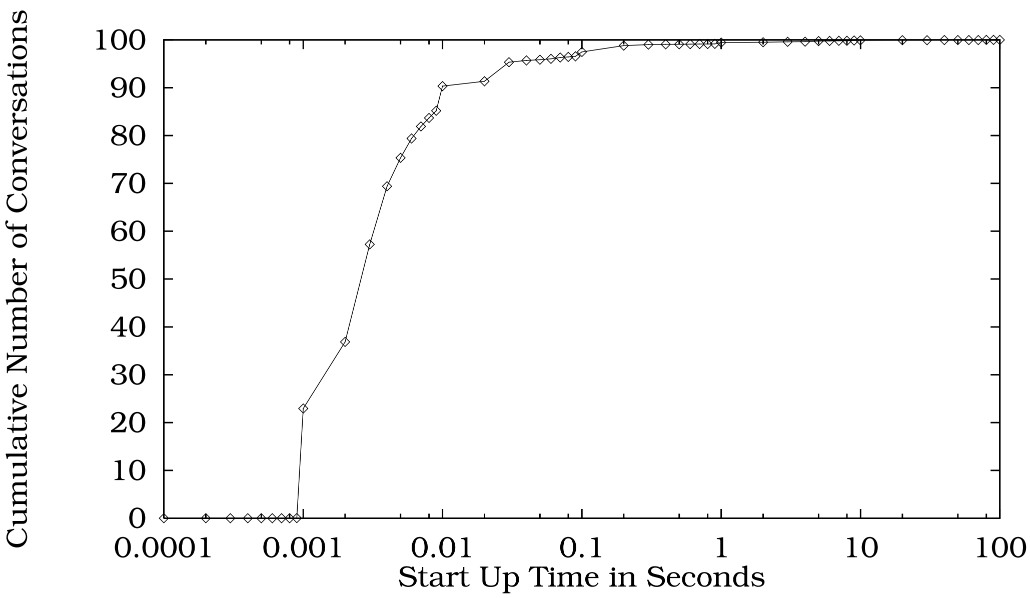

Figure 5.3: Setup Latency of Conversations (Internet-initiated), All Protocols

5.5 Analysis of TCP Conversations

This section presents per-conversation statistics. These include statistics on conversation size, conversation duration, the relationship between size and duration, conversation bandwidth, interarrival time of conversations, number of concurrent conversations as defined in Section 2.8, and direction of data flow in conversations. An analysis of this data can provide insight into network performance and the characteristics of conversation types that the current Internet is exposed to.

First, an overview of the basic properties of these conversations is presented. Next, histograms of the distributions of the quantities of these properties are presented to give a first impression of the variability of the data. Finally, statistics are presented to show how conversation activity varies over the course of four days.

Section 5.5.1 presents per-conversation statistics in terms of all complete TCP conversations. Section 5.5.2 to Section 5.5.9 deal with particular widely used application protocols. Section 5.5.10 gives some ideas about direction of data flow of these application protocols.

Some analyses in this section follow the style used in [Paxson 91, Schmidt 92]. Where appropriate, results are compared with those presented in [Paxson 91] and [Wakeman 93, Schmidt 92].

5.5.1 All TCP Traffic

A statistical analysis of the TCP traffic from all protocols is given in Table 5.14. Figure 5.4 gives histograms of the number of conversations of all protocols whose size, duration, bandwidth, and interarrival time lie in a particular range. Since the difference between the maximum and the minimum of the conversation size, duration, bandwidth, and interarrival time is large, the x-axis uses a logarithmic scale (base 2).

Table 5.14: Statistics of TCP Conversations, All Protocols

	Max	Min	Average	Cov	% smaller than average
Conversation size(MB)	20.849	0.000	0.018	9.987	90.32
Duration(sec)	55114.195	0.042	74.908	8.983	91.13
Bandwidth(KB/sec)	203.788	0.000	0.951	2.462	77.59
Interarrival time(sec)	108.902	0.000	3.354	1.395	68.41

The size of a conversation refers to the amount of data transferred. 47,729 conversations out of 103,016 conversations were in the range of 1 Kbyte to 4 Kbytes

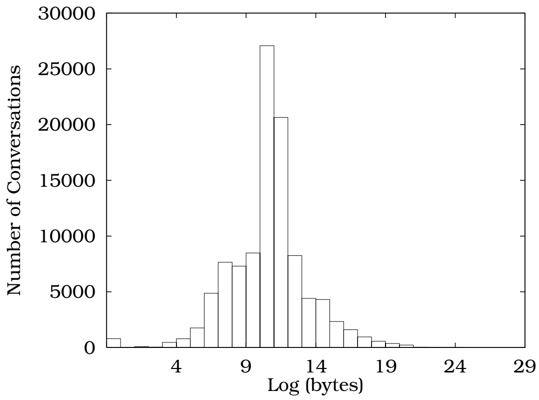

(a) Conversation Size

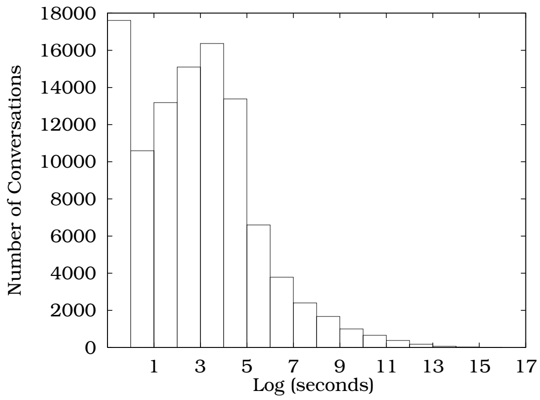

(b) Conversation Duration

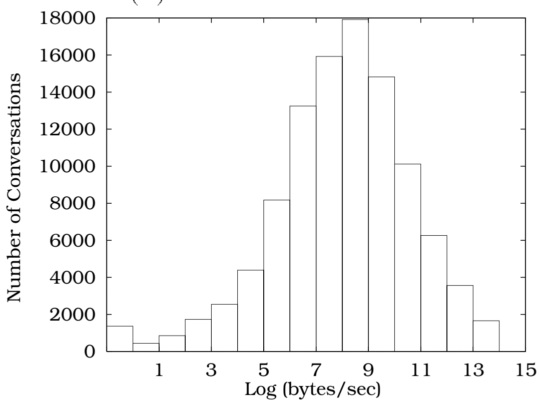

(c) Conversation Bandwidth

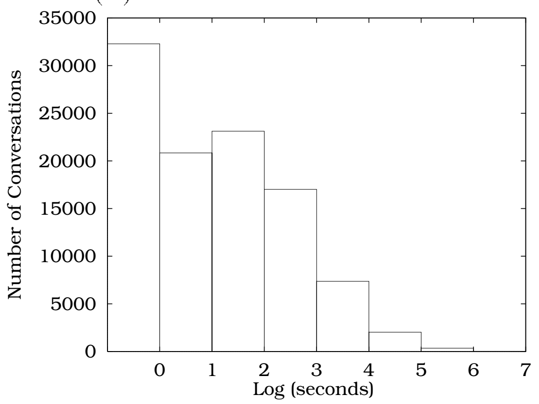

(d) Conversation Interarrival Time

Figure 5.4: Characteristics of TCP Conversations

in size. 90.32% of the conversation sizes were smaller than the average size, which was 18 Kbytes. This shows that some conversations were very big, with the biggest being 21 MBytes. Further study of the trace shows that the biggest conversation resulted from one NNTP conversation initiated by host *tribune* at the University of Saskatchewan to news server *sue.cc.uregina.ca*. The conversation lasted 2.74 hours with data rate 2.1 Kbytes/sec. The average conversation size of 18 Kbytes on the University of Saskatchewan Campus Ethernet is smaller than the 25 Kbytes reported by Paxson [Paxson 91].

Although most conversations were reasonably short-lived (83.7% of conversations were shorter than 32 seconds) there are some very long conversations. The longest conversation observed lasted about 15 hours. It was a Telnet conversation from host *access* at the University of Saskatchewan to host *bee08* in the *uspnet.br* domain. The average conversation duration was around 75 seconds, shorter than the average conversation duration of 113 seconds reported by Wakeman *et al.* [Wakeman 93] and the 209 seconds reported by Paxson [Paxson 91].

The maximum conversation bandwidth observed was 1.6 Mbits/sec. It was a SMTP conversation from host *skdad* on the Campus Ethernet to host *george* in *sedsystems.ca* domain. It transmitted 13.26 Mbytes in 65 seconds. About 60% of all conversations had data rates between 64 bytes/sec and 1K bytes/sec. The average conversation bandwidth was about 951 bytes/sec which is much higher than the 127 bytes/sec reported in Paxson's study [Paxson 91].

The minimum conversation interarrival time was 88 μsec and the maximum was 108.9 seconds. 32,287 conversations out of 103,016 had interarrival times shorter than 1 second. 68.41% of the conversation interarrival times were shorter than the average, which was about 3.35 seconds.

Figure 5.5 shows a histogram of the number of conversations that were active over 15 minute intervals during the four observation days. This information tells whether or not the network is busy at particular times of day. In general, the higher the number of conversations active in the same length of time, the busier the network is. The bandwidth of the conversations is clearly a factor as well. It can be seen that normally the network is relatively very busy from 8 AM to 6 PM and from 8 PM to 10 PM. There is a drop during lunch time from noon to 1 PM. This pattern is consistent with the regular working hours of people at universities given the fact that some students work in the evenings.

It has been shown that some applications open TCP conversations for relatively long periods of time. It is interesting to study the correlation between the conversation duration and the amount of data transmitted. Figure 5.6 shows the relationship between conversation sizes and conversation durations. Although the conversation size does not necessarily increase as the conversation duration increases, there is a tendency in that direction. In general, long conversation duration does not necessary create a great deal of traffic. In fact, the maximum conversation size was not associated with the longest conversation duration. The same conclusion was reached by Schmidt *et al.* [Schmidt 92].

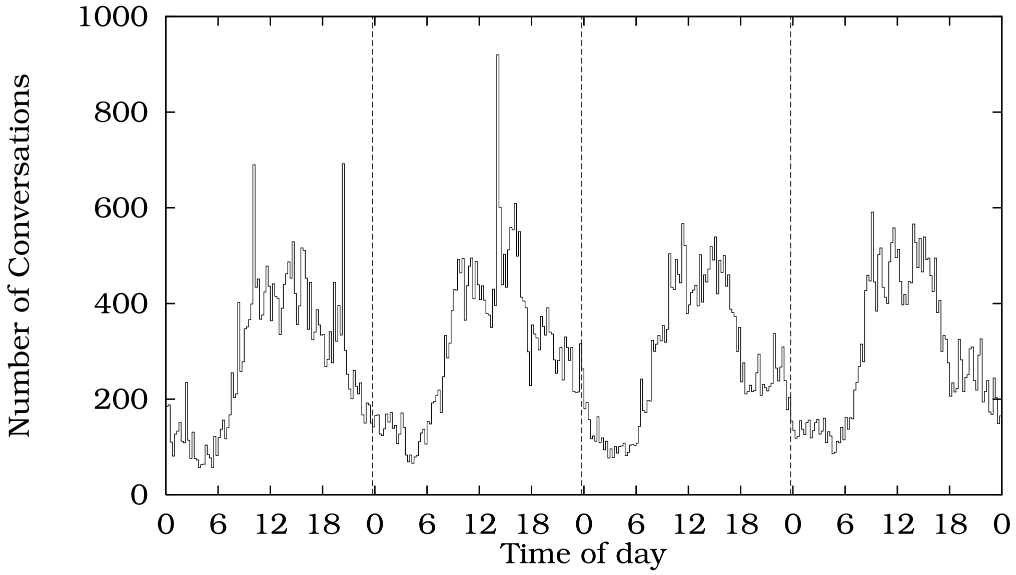

Figure 5.5: Number of Active Conversations, All Protocols

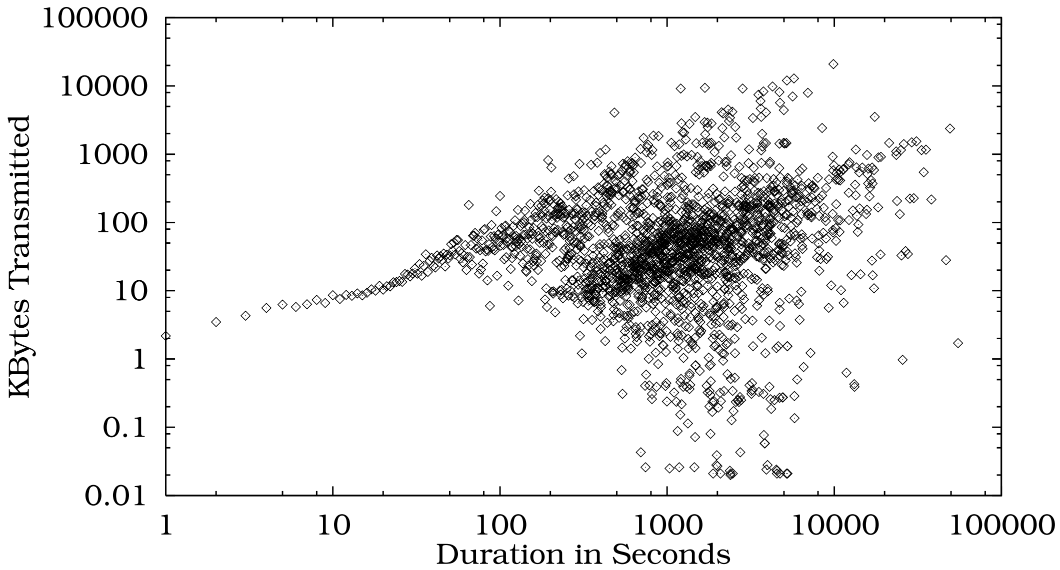

Figure 5.6: Conversation Size vs Conversation Duration, All Protocols

5.5.2 File Transfer Protocol Traffic (FTP and FTP-Data)

FTP is used for sending files between hosts. Each FTP session initiates one FTP (FTP-Control) and zero or more FTP-Data connections (see Section 2.7.2). A statistical analysis of FTP (FTP-Control) traffic is given in Table 5.15. Figure 5.7 shows histograms of the number of FTP conversations whose size, duration, bandwidth, and interarrival time lie in a particular range. Again, the x-axis uses a logarithmic scale (base 2).

Table 5.15: Statistics of FTP (FTP-Control) Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(MB)	0.083	0.000	0.002	1.911	66.65
Duration(sec)	25689.795	0.199	272.522	2.865	76.15
Bandwidth(KB/sec)	1.325	0.000	0.036	3.161	83.99
Interarrival time(sec)	3045.634	0.004	163.182	1.471	69.91

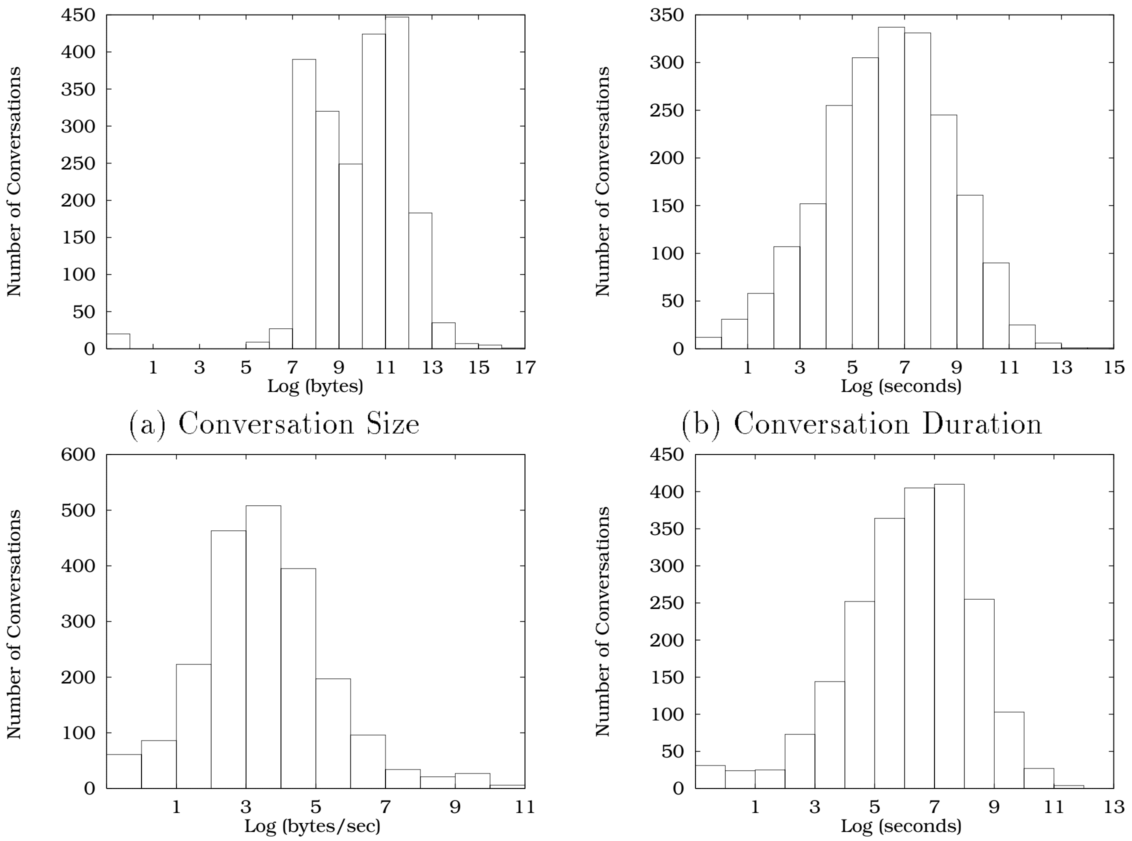

(a) Conversation Size (b) Conversation Duration

(c) Conversation Bandwidth (d) Conversation Interarrival Time

Figure 5.7: Characteristics of FTP (FTP-Control) Conversations

About 18% of the FTP conversations were from 128 to 256 bytes in size and

41% were from 1K to 4K bytes in size. The average FTP conversation size was 2 Kbytes. This is similar to the average of 2.4 Kbytes reported by Wakeman *et al.* [Wakeman 93] but is higher than the 1.3 Kbytes reported by Paxson [Paxson 91]. About 46% of the FTP conversations were from 32 seconds to 256 seconds in duration. The average duration of a FTP conversation was about 273 seconds, which is shorter than Paxson's 420 seconds and Wakeman *et al.*'s 343 seconds. The average bandwidth was about 36 bytes/sec, which is much larger than Paxson's 3 bytes/sec. About 70% of the interarrival times of the FTP conversations were shorter than 163 seconds, the average interarrival time.

The actual file transfer of a FTP session is done using a separate FTP-Data connection and the FTP site always sends the FTP-Data connection request. In this study, the average number of FTP-Data connections within one FTP (FTP-Control) connection was 7. The average quantity of data transmitted by a FTP session (FTP and FTP-Data) was 373 Kbytes.

A statistical analysis of the FTP-Data traffic is given in Table 5.16. Figure 5.8 shows histograms of the number of FTP-Data conversations whose size, duration, bandwidth, and interarrival time lie in a particular range. The x-axis still uses a logarithmic scale.

Table 5.16: Statistics of FTP-Data Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(MB)	10.737	0.000	0.054	5.355	86.92
Duration(sec)	5183.157	0.042	20.722	6.004	87.60
Bandwidth(KB/sec)	106.431	0.000	2.443	1.662	70.90
Interarrival time(sec)	3923.103	0.001	23.656	3.670	79.91

FTP-Data conversations transmitted a large amount of data on the network with an average conversation size of 54 Kbytes. Nevertheless the data from file *CAMPUS-TCP-1* shows that about 70% of FTP-Data conversations transmitted less than 4 Kbytes of data. This is because many FTP sites restrict the size of files that can be sent electronically. As a result, some FTP sites break their huge files into several smaller ones. 1.13% of the FTP-Data conversations sent zero bytes of data. No data

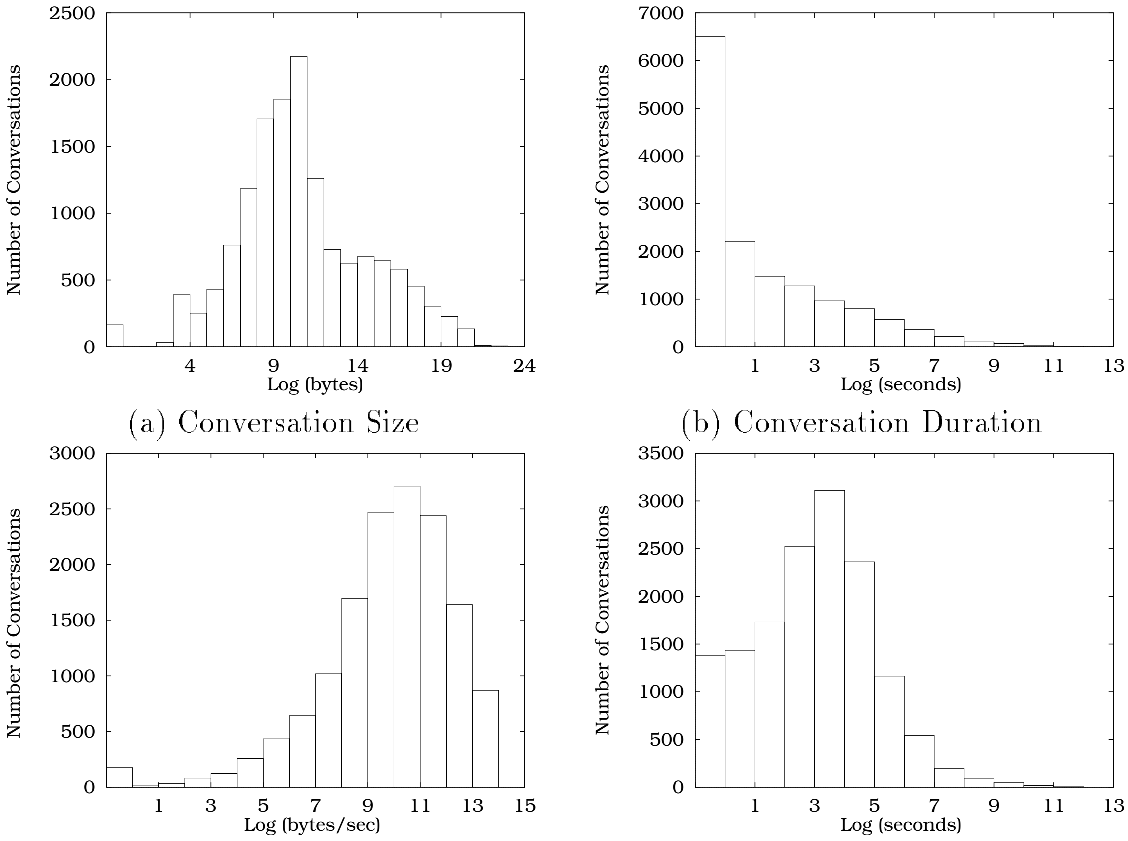

(a) Conversation Size
(b) Conversation Duration
(c) Conversation Bandwidth
(d) Conversation Interarrival Time

Figure 5.8: Characteristics of FTP-Data Conversations

are transmitted in the FTP-Data connections caused by operations such as *dir* to list the file names and directories on the FTP site. 45% of the FTP-Data conversations were shorter than 1 second in duration. The average duration was 21 seconds, longer than Paxson's 17 seconds. The average bandwidth was 2.4 Kbytes/sec, lower than Paxson's 4.7 Kbytes/sec. The minimum interarrival time was 1 msec. 80% of the FTP-Data conversations' interarrival times were shorter than the average interarrival time, which was 24 seconds.

Figure 5.9 and Figure 5.10 show histograms of the number of active FTP and FTP-Data conversations during the four observation days. The maximum number of active FTP conversations in any 15 minute interval was 30 and the maximum number of active FTP-Data conversations was 547. At the period when the number of active FTP-Data conversations reached the maximum, the number of active FTP conversations was 12.

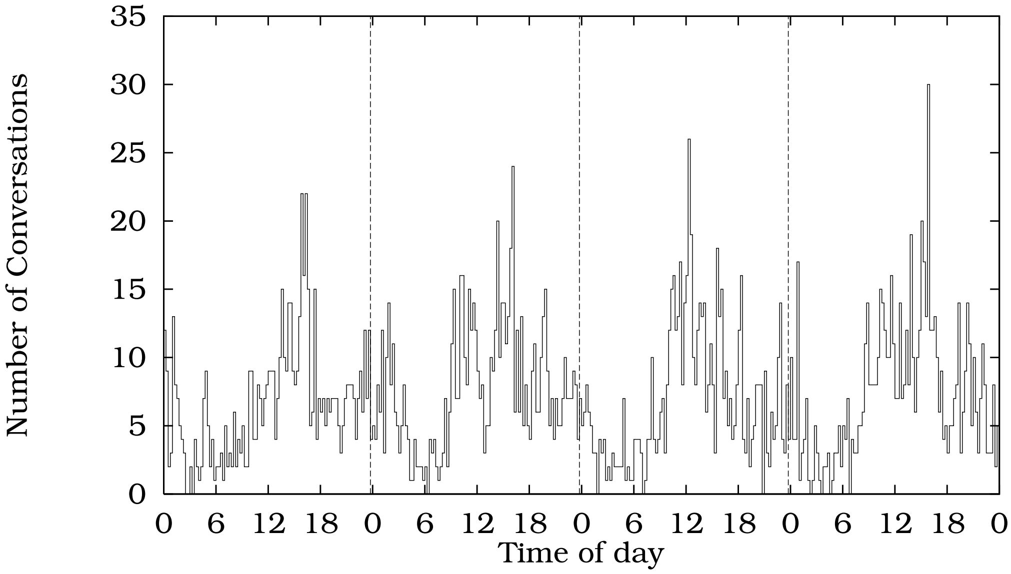

Figure 5.9: Number of Active FTP (FTP-Control) Conversations

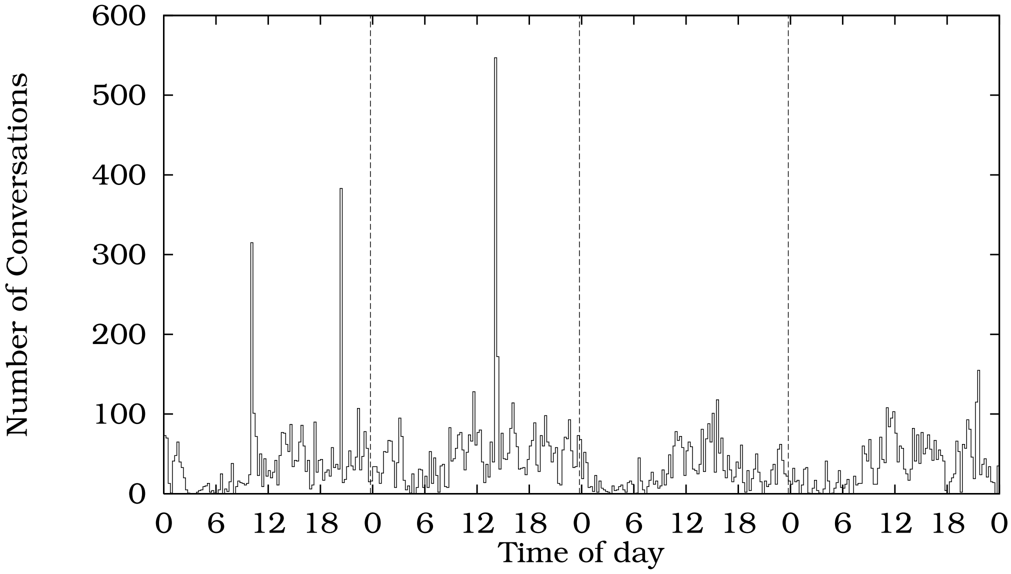

Figure 5.10: Number of Active FTP-Data Conversations

5.5.3 SMTP Traffic

SMTP is used for sending electronic mail. Table 5.17 gives some statistics of SMTP conversations. Figure 5.11 shows similar histograms of the number of SMTP conversations whose size, duration, bandwidth, and interarrival time lie in a particular range.

Table 5.17: Statistics of SMTP Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(MB)	13.258	0.000	0.005	15.165	86.10
Duration(sec)	2380.327	0.045	23.138	2.078	75.00
Bandwidth(KB/sec)	203.788	0.000	0.445	3.321	76.94
Interarrival time(sec)	246.722	0.000	7.942	1.418	68.86

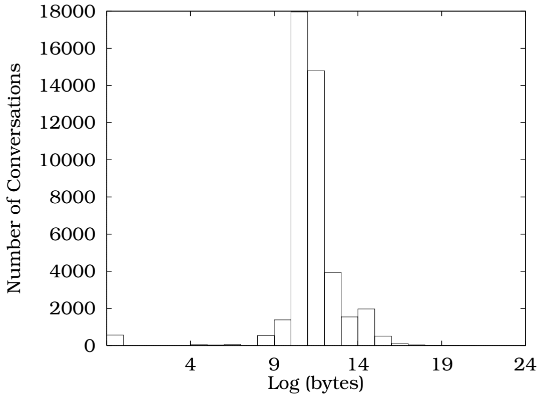

(a) Conversation Size

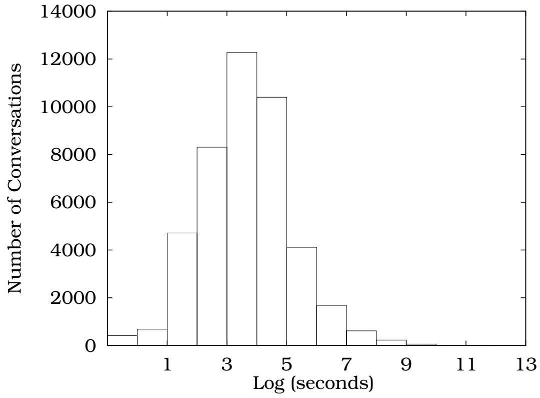

(b) Conversation Duration

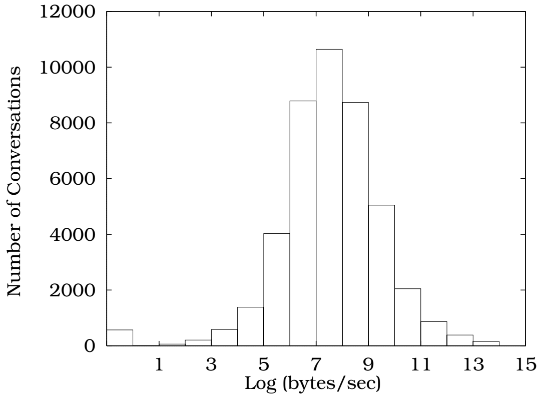

(c) Conversation Bandwidth

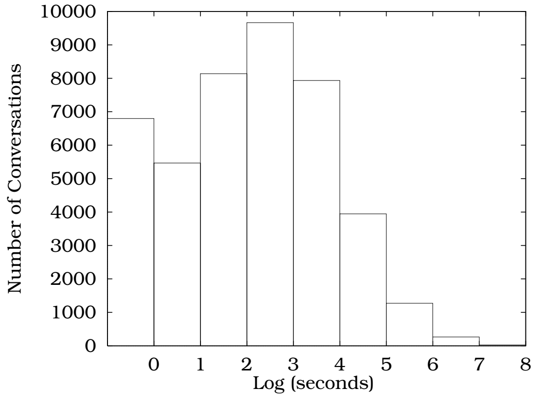

(d) Conversation Interarrival Time

Figure 5.11: Characteristics of SMTP Traffic

75% of SMTP conversations were from 1K to 4K bytes in size. The average conversation size was 5 Kbytes which is close to Paxson's 3.4 Kbytes. The average conversation duration was 23 seconds which is higher than Paxson's 13 seconds but shorter than Wakeman *et al.*'s 43 seconds. About 65% of the conversation bandwidths were in the range from 64 to 512 bytes/sec. The average conversation bandwidth was 445 bytes/sec, which is much higher than Paxson's 268 bytes/sec. The minimum interarrival time was 164 μsec. 69% of all the interarrival times of the SMTP conversations were shorter than the average interarrival time, which was 8 seconds.

The distribution of the number of active SMTP conversations is shown in Figure 5.12. The number of active conversations was larger than zero at almost all periods. It peaked during the regular working hours. Interestingly, there were still many SMTP conversations during "off hours". Many of these might be generated by international hosts in different timezones.

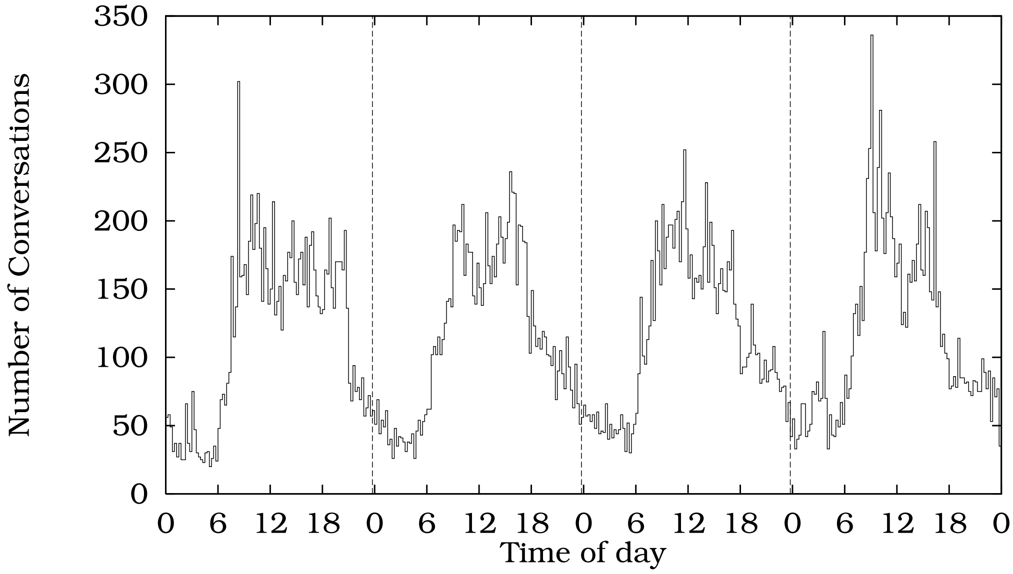

Figure 5.12: Number of Active SMTP Conversations

5.5.4 Telnet Traffic

Telnet is used for remote login. Table 5.18 summarizes some statistics of Telnet traffic. Figure 5.13 shows the same sort of histograms of the number of Telnet conversations whose size, duration, bandwidth, and interarrival time lie in a particular range.

Table 5.18: Statistics of Telnet Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(MB)	3.524	0.000	0.022	3.898	79.27
Duration(sec)	55114.195	0.127	561.801	3.172	79.16
Bandwidth(KB/sec)	1.412	0.000	0.044	1.200	61.73
Interarrival time(sec)	2673.305	0.009	80.646	1.644	70.71

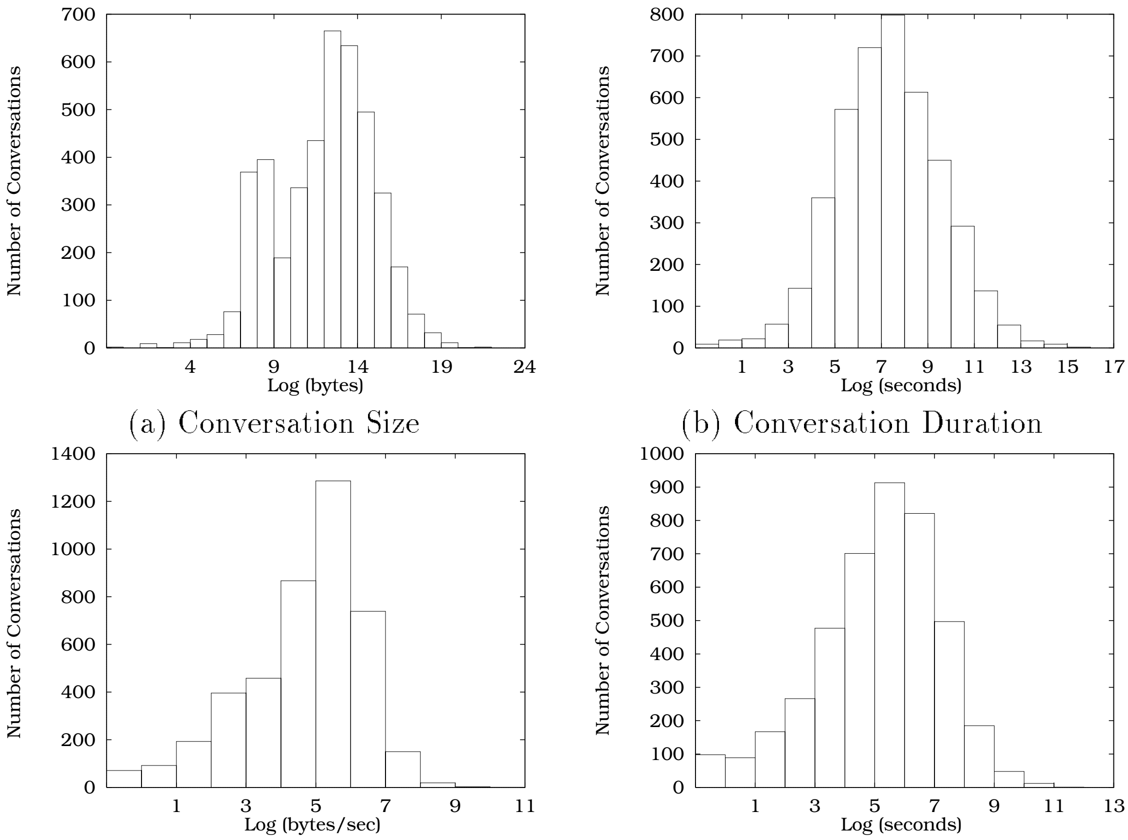

(a) Conversation Size (b) Conversation Duration

(c) Conversation Bandwidth (d) Conversation Interarrival Time

Figure 5.13: Characteristics of Telnet Traffic

18% of Telnet conversations were from 128 to 512 bytes and 30% were from 4K to 16K bytes in size. The average conversation size was 22 Kbytes, which is smaller than Paxson's 31.7 Kbytes. 63% of the conversations were from 32 seconds to 512 seconds in duration. The average conversation duration was 562 seconds, which is much shorter than Paxson's 1610 seconds. 30% of the conversations had a data rate between 32 bytes/sec and 64 bytes/sec. The average conversation bandwidth was 44 bytes/sec, which is higher than Paxson's 20 bytes/sec. 70.7% of the interarrival times of Telnet conversations were shorter than 81 seconds, the average interarrival time.

Figure 5.14 shows the distribution of number of active Telnet conversations. It can be seen that there were many Telnet conversations from 8 AM to 6 PM. The maximum number of active Telnet conversations was 103.

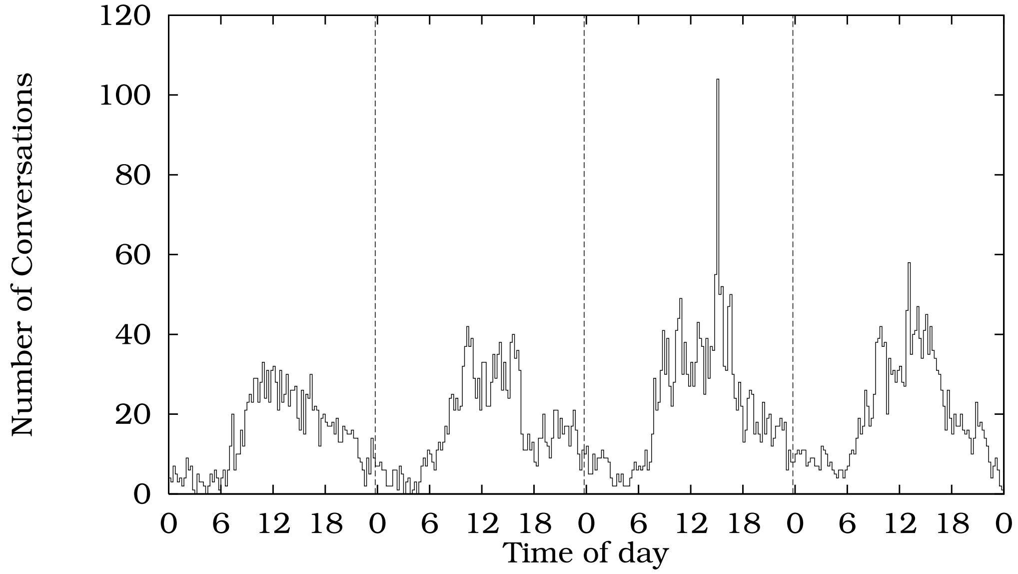

Figure 5.14: Number of Active Telnet Conversations

5.5.5 Gopher Traffic

Gopher is used for distributed document search and retrieval. A statistical analysis of the Gopher traffic is presented in Table 5.19. Figure 5.15 plots the similar histograms of the number of Gopher conversations whose size, duration, bandwidth, and interarrival time lie in a particular range.

60% of Gopher conversations were from 256 to 2K bytes in size. The average conversation size was 10 Kbytes. 81% of Gopher conversations were shorter than 8 seconds in duration. The average conversation duration was 25 seconds and the average conversation bandwidth was 1.2 Kbytes/sec. The minimum interarrival time was 196 μsec and the average interarrival time was 19 seconds with 77% of the interarrival times of the Gopher conversations being shorter than the average.

Table 5.19: Statistics of Gopher Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(MB)	5.708	0.000	0.010	6.586	88.47
Duration(sec)	13218.567	0.165	24.739	9.437	92.47
Bandwidth(KB/sec)	51.117	0.000	1.219	2.009	74.97
Interarrival time(sec)	1976.582	0.000	18.838	2.610	76.90

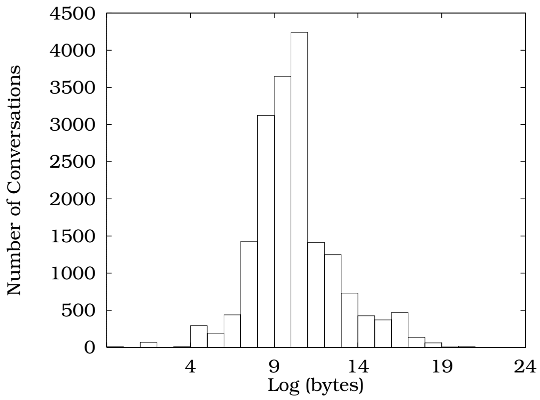

(a) Conversation Size

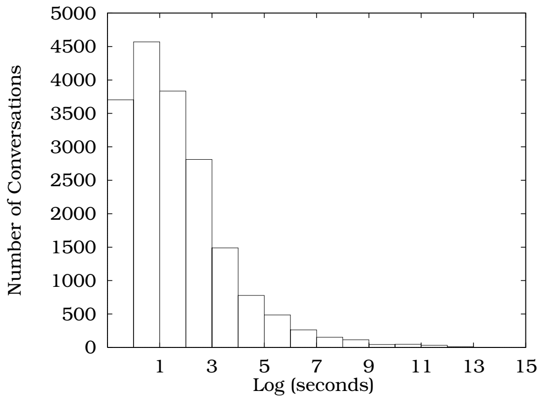

(b) Conversation Duration

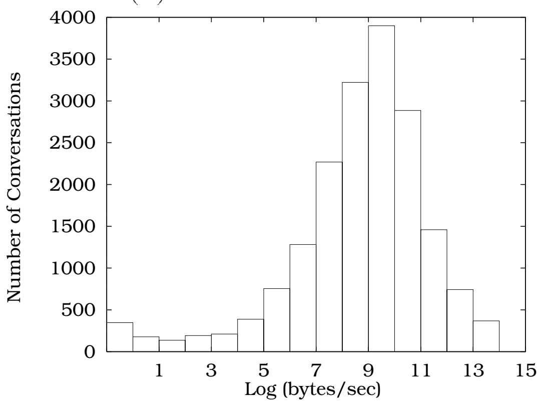

(c) Conversation Bandwidth

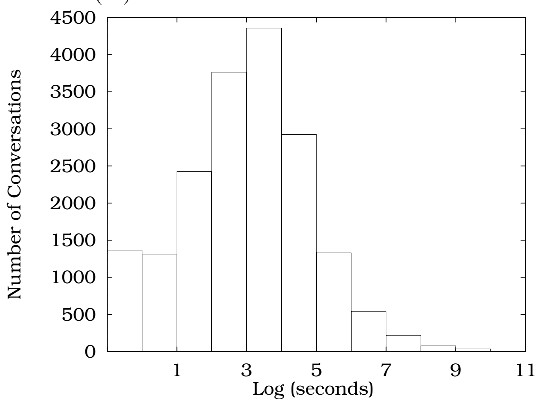

(d) Conversation Interarrival Time

Figure 5.15: Characteristics of Gopher Traffic

Figure 5.16 shows number of active Gopher conversations as a function of time of day. Since each service request of the client uses one Gopher conversation and the Gopher server terminates the connection after sending the response, the number of active Gopher conversations increases as the number of requests from clients increase within a time period.

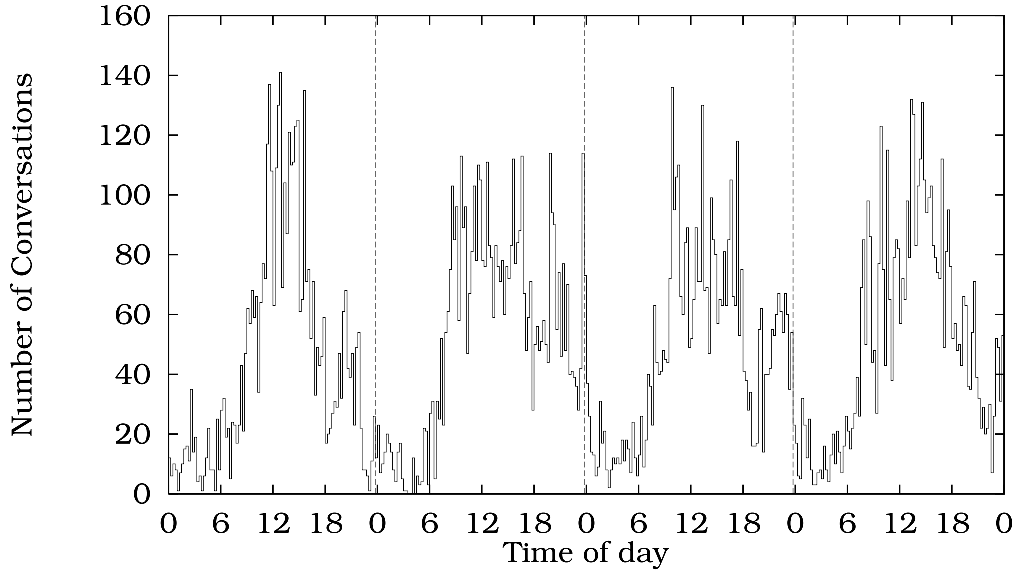

Figure 5.16: Number of Active Gopher Conversations

5.5.6 Finger Traffic

Finger is used for the exchange of user information. A statistical analysis of Finger traffic is shown in Table 5.20. Figure 5.17 shows histograms of the number of Finger conversations whose size, duration, bandwidth, and interarrival time lie in a particular range. As usual, the x-axis uses a logarithmic scale.

Table 5.20: Statistics of Finger Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(MB)	0.070	0.000	0.001	2.709	78.52
Duration(sec)	2024.645	0.128	8.250	5.932	82.40
Bandwidth(KB/sec)	12.975	0.000	0.361	2.729	77.27
Interarrival time(sec)	13085.709	0.001	138.421	2.841	76.31

Sometimes users use Finger to check out whether a friend or colleague is currently logged in, while sometimes they want to see a list of all users presently logged into a remote site. The size of messages sent by Finger conversations depends on each site and what the user wants to see. 58% of Finger conversations were from 128 to 512 bytes in size. The average conversation size was 1 Kbyte, which doubles Paxson's 0.5 Kbytes. 669 out of the 2,495 Finger conversations were from 1 second to 2 seconds in duration. The average conversation duration was 8 seconds, which is shorter than Paxson's 34 seconds. The average conversation bandwidth was 361 bytes/sec, which is much larger than Paxson's 16 bytes/sec. 76% of all the interarrival times of the Finger conversations were shorter than the average interarrival time, which was 138 seconds.

Figure 5.18 shows the distribution of the number of active Finger conversations for the four days. Obviously, there were many Finger conversations during the regular working hours. The maximum number of active Finger conversations observed over any 15 minute period was 66.

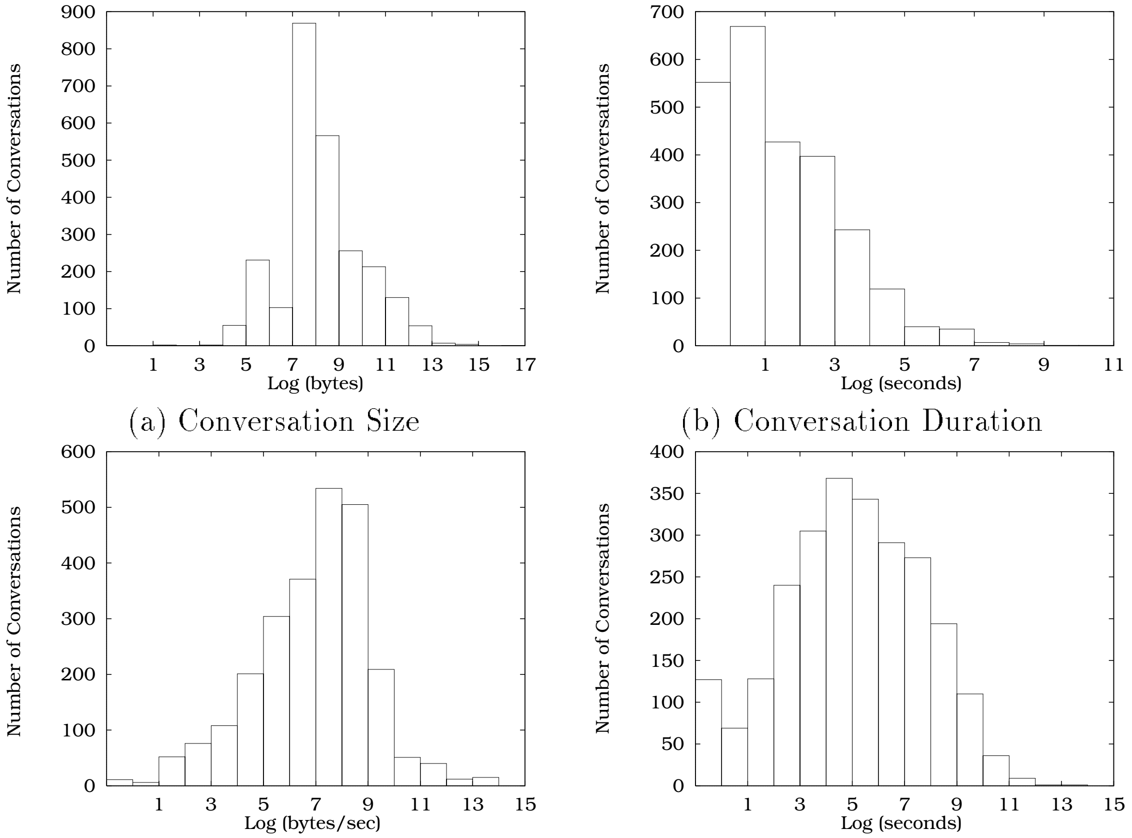

(a) Conversation Size (b) Conversation Duration

(c) Conversation Bandwidth (d) Conversation Interarrival Time

Figure 5.17: Characteristics of Finger Traffic

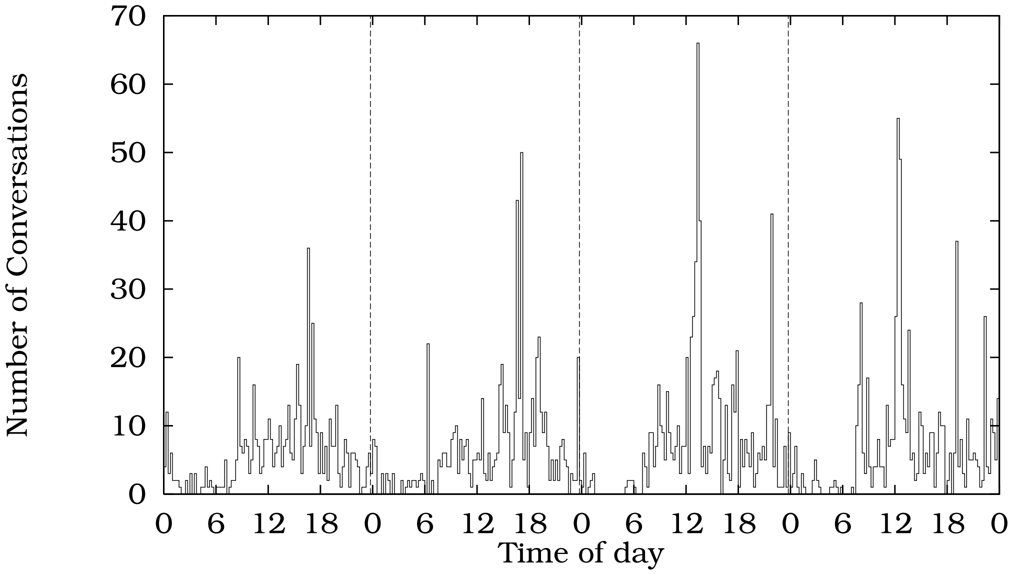

Figure 5.18: Number of Active Finger Conversations

5.5.7 Domain Traffic

The Domain protocol is used for exchanging hostname information among the distributed name servers on the Internet. A statistical analysis of Domain traffic is given in Table 5.21. Figure 5.19 shows histograms of the number of Domain conversations whose size, duration, bandwidth, and interarrival time lie in a particular range. Again, the x-axis uses a logarithmic scale (base 2).

Table 5.21: Statistics of Domain Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(Bytes)	8701.0	111.0	137.626	1.028	44.66
Duration(sec)	141.929	0.120	1.417	3.727	84.86
Bandwidth(KB/sec)	1.446	0.001	0.521	0.645	63.31
Interarrival time(sec)	315.434	0.000	63.322	0.909	64.96

Domain conversations are very small in size and short in duration in this study. Almost all of the Domain conversations (99.89% of them) were between 111 and 256 bytes in size. The average conversation size was 138 bytes, which is much smaller than Paxson's 5.5 Kbytes. The reason for this big difference might be that the Do-

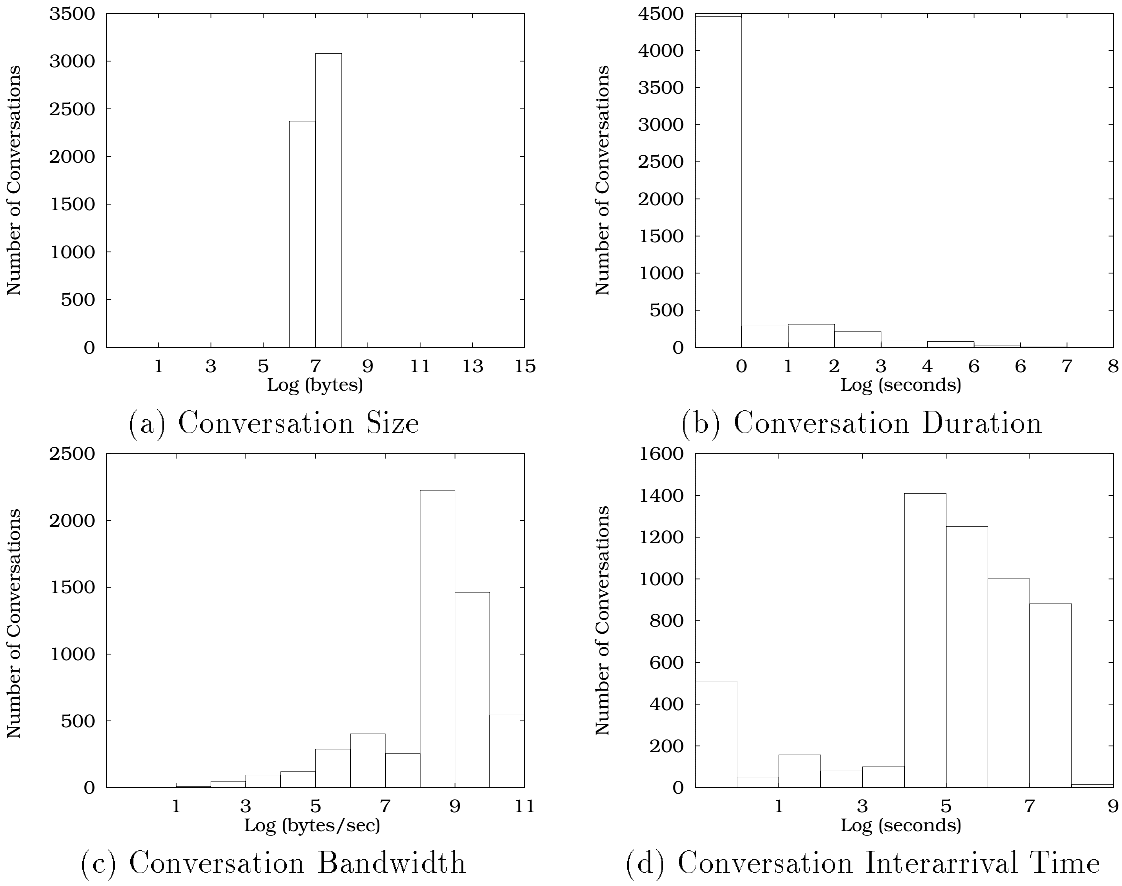

Figure 5.19: Characteristics of Domain Traffic

main Name System used at the University of Saskatchewan is configured to use TCP for looking up an address more often than does the system at Lawrence Berkeley Laboratory on which Paxson's results were obtained. On the latter system a failure to resolve an address locally usually results in a zone transfer that transfers considerably more data [Paxson 94]. Failed zone transfers, which means the DNS got bad information, could lead to small Domain conversations too. About 82% of the Domain conversations were shorter than 1 second in duration. The average duration of Domain conversations was about 1.4 seconds, which is shorter than Paxson's 5 seconds. The average bandwidth was about 521 bytes/sec, which is lower than Paxson's 1211 bytes/sec. The minimum interarrival time was 235 μsec. About 65% of the interarrival times of the Domain conversations were shorter than 63 seconds, the average interarrival time.

The distribution of the number of active Domain conversations is shown in Figure 5.20. The number of active conversations was larger than zero at all periods. This is because as long as there are wide area conversations, there might be wide area Domain conversations to map the domain name to an IP address.

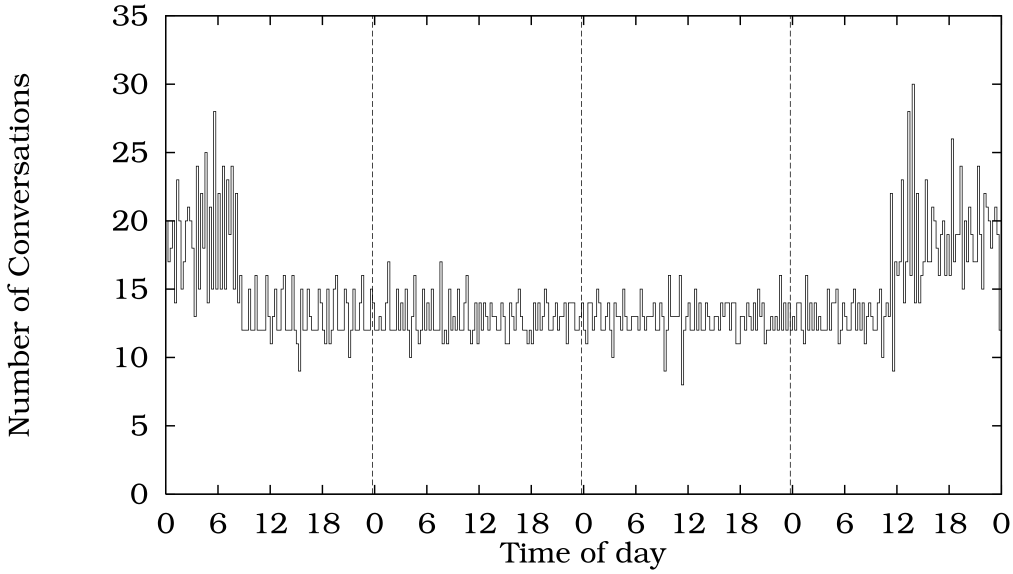

Figure 5.20: Number of Active Domain Conversations

5.5.8 NNTP Traffic

NNTP is used for the distribution, inquiry, retrieval, and posting of news articles over the Internet. Table 5.22 presents summary statistics of NNTP traffic. Figure 5.21 shows the same sort of histograms of the number of NNTP conversations whose size, duration, bandwidth, and interarrival time lie in a particular range.

Table 5.22: Statistics of NNTP Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(MB)	20.849	0.000	0.268	4.374	85.23
Duration(sec)	9893.933	0.103	268.401	2.819	75.00
Bandwidth(KB/sec)	24.826	0.000	1.016	1.979	72.59
Interarrival time(sec)	1993.623	0.666	333.320	0.894	58.78

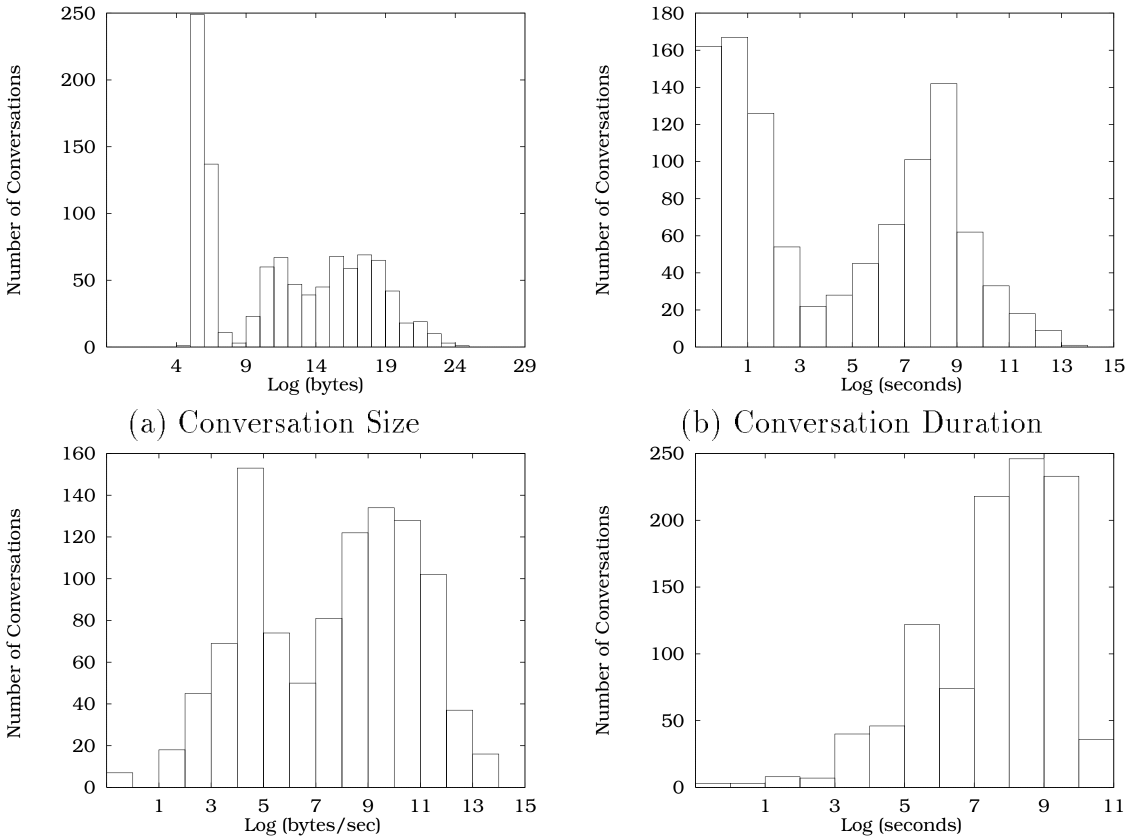

(a) Conversation Size

(b) Conversation Duration

(c) Conversation Bandwidth

(d) Conversation Interarrival Time

Figure 5.21: Characteristics of NNTP Traffic

Usually, NNTP conversations can have very big conversation size since clients can request as many articles as they want in one connection. Although 386 out of 1036 NNTP conversations (37%) were between 32 and 128 bytes in size, the average conversation size, 268 Kbytes, was relatively larger than other wide area conversations in this study. Compared to Paxson's 10.8 Kbytes, the average in this study is also much larger. About 44% of the NNTP conversations were shorter than 4 seconds in duration and about 14% were in the range from 256 to 512 seconds in duration. The average duration of NNTP conversations was about 268 seconds, which is longer than Paxson's 96 seconds. The average bandwidth was about 1 Kbytes/sec, which is 9 times higher than Paxson's 113 bytes/sec. The minimum interarrival time was 0.7 second. About 59% of the interarrival times of the NNTP conversations were shorter than the average interarrival time, which was 333 seconds.

The distribution of the number of active NNTP conversations is shown in Figure 5.22. In general the numbers are small at this point in time. The number of active conversations was larger than zero at almost all periods. Many of those between 1 AM to 7 AM were generated by international hosts in different timezones using the local NNTP server *tribune* on the University of Saskatchewan Campus Ethernet.

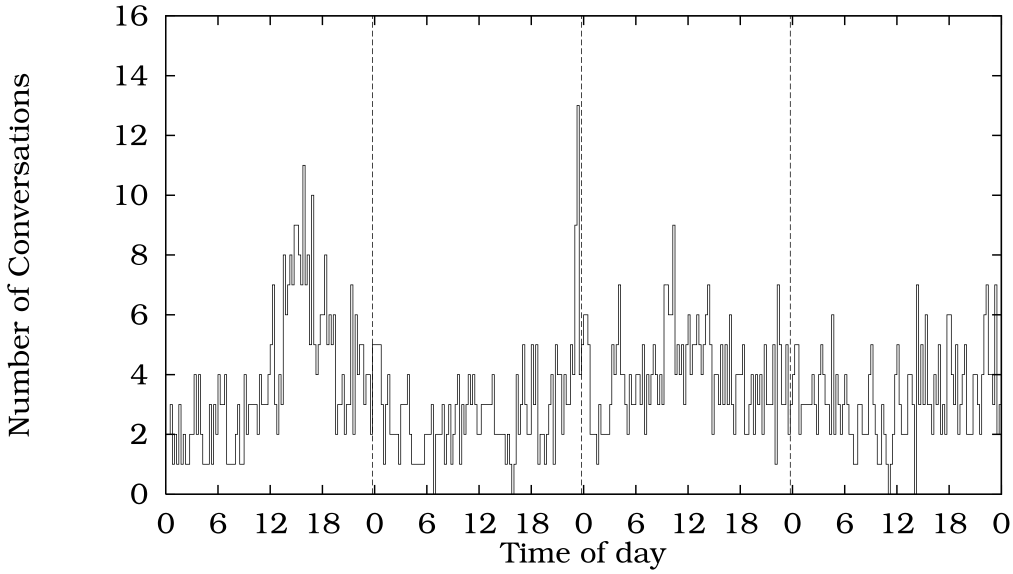

Figure 5.22: Number of Active NNTP Conversations

5.5.9 WWW Traffic

WWW can be used to view all the browsable online information available on the Internet. Table 5.23 summarizes some statistics of WWW traffic. Figure 5.23 shows the same sort of histograms of the number of WWW conversations whose size, duration, bandwidth, and interarrival time lie in a particular range.

About 67% of the WWW conversations were between 1K and 8K bytes in size. The average conversation size was 14 Kbytes. About 65% of the WWW conversations were in the range from 1 to 8 seconds in duration. The average duration of WWW conversations was about 13 seconds. The WWW conversations had a relatively high bandwidth. The average bandwidth was about 1.1 Kbytes/sec. About 83% of the interarrival times of the WWW conversations were shorter than the average interarrival time, which was 59 seconds.

Table 5.23: Statistics of WWW Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(MB)	1.664	0.000	0.014	4.135	83.29
Duration(sec)	2404.696	0.240	13.147	5.511	85.81
Bandwidth(KB/sec)	24.303	0.000	2.038	1.222	67.56
Interarrival time(sec)	4810.051	0.013	59.000	3.616	82.50

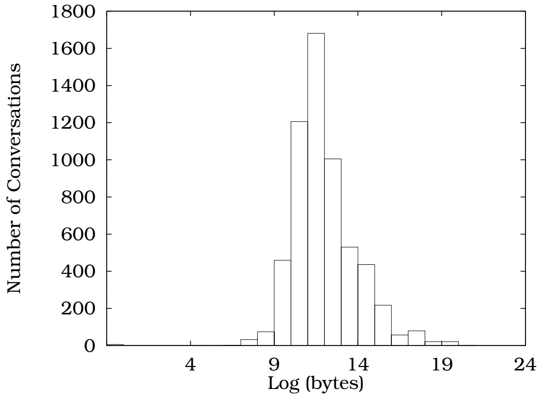

(a) Conversation Size

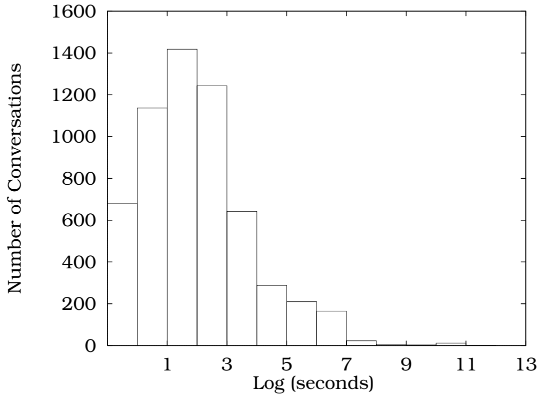

(b) Conversation Duration

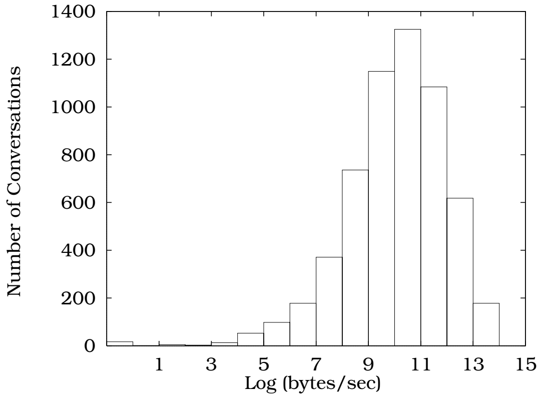

(c) Conversation Bandwidth

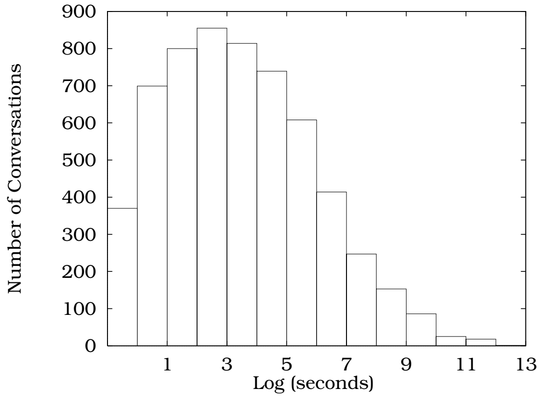

(d) Conversation Interarrival Time

Figure 5.23: Characteristics of WWW Traffic

The distribution of the number of active WWW conversations is shown in Figure 5.24. It can be seen that there are many WWW conversations during the regular working hours.

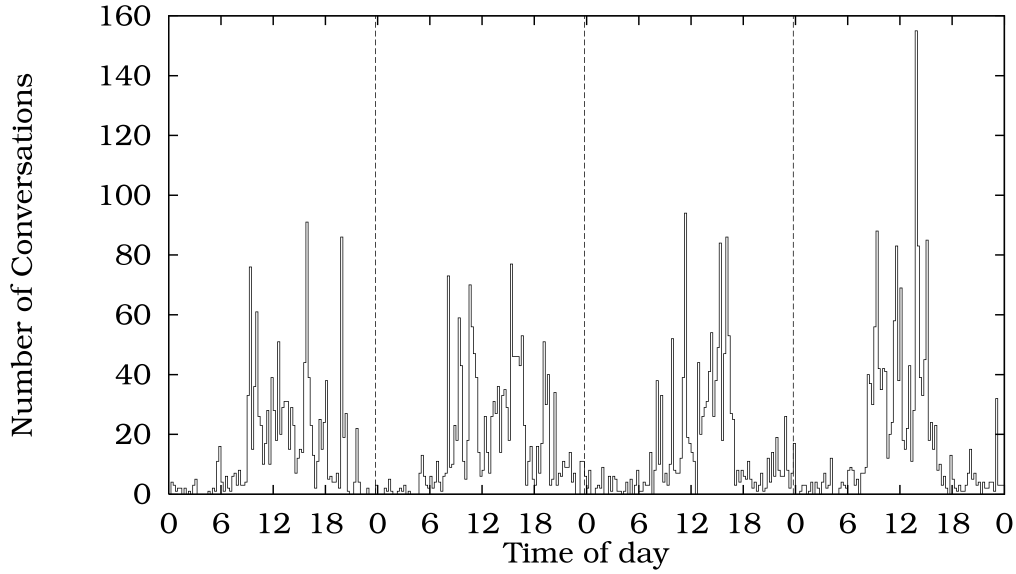

Figure 5.24: Number of Active WWW Conversations

5.5.10 Direction of Data Flow

This section focuses on the distinction between unidirectional and bidirectional traffic in the context of wide area network traffic data flow. In FTP-Data conversations, FTP sites always send the FTP-Data connection request. It doesn't matter whether the command is a *get* or *put* operation, FTP servers are always the source host. By studying the file *CAMPUS-TCP-1*, it can be seen that FTP-Data conversations are always unidirectional. When the command is a *get* operation, data are sent from the FTP server to the client. When the command is a *put* operation, data are sent from the client to the FTP server. FTP, Finger, Gopher, SMTP, Telnet, WWW, Domain, and NNTP conversations are bidirectional. Table 5.24 shows the total data transmitted from source hosts to destination hosts and from destination hosts to source hosts by wide area conversations. Only the 9 most popular applications are listed. Each row of the table consists of protocol name, total data bytes from source hosts to destination hosts and vice versa, the ratio of total data bytes from source hosts to destination hosts to total data bytes from destination hosts to source hosts, and average data bytes per conversation from source hosts to destination hosts and vice versa. Since FTP-Data conversations are unidirectional, the total data bytes from source hosts to destination hosts include all the data transmitted by the FTP-Data conversations in order to show the difference of unidirectional and bidirectional conversations.

Although the applications listed in Table 5.24 are all bidirectional except for FTP-Data, some applications, such as Gopher and Finger, send much more data in one direction than in the other direction.

Table 5.24: Data Flow From Source To Destination and Vice Versa

Protocol	Mbytes Src to Dst	Mbytes Dst to Src	Ratio S-D:D-S	Kbytes/con Src to Dst	Kbytes/con Dst to Src
SMTP	207.758	27.220	8:1	4.775	0.626
Gopher	0.790	173.978	1:220	0.043	9.487
FTP-Data	785.345	0.000		53.769	0.000
WWW	4.202	79.509	1:19	0.721	13.640
Domain	0.220	0.531	1:2	0.040	0.097
Telnet	1.516	92.476	1:61	0.355	21.632
Finger	0.016	1.809	1:113	0.007	0.725
FTP	0.684	3.477	1:5	0.323	1.643
NNTP	262.358	14.997	17:1	253.241	14.475

5.6 Summary of Results

In this chapter, detailed measurements of wide area Internet conversations involving hosts connected to the Campus Ethernet at the University of Saskatchewan have been presented. The study was carried out in the following categories: wide area IP traffic breakdown by transport level protocols; wide area TCP traffic breakdown by department subnets, protocols, and domains; conversation setup latency; per-conversation characteristics of all TCP conversations, such as conversation size, duration, bandwidth, interarrival time, number of simultaneous conversations, and direction of data flow; and characteristics of TCP conversations by types of application protocols. The following results have been identified:

- TCP packets made up 96.01% of all wide area network IP traffic on the University of Saskatchewan Campus Ethernet.

- The Department of Computing Services subnet contributed 68.15% of all the wide area conversations and accounted for 53.88% of all the wide area data on the Campus Ethernet. The Department of Computational Science contributed 8.63% of all the wide area conversations and accounted for 18.12% of all the wide area data.

- A total of 419 application protocols were identified. Nine applications (SMTP, Gopher, FTP-Data, WWW, Domain, Telnet, Finger, FTP-Control, and NNTP) were responsible for 94.8% of all the TCP conversations.

- *Skycat* and *jester* were the busiest hosts at the University of Saskatchewan. Local area networks *trlabs.ca* and *sk.ca* were the busiest local area networks on the Internet that communicated with hosts at the University of Saskatchewan.

- A total of 61 high level Internet domains, including 53 countries, were identified. 27.02% of all the conversations involved hosts within Canada, 44.23% involved hosts in the EDU domain, and 9.49% involved hosts in the COM domain.

- Conversations with a small number of local area networks dominated since only 12% of external LANs were responsible for roughly 80% of the conversations observed over the four-day period.

- The average setup latency for the conversations that were initiated by hosts at the University of Saskatchewan was 0.85 seconds and the average setup latency for the conversations that were initiated by hosts on the Internet was 0.1 seconds.

- FTP-Data, NNTP, SMTP, and Gopher conversations transmitted a large amount of data, comprising about 80% of all the data bytes transmitted by the wide area conversations during the four days.

- The average overall conversation size was 18 Kbytes, the average duration was 75 seconds, the average bandwidth was 951 bytes/sec, and the average interarrival time was 3.35 seconds. The average data rate of the Campus Ethernet for transmitting wide area conversations was 5.3 Kbytes/sec when this study was carried out (April 1994).

Table 5.25: Statistics of the Nine Most Frequently Used Applications

Protocol	Ave. Size (Kbytes)	Ave. Duration (Seconds)	Ave. Bandwidth (KBytes/sec)	Ave. Interarrival Time (Seconds)
SMTP	5	23	0.445	8
Gopher	10	25	1.219	19
FTP-Data	54	21	2.443	24
WWW	14	13	2.038	59
Domain	0.1	1	0.521	63
Telnet	22	562	0.044	81
Finger	1	8	0.361	138
FTP	2	273	0.036	163
NNTP	268	268	1.016	333

- Longer conversations did not necessarily create the most traffic.
- The data flow of FTP-Data traffic is unidirectional while the data flow of SMTP, Gopher, WWW, Domain, Telnet, Finger, FTP, and NNTP is bidirectional (but asymmetric).

5.7 Comparison of the Campus Data and the CS Data

Measurement results from the Campus data and the CS data are similar in many respects, but there are some differences between them.

Table 5.26 summarizes an interesting difference. It is striking that there were no NNTP or Domain conversations observed in the CS Data. This seems very strange. The reason for this is that the campus news server and domain name server are both located on the subnet in the Department of Computing Services. When clients on campus want to have news services or need to find the IP address of a domain, the requests will be sent to the local NNTP server (*tribune*) or the local domain name server (*access*) on campus. The servers will search information for the clients locally first. If they cannot find it, they will send the information request to other servers on the Internet on behalf of these clients. At this time, these two servers present themselves as clients to communicate with servers on the Internet. So all the wide area NNTP and Domain conversations at the University of Saskatchewan are done by machines *tribune* and *access*.

Table 5.26: Summary of Applications in the Campus Data and the CS Data

Applications	Campus Data	CS Data
NNTP, Domain	7.31%	0%
Others	92.69%	100%

There were some other noticeable differences as well. Table 5.27 shows the difference of applications in the Campus data and the CS data after removing all the NNTP and Domain conversations from the Campus data. There were relatively fewer SMTP, Gopher, Telnet conversations in the CS data, and more FTP-Data, WWW, Finger, and FTP conversations. Although the browsing tool *Mosaic* for the WWW application was only installed on hosts in the Department of Computational Science in December 1993, this application is already responsible for the third most wide area conversations on the Computational Science Ethernet.

Table 5.27: Renormalized Comparison of the Campus Data and the CS Data

Applications	Campus Data	CS Data
SMTP	45.08%	36.63%
Gopher	19.00%	4.65%
FTP-Data	15.13%	24.34%
WWW	6.04%	13.87%
Telnet	4.43%	2.81%
Finger	2.58%	6.58%
FTP	2.19%	3.32%
Others	5.55%	7.80%

The average sizes of SMTP and FTP-Data conversations derived from the CS data were much larger than that seen in the Campus data. Table 5.28 summarizes these differences. This means that the average sizes of the mail messages and FTP files from CS users tend to be larger. Users in the Department of Computational Science use these tools mostly to exchange information such as technical reports, image files, and application software packages across the Internet. This usage results in files that are relatively large.

Table 5.28: Summary of Average Sizes of Application Conversations in Kbytes

Applications	Campus Data	CS Data
SMTP	5	13
FTP-Data	54	109

There were some significant differences observed in Gopher usage as shown in Table 5.29. In the CS data, the volume of data imported into the CS Ethernet by the Gopher application was 219 times that exported from the CS Ethernet, while in the Campus data, the volume of data exported from the Campus Ethernet by Gopher conversations was 3 times that imported into the Campus Ethernet. This is because there was no Gopher server on the CS Ethernet. Users in the Department of Computational Science used Gopher servers on the Internet and brought files from these servers to the CS Ethernet. Overall, the volume of data was obtained from Gopher servers at the University of Saskatchewan by users from the Internet was more than that fetched from Gopher servers on the Internet by users at the University of Saskatchewan.

Table 5.29: Summary of Gopher Usage

Data	Ratio (Outgoing Data : Incoming Data)
Campus data	3.27:1
CS data	1:218.85

A total of 98 different application protocols were observed in the CS data while a total of 419 different application protocols were observed in the Campus data. The group of application protocols in the CS data was only one-fourth of that seen in the Campus data. This suggests that CS users make heavier use of a more restricted set of applications than does the campus community as a whole. Isolating other departments may well show similar concentration to a (possibly different) subset of applications.

Chapter 6

Summary and Conclusions

6.1 Thesis Summary

This thesis has provided general information about the Internet and detailed introductions to some of the most popular protocols in the TCP/IP suite. It has also surveyed measurement studies of the Internet over the world in terms of growth of the Internet, end-to-end behaviour over the Internet, and traffic characteristics of the Internet. Most importantly, it has presented results showing how users at the University of Saskatchewan use the Internet. It has provided characteristics of the wide area TCP conversations involving not only the full campus network but also the subnet in the Department of Computational Science as well. The following is a summary of the work that was done.

Traces for this research were captured using the network monitoring and packet acquisition tool TCPDUMP. Two traces formed the basis for the analysis. A complete trace of IP packets was captured for 24 hours on April 22 (Friday), 1994. This trace was about 594 Mbytes in size, with a total of 6,556,683 packets in it. Analysis of this trace confirmed that TCP packets are dominant. Subsequently, a second trace was collected over four continuous days from April 25 to 28 (Monday to Thursday), 1994. This trace was limited to those packets marking the beginning or end of a TCP conversation. It was about 64 Mbytes in size, with a total of 617,806 packets in it (about 154,452 packets per day). After discarding packets from incomplete conversations, 103,016 complete wide area TCP conversations were derived from this second trace. This represents an average of 25,754 wide area TCP conversations per day on the Campus Ethernet. These complete wide area TCP conversations formed the basis for a detailed analysis of the use of various application protocols.

It was found that TCP packets made up a large percentage (96.01%) of all wide area IP packets involving the University of Saskatchewan over the period of this study, and that SMTP, Gopher, FTP-Data, WWW, Domain, Telnet, Finger, FTP-Control, and NNTP accounted for the majority of the conversations (and the majority of the bytes). People at the University of Saskatchewan use the Internet mainly for information searching and sharing with sites over the world. The major tools supporting this are electronic mail, network news, Gopher, and World-Wide Web services. FTP provides a mean to exchange information while Telnet lets users use resources on a remote site. These constituted the applications used most frequently.

A total of 461 Mbytes of data were transmitted by wide area TCP conversations per day on the Campus Ethernet, with 156 Mbytes transmitted from the Campus Ethernet to the Internet and 305 Mbytes transmitted from the Internet to the Campus Ethernet per day.

Many hosts in the Department of Computing Services are used generally by faculty, staff, and students to access Internet resources. The subnet in this department accounted for a significant portion of the wide area traffic, up to 68.15% of all the wide area conversations on the Campus Ethernet. Hosts in the Department of Computational Science are used by faculty, staff, and students in the Department. As a group they accounted for 8.63% of all the wide area conversations on the Campus Ethernet during the study. There was an average of 2,222 wide area TCP conversations per day on the Computational Science subnet. A total of 84 Mbytes of data were transmitted by wide area TCP conversations per day on the Computational Science subnet, with 18 Mbytes transmitted from the Computational Science Ethernet to the Internet and 66 Mbytes transmitted from the Internet to the Computational Science Ethernet. No wide area Domain and NNTP traffic were contributed by the subnet in the Department of Computational Science since all these wide area conversations were implemented by the local domain name server and local network news server in the Department of Computing Services.

The interconnectivity between the Campus Ethernet and the rest of the world was quite rich. During the measurement period, the University of Saskatchewan Campus Ethernet reached 53 countries. A total of 61 highest level Internet domains were identified, of which CA, EDU, and COM were the most frequently observed.

Over the period of observation, conversations tended to be concentrate to small subsets of local area networks on the Internet. About 80% of the campus wide area conversations involved only 12% of the external LANs, while 80% of the wide area conversations involving the Computational Science Ethernet involved 25% of the external LANs.

The setup latency for the wide area conversations was reasonably short. The average setup latency for the conversations that were initiated by hosts at the University of Saskatchewan was 0.85 seconds and the average setup latency for the conversations initiated by hosts on the Internet was 0.1 seconds.

Per-conversation characteristics of some particular TCP conversations in terms of average size, duration, bandwidth, interarrival time were summarized in Table 5.25 and Table A.22. In general, FTP-Data, NNTP, SMTP, and Gopher conversations transmit a large amount of data. NNTP and FTP-Data conversations are especially large in size. Telnet, FTP-Control, NNTP, and Login conversations have very long conversation durations. FTP-Data, WWW, Gopher, and NNTP conversations have relatively high data rate. Domain and Ident conversations are small in size and short in duration, while FTP-Control conversations have very low data rates. FTP-Data generates unidirectional traffic while SMTP, Gopher, WWW, Domain, Telnet, Finger, FTP, NNTP, Login, and Ident generate bidirectional traffic.

6.2 Directions for Future Research

There are a number of possible directions for future work in this area. It would be interesting to monitor the wide area traffic at the University of Saskatchewan on a periodic basis to see how network use changes and grows with time. Periodic usage updates are important for network management and planning since the Internet is growing rapidly, both in size and in services, and applications involving audio and video traffic (with different characteristics from current traffic) will emerge strongly in future networks. Workload models to describe the behaviour of different protocols based on such measurements could be developed for use in performance and capacity planning studies.

There are a wide variety of operating system platforms available on campus, including SunOS, Ultrix, HP-UX, VMS, etc. It might be interesting to analyze the Internet activity by platform.

Data about how users at the University of Saskatchewan use the Internet has been lacking. The goal of this thesis was to provide an analysis of wide area Internet conversations involving hosts at the University of Saskatchewan. It is hoped that this measurement and analysis will contribute to a better understanding of usage of the Internet by people at the University of Saskatchewan and will contribute to an accumulating body of knowledge on Internet usage more generally.

References

[Anklesaria 93] F. Anklesaria, M. McCahill, P. Lindner, D. Johnson, D. Torrey, and B. Alberti, *The Internet Gopher Protocol*, RFC 1436, Network Information Center, SRI International, Menlo Park, California, March 1993.

[Asaba 92] T. Asaba, K. Claffy, O. Nakamura, and J. Murai, "An Analysis of International Academic Research Network Traffic between Japan and Other Nations", *Proceedings of Inet '92*, Kobe, Janpan, June, 1992, pp. 431-440.

[Berners-Lee 92] T. Berners-Lee, R. Cailliau, J-F. Groff, and B. Pollermann, "World-Wide Web: The Information Universe", *Electronic Networking: Research, Applications and Policy*, Vol. 1, No. 2, Spring 1992, pp. 52-58.

[Berners-Lee 93] T. Berners-Lee, R. Cailliau, N. Rellow, and A. Secret, "The World-Wide Web Initiative", *Proceedings of Inet '93*, San Francisco, August, 1993, pp. DBC1-DBC4.

[Caceres 91] R. Caceres, P. B. Danzig, S. Jamin, and D. J. Mitzel, "Characteristics of Wide-Area TCP/IP Conversations", *Proceedings of ACM SIGCOMM '91*, Zurich, Switzerland, September 1991, pp. 101-112.

[Caceres 92] R. Caceres, *Multiplexing Traffic at the Entrance to Wide-Area Networks*, Ph. D. Dissertation, Computer Science Division, Department of EECS, University of California, Berkeley, 1992.

[Caceres 93] R. Caceres, through e-mail connection, August 1993.

[Cerf 83] V. Cerf and E. Cain, “The DoD Internet Architecture Model”, *Computer Networks*, Vol. 7, No. 5, October 1983, pp. 307-318.

[Claffy 93] K. C. Claffy, H. -W. Braun, and G. C. Polyzos, “Traffic Characteristics of the T1 NSFNET Backbone”, *Proceedings of IEEE INFOCOM ’93*, Vol. 2, San Francisco, California, July 1993, pp. 885-892.

[Comer 91] D. E. Comer, *Internetworking with TCP/IP, Volume I*, Prentice Hall, 1991.

[Danzig 92a] P. B. Danzig, S. Jamin, R. Caceres, D. J. Mitzel, and D. Estrin, “An Empirical Workload Model for Driving Wide-Area TCP/IP Network Simulations”, *Internetworking: Research and Experience*, Vol. 3, 1992, pp. 1-26.

[Danzig 92b] P. B. Danzig, K. Obraczka, and A. Kumar, “An Analysis of Wide-Area Name Server Traffic”, *Proceedings of ACM SIGCOMM ’92*, Baltimore, Maryland, September 1992, pp. 281-292.

[Danzig 93] P. B. Danzig, K. Obraczka, and S. Li, “Internet Resource Discovery Services”, *IEEE Computer Magazine*, September 1993, pp. 8-22.

[Deering 91] S. Deering, *ICMP Router Discovery Messages*, RFC 1256, Network Information Center, SRI International, Menlo Park, California, 1991.

[Gulati 93] N. Gulati, C. Williamson, and R. Bunt, “Local Area Network Traffic Locality: Characteristics and Application”, in *Local Area Network Interconnection (R.O. Onvural and A. Nilsson, eds.)*, Plenum Press, 1993, pp. 233-250.

[Heimlich 90] S. A. Heimlich, “Traffic Characterization of the NSFNET National Backbone”, *Proceedings of the 1990 Winter USENIX Conference*, Berkeley, California, pp. 207-227.

[Jacobson 89] V. Jacobson, C. Leres, and S. McCanne, *TCPDUMP (1)*, Unix Manual page, June 1989.

[Jain 86] R. Jain, and S. A. Routhier, "Packet Trains – Measurements and a New Model for Computer Network Traffic", *IEEE Journal on Selected Areas in Communications*, Vol. SAC-4, No. 6, September 1986, pp. 986-995.

[Johns 93] M. Johns, *Identification Protocol (Ident)*, RFC 1413, Network Information Center, SRI International, Menlo Park, California, 1993.

[Kantor 86] B. Kantor and P. Lapsley, *Network News Transfer Protocol*, RFC 977, Network Information Center, SRI International, Menlo Park, California, 1986.

[Kantor 91] B. Kantor, *BSD Rlogin*, RFC 1282, Network Information Center, SRI International, Menlo Park, California, 1991.

[Krol 92] E. Krol, *The Whole Internet User's Guide & Catalog*, O'Reilly & Associates, 1992.

[Leres 93] C. Leres, *HF (1)*, Unix Manual page, available via anonymous ftp to ftp.ee.lbl.gov, November 23, 1993.

[Lottor 92] M. Lottor, *Internet Growth (1981-1991)*, RFC 1296, Network Information Center, SRI International, Menlo Park, California, 1992.

[Lottor 94] M. Lottor, SRI International, statistics available via anonymous ftp to ftp.nw.com, directory pub/zone, January, 1994.

[Lynch 93] D. C. Lynch and M. T. Rose, *Internet System Handbook*, Addison-Wesley, January 1993.

[Maffeis 93] S. Maffeis, "File Access Patterns in Public FTP Archives and an Index for Locality of Reference", *ACM SIGMETRICS Performance Evaluation Review*, Vol. 20, No. 3, March 1993, pp. 22-35.

[McCanne 93] S. McCanne and V. Jacobson, "The BSD Packet Filter: A New Architecture for User-level Packet Capture", *Proceedings of the 1993 Winter USENIX Conference*, San Diego, California, January 1993, pp. 259-269.

[Merit 94] Merit Network Information Center Services, NSFnet Statistics available via anonymous ftp to nic.merit.edu, directory statistics/nsfnet, January 1994.

[Mills 83] D. L. Mills, *Internet Delay Experiments*, RFC 889, Network Information Center, SRI International, Menlo Park, California, 1983.

[Mills 92] D. L. Mills, *Network Time Protocol*, RFC 1305, Network Information Center, SRI International, Menlo Park, California, 1992.

[Mockapetris 87] P. Mockapetris, *Domain Names – Implementation and Specification)*, RFC 1035, Network Information Center, SRI International, Menlo Park, California, 1987.

[Mogul 91] J. C. Mogul, "Network Locality at the Scale of Processes", *Proceedings of ACM SIGCOMM '91*, Zurich, Switzerland, September 1991, pp. 273-284.

[Mukherjee 92] A. Mukherjee, *On the Dynamics and Significance of Low Frequency Components of Internet Load*, Technical Report MS-CIS-92-83, Department of Computer and Information Science, University of Pennsylvania, 1992.

[Paxson 91] V. Paxson, *Measurements and Models of Wide Area TCP Conversations*, Technical Report LBL30840, Computer Systems Engineering Department, Lawrence Berkeley Laboratory, University of California, Berkeley, California, May 1991.

[Paxson 93a] V. Paxson, "Growth Trends in Wide-Area TCP Connections", to appear in *IEEE Network*, 1994.

[Paxson 93b] V. Paxson, "Empirically-Derived Analytic Models of Wide-Area TCP Connections", submitted to *IEEE/ACM Transactions on Networking*, 1993.

[Paxson 94] V. Paxson, through e-mail connection, July 1994.

[Postel 80] J. B. Postel, *User Datagram Protocol*, RFC 768, Network Information Center, SRI International, Menlo Park, California, 1980.

[Postel 81a] J. B. Postel, *Internet Protocol*, RFC 791, Network Information Center, SRI International, Menlo Park, California, 1981.

[Postel 81b] J. B. Postel, *Transmission Control Protocol*, RFC 793, Network Information Center, SRI International, Menlo Park, California, 1981.

[Postel 82] J. B. Postel, *Simple Mail Transfer Protocol*, RFC 821, Network Information Center, SRI International, Menlo Park, California, 1982.

[Postel 83] J. B. Postel and J. K. Reynolds, *Telnet Protocol Specification*, RFC 854, Network Information Center, SRI International, Menlo Park, California, 1983.

[Postel 85] J. B. Postel and J. K. Reynolds, *File Transfer Protocol*. RFC 959, Network Information Center, SRI International, Menlo Park, California, 1985.

[Reynolds 92] J. Reynolds and J. Postel, *Assigned Numbers*, RFC 1340, Network Information Center, SRI International, Menlo Park, California, 1992.

[Rutkowski 94] T. Rutkowski, *The NSFnet Statistics*, available via anonymous ftp to ftp.isoc.org, directory isoc/charts, 1994.

[Sanghi 93] D. Sanghi, A. K. Agrawala, O. Gudmundsson, and B. N. Jain, "Experimental Assessment of End-to-end Behavior on Internet", *Proceedings of IEEE INFOCOM '93*, Vol. 2, San Francisco, California, March 1993, pp. 867-874.

[Schmidt 92] A. Schmidt and R. Campbell, "Internet Protocol Traffic Analysis with Applications for ATM Switch Design", *Computer Communication Review*, Vol. 23, No. 2, April 1993, pp 39-52.

[Schwartz 93] M. Schwartz and J. Quarterman, "A Measurement Study of Changes in Service-Level Reachability in the Global Internet", *Proceedings of Inet '93*, San Francisco, August, 1993, pp. CBB1-CBB10.

[Tanenbaum 89] A. Tanenbaum, *Computer Networks*, Second edition, Prentice Hall, 1989.

[Wakeman 93] I. Wakeman, D. Lewis, and J. Crowcroft, "Traffic Analysis of Trans-Atlantic Traffic", *Computer Communications*, Vol. 16, No. 6, June 1993, pp. 376-388.

[Zimmermann 80] H. Zimmermann, "OSI Reference Model - the ISO Model of Architecture for Open Systems Interconnection", *IEEE Transactions on Communications*, Vol. COM-28, No. 4, April 1980, pp. 425-432.

[Zimmerman 90] D. Zimmerman, *The Finger User Information Protocol* , RFC 742, Network Information Center, SRI International, Menlo Park, California, November 1990.

Appendix A

Measurement Results for the CS Ethernet

This appendix presents a detailed analysis of wide area traffic on the CS Ethernet using file *CS-TCP-1* (see Section 5.2). The characteristics are described under five categories. Section A.1 presents a breakdown of wide area TCP traffic on the CS Ethernet by types of application protocols. Section A.2 provides information on how wide conversations involving hosts on the CS Ethernet reach. Section A.3 measures the wide area conversation setup latency on the CS Ethernet. Section A.4 analyzes the wide area conversations in several subcategories and particular application protocols as well. Section A.5 gives conclusions reached about the characteristics of wide area traffic on the CS Ethernet.

A.1 TCP Traffic Breakdown

Table A.1: Breakdown of Complete TCP Conversations by Protocol

Protocol	Number	Percent	Total MB	Data Rate (Bytes/sec)	Con. Size (KB/con)	Duration (Secs/con)
SMTP	3256	36.63	43.661	558.403	13.410	24.014
FTP-Data	2163	24.34	236.475	3973.986	109.327	27.511
WWW	1233	13.87	17.979	2284.357	14.582	6.383
Finger	585	6.58	0.299	75.192	0.511	6.789
Gopher	413	4.65	2.858	1267.558	6.920	5.460
FTP	295	3.32	0.584	4.462	1.981	443.885
Ident	258	2.90	0.012	25.414	0.045	1.781
Telnet	250	2.81	9.346	57.745	37.383	647.378
Login	59	0.66	4.905	22.086	83.131	3763.964
Others	376	4.24	17.854	53.137	47.485	893.632
All	8888	100.00	333.973	332.933	37.576	112.863

A total of 98 application protocols were identified in file *CS-TCP-1*. 334 Mbytes of wide area data were transmitted by the CS subnet, with 209 Mbytes to and 125 Mbytes from the CS Ethernet. Among the 8,888 TCP conversations, 52% of them were initiated by hosts on the Internet while 48% were initiated by hosts on the CS Ethernet.

Table A.2: Volumes of Data from the CS Ethernet to the Internet and Vice Versa

Protocol	CS Ethernet to Internet (Mbytes, total)	Internet to CS Ethernet (Mbytes, total)	Ratio (Out:In)
SMTP	6.505	37.157	1:5.71
FTP-Data	49.346	187.129	1:3.79
WWW	1.364	16.615	1:12.18
Finger	0.132	0.167	1:1.27
Gopher	0.013	2.845	1:218.85
FTP	0.170	0.414	1:2.44
Ident	0.008	0.003	2.67:1
Telnet	7.946	1.399	5.68:1
Login	3.334	1.570	2.12:1
Others	2.608	15.247	1:5.85
All	71.426	262.546	1:3.68

Table A.3: TCP Conversations Initiated by Hosts on the Internet

Protocol	Number	Percent	Total MB	Data Rate (Bytes/sec)	Con. Size (KB/con)	Duration (Secs/con)
SMTP	2127	46.00	24.234	397.431	11.393	28.667
FTP-Data	1639	35.45	180.798	4940.416	110.310	22.328
Finger	312	6.75	0.132	54.809	0.422	7.691
Ident	222	4.80	0.010	26.620	0.046	1.720
Telnet	194	4.20	8.023	60.440	41.354	684.221
Others	130	2.81	3.578	21.672	27.519	1269.815
All	4624	100.00	216.773	544.427	46.880	86.109

Table A.4: TCP Conversations Initiated by Hosts on the CS Ethernet

Protocol	Number	Percent	Total MB	Data Rate (Bytes/sec)	Con. Size (KB/con)	Duration (Secs/con)
WWW	1233	28.92	17.979	2284.357	14.582	6.383
SMTP	1129	26.48	19.428	1128.589	17.208	15.247
FTP-Data	524	12.29	55.677	2430.252	106.254	43.722
Gopher	413	9.69	2.858	1267.558	6.920	5.460
Finger	273	6.40	0.167	106.304	0.612	5.759
Other	692	16.23	21.090	38.129	30.477	799.324
All	4264	100.00	117.200	193.733	27.486	141.875

A.2 How Wide do Wide Area TCP Conversations Reach?

A total of 37 highest Internet domains were observed in file *CS-TCP-2* (see Section 5.3). 31 countries out of 78 that have access to the Internet or general electronic mail services were captured (including Canada). 45 hosts on the CS subnet in the Department of Computational Science and 454 networks on the Internet were observed.

Table A.5: Number of Conversations as a Function of Domains

Domain	Description	No. of Conversations	Percent of Total	No. of Subdomains	No. of Protocols
AT	Austria	3	0.03	1	2
AU	Australia	73	0.82	3	20
BE	Belgium	15	0.17	2	2
BR	Brazil	7	0.08	3	1
CA	Canada	2188	24.70	52	21
CH	Switzerland	73	0.82	4	3
DE	Germany	66	0.74	19	6
DK	Denmark	4	0.05	2	1
EG	Egypt	2	0.02	1	2
ES	Spain	1	0.01	1	1
FI	Finland	59	0.67	7	6
FR	France	27	0.30	8	5
HK	Hong Kong	30	0.34	4	4
HR	Croatia	177	2.00	4	10
HU	Hungary	7	0.08	1	3
IE	Ireland	5	0.06	1	2
IN	India	41	0.46	2	5
IT	Italy	9	0.10	6	3
JP	Japan	26	0.29	4	3
KR	Korea (South)	8	0.09	1	3
MX	Mexico	7	0.08	1	2
NL	Netherlands	142	1.60	9	7
NO	Norway	35	0.40	2	2
NZ	New Zealand	8	0.09	2	2
PT	Portugal	1	0.01	1	1
SE	Sweden	232	2.62	9	6
SG	Singapore	11	0.12	3	3
TW	Taiwan	36	0.41	1	7
UK	United Kingdom	221	2.49	2	9
VE	Venezuela	5	0.06	1	1
ZA	South Africa	5	0.06	1	1
COM	Commercial	1293	14.59	93	16
EDU	Educational	3390	38.26	137	44
GOV	Government	161	1.82	12	9
MIL	US Military	18	0.20	3	4
NET	Network	336	3.79	20	23
ORG	Non-Profit Organization	138	1.56	15	7

Table A.6: Ten Most Popular Hosts on the CS Ethernet

Hosts	Number of Conversations	Percent of Total	Number of Internet LANs
Skdad	2781	31.39	271
Skorpio	1354	15.28	124
Sparkle	1010	11.40	77
Batoche	651	7.35	60
Skuncle	480	5.42	21
Piapot	360	4.06	38
Spr	276	3.12	49
Fvr	236	2.66	51
Graph	231	2.61	26
Waskesiu	168	1.90	18

Conversations with a small number of local area networks dominated as can be seen in Figure A.1. Roughly 80% of the conversations involved only 25% of the external LANs. In fact, only 10 external LANs (out of a total of 454) were responsible for 22.79% of the conversations.

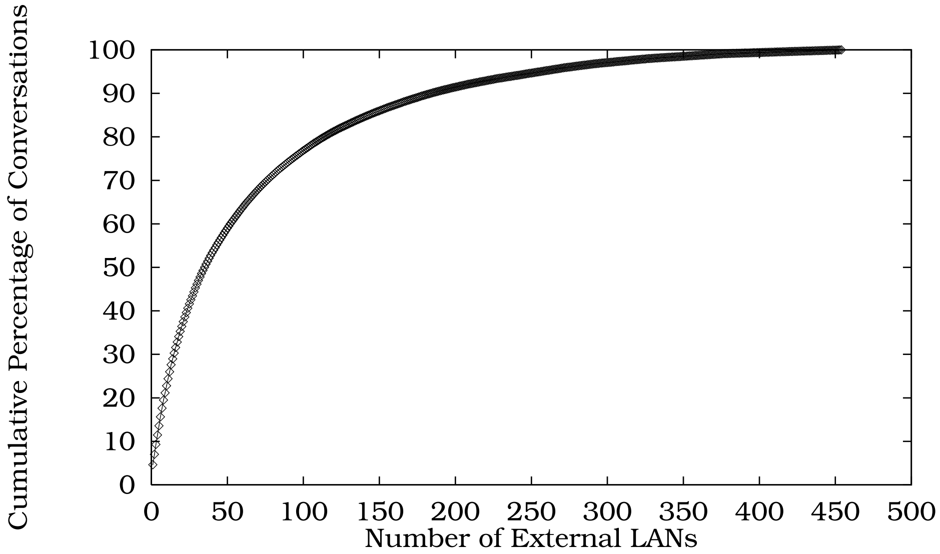

Figure A.1: Cumulative Percentage of Number of Conversations with External LANs

Table A.7: Ten Most Popular Local Area Networks on the Internet

LANs	Number of Conversations	Percent of Total	Number of UofS hosts
Umich.edu	410	4.63	12
Uiuc.edu	211	2.38	16
Uu.net	201	2.27	4
Sedsystems.ca	193	2.18	8
Nrc.ca	189	2.13	15
Ac.uk	182	2.05	12
Jhu.edu	177	2.00	2
Trlabs.ca	164	1.85	7
Mit.edu	147	1.66	7
Berkeley.edu	145	1.64	7

Table A.8: Five Most Popular FTP Servers on the Internet (Total 124)

Hosts	Num. of Con.	Percent of Total
Wasp.eng.ufl.edu	9	3.06
Klingon.epas.utoronto.ca	7	2.38
Rudolf.ethz.ch	7	2.38
Ftp.uu.net	5	1.70
Dg-rtp.rtp.dg.com	5	1.70

Table A.9: Five Most Popular Gopher Servers on the Internet (Total 77)

Hosts	Num. of Con.	Percent of Total
Ftp.sunet.se	96	23.24
Wired.com	26	6.30
Phoebe.albany.edu	23	5.57
Mach1.wlu.ca	20	4.84
Nass.ucs.ubc.ca	15	3.63

Table A.10: Five Most Popular WWW Servers on the Internet (Total 75)

Hosts	Num. of Con.	Percent of Total
Sils.umich.edu	316	25.63
Www.sun.com	106	8.60
Thrall.cm.cf.ac.uk	65	5.27
S1.arc.umn.edu	60	4.87
Harpo.ncsa.uiuc.edu	54	4.38

A.3 Conversation Setup Latency

The average setup latency for the conversations initiated by hosts on the CS Ethernet was 0.99 seconds, with a maximum of 60.05 seconds and a minimum of 0.02 seconds. The average setup latency for the conversations initiated by hosts on the Internet was 0.01 seconds, with a maximum of 2.97 seconds and a minimum of 0.95 msec.

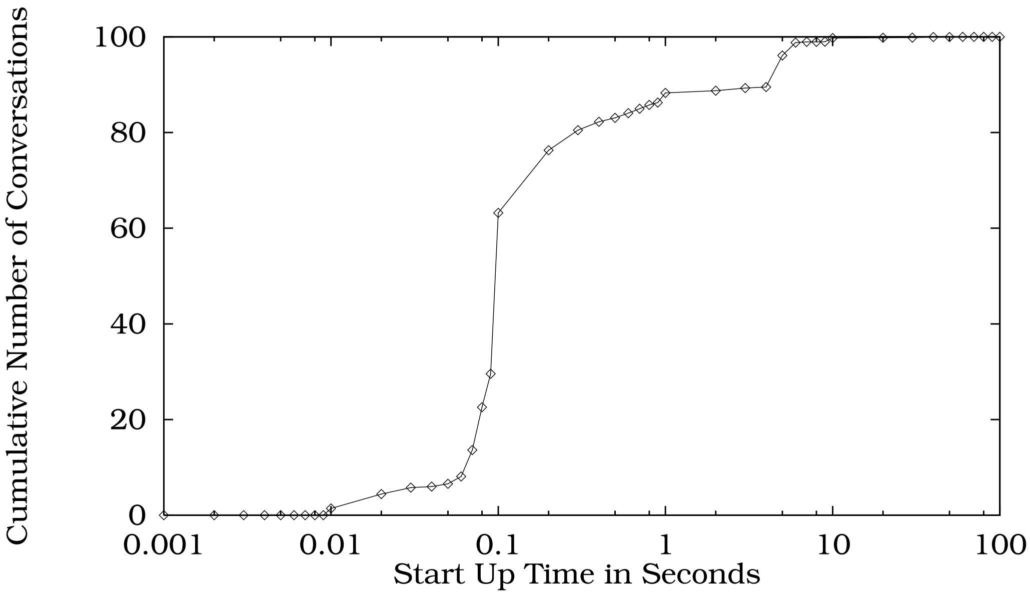

Figure A.2: Setup Latency of Conversations (CS-initiated), All Protocols

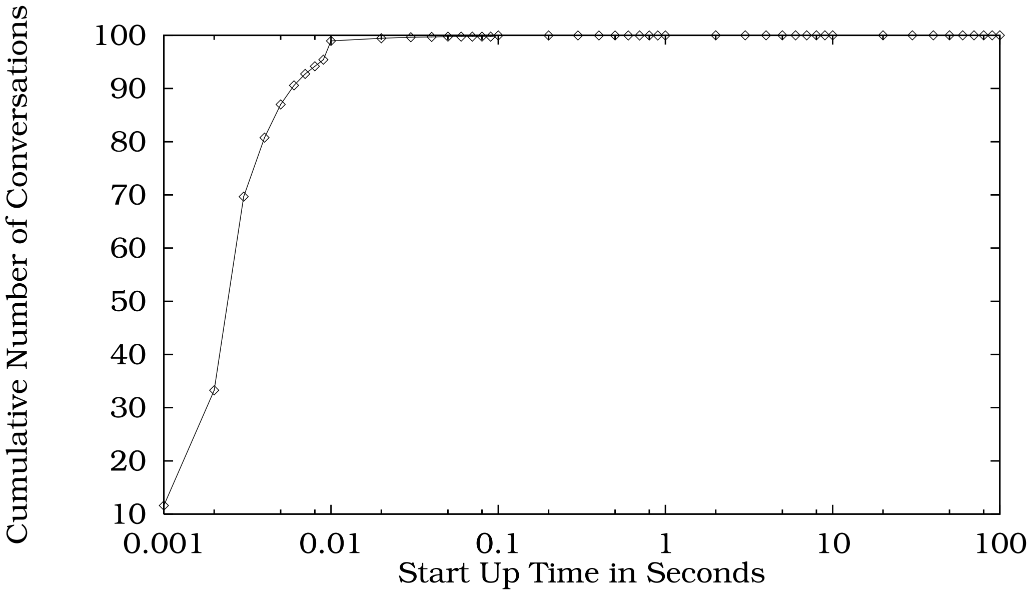

Figure A.3: Setup Latency of Conversations (Internet-initiated), All Protocols

A.4 Analysis of TCP Conversations

This section presents per-conversation statistics of the Department of Computational Science data. The methodology of analysis used is the same with that in Section 5.5.

Section A.4.1 presents per-conversation statistics in terms of all complete TCP conversations retrieved from file *CS-TCP-1*. These statistics include conversation size, conversation duration, the relationship between size and duration, conversation bandwidth, interarrival time of conversations, and number of concurrent conversations. Section A.4.2 to Section A.4.9 analyze particular application protocols widely used in the Department. Section A.4.10 gives some ideas about direction of data flow of these application protocols.

A.4.1 All TCP Traffic

Table A.11: Statistics of TCP Conversations, All Protocols

	Max	Min	Average	Cov	% smaller than average
Conversation size(MB)	13.258	0.000	0.038	8.724	91.58
Duration(sec)	38115.543	0.067	112.863	8.803	91.88
Bandwidth(KB/sec)	203.788	0.000	1.594	2.612	75.52
Interarrival time(sec)	2155.529	0.001	38.842	2.177	75.17

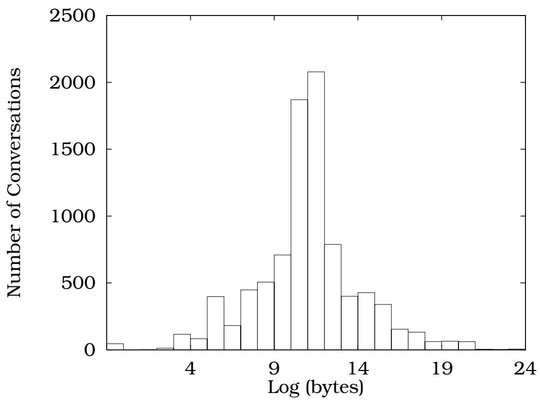

(a) Conversation Size

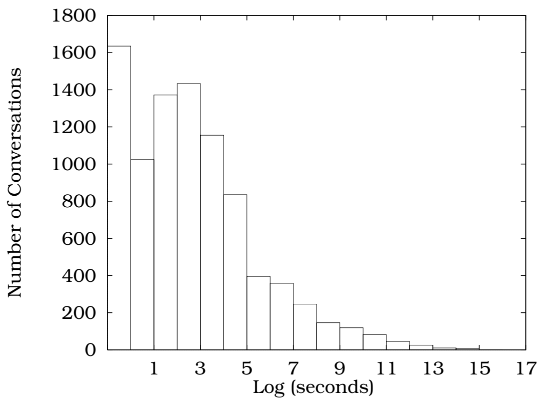

(b) Conversation Duration

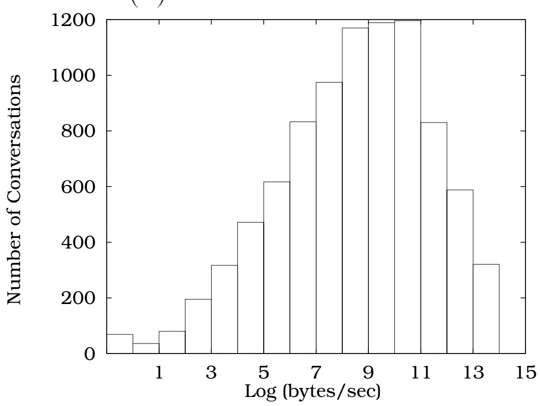

(c) Conversation Bandwidth

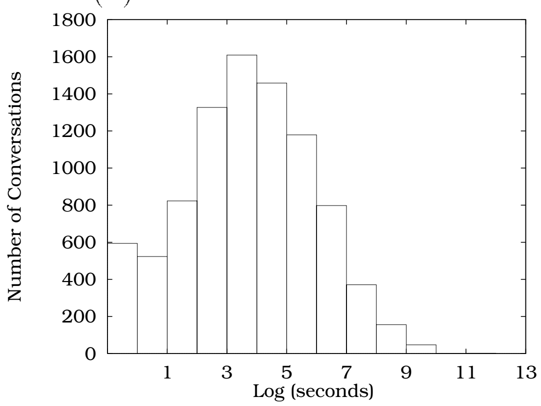

(d) Conversation Interarrival Time

Figure A.4: Characteristics of TCP Conversations

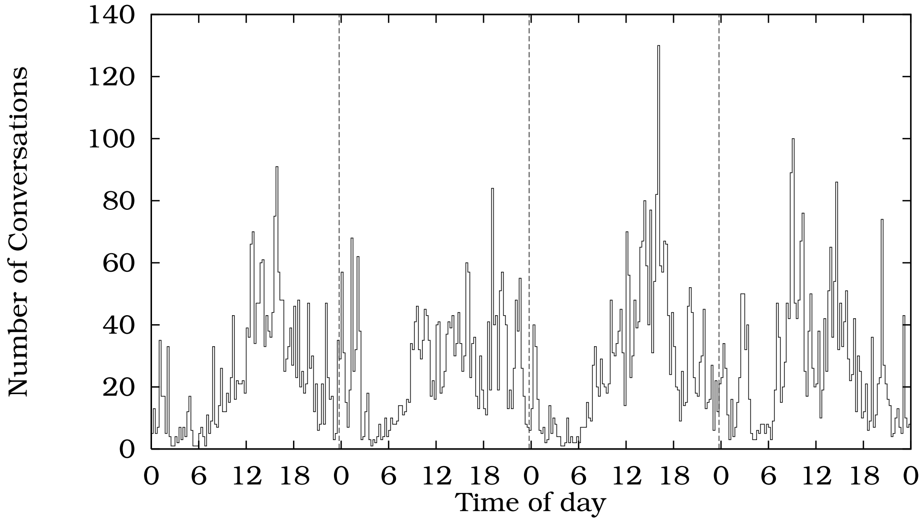

Figure A.5: Number of Active Conversations, All Protocols

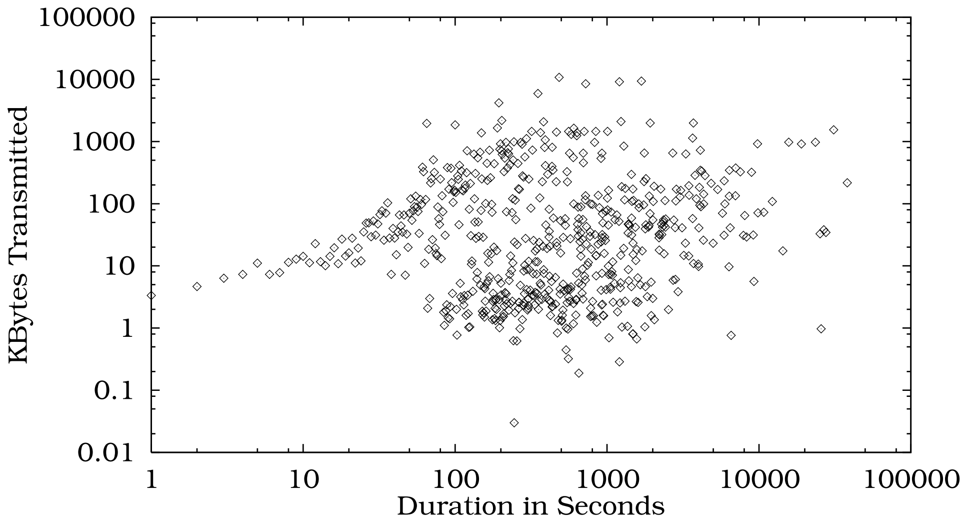

Figure A.6: Conversation Size vs Conversation Duration, All Protocols

A.4.2 File Transfer Protocol Traffic (FTP and FTP-Data)

Table A.12: Statistics of FTP (FTP-Control) Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(KB)	12.195	0.000	1.981	1.022	63.73
Duration(sec)	25689.795	0.641	443.885	3.589	76.61
Bandwidth(KB/sec)	0.462	0.000	0.017	2.188	73.56
Interarrival time(sec)	17241.836	5.881	1154.592	1.763	71.19

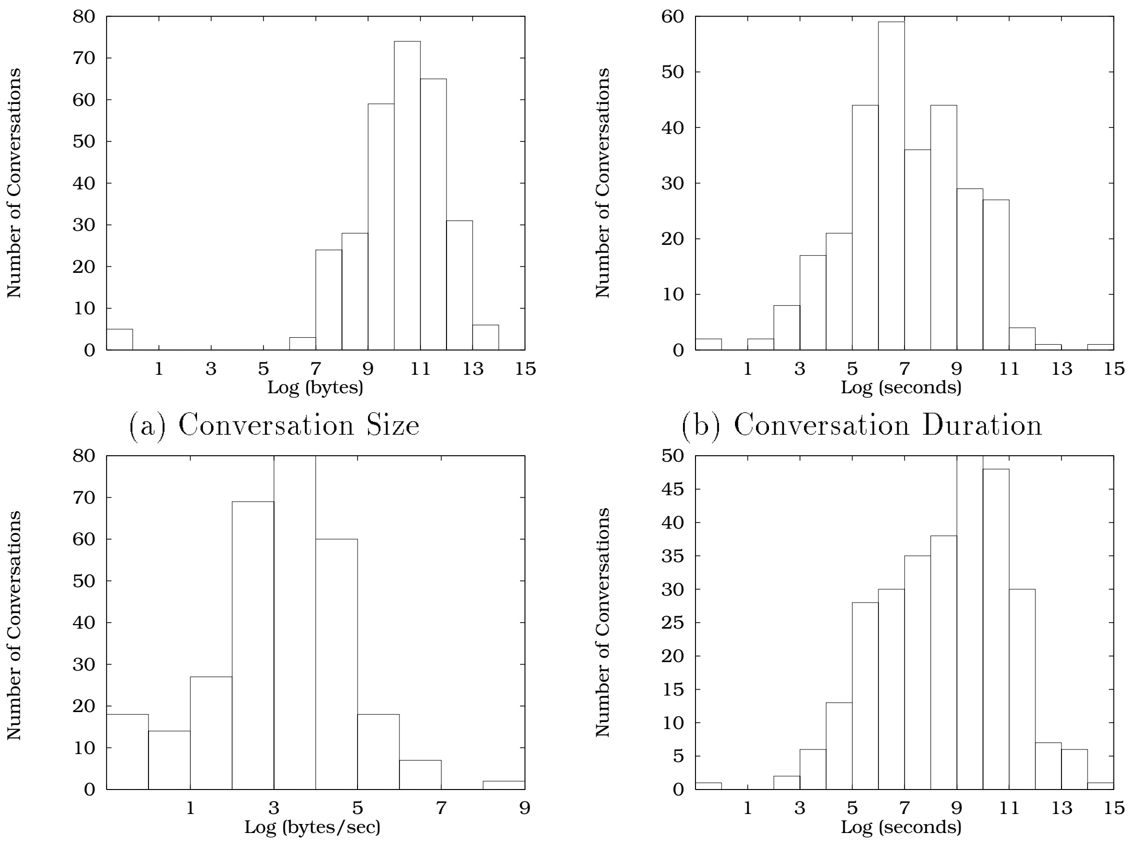

(a) Conversation Size (b) Conversation Duration

(c) Conversation Bandwidth (d) Conversation Interarrival Time

Figure A.7: Characteristics of FTP (FTP-Control) Conversations

Table A.13: Statistics of FTP-Data Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(MB)	10.737	0.000	0.109	5.022	87.79
Duration(sec)	3701.490	0.067	27.511	5.271	88.16
Bandwidth(KB/sec)	73.772	0.000	3.278	1.848	73.51
Interarrival time(sec)	13490.359	0.026	156.851	4.407	88.12

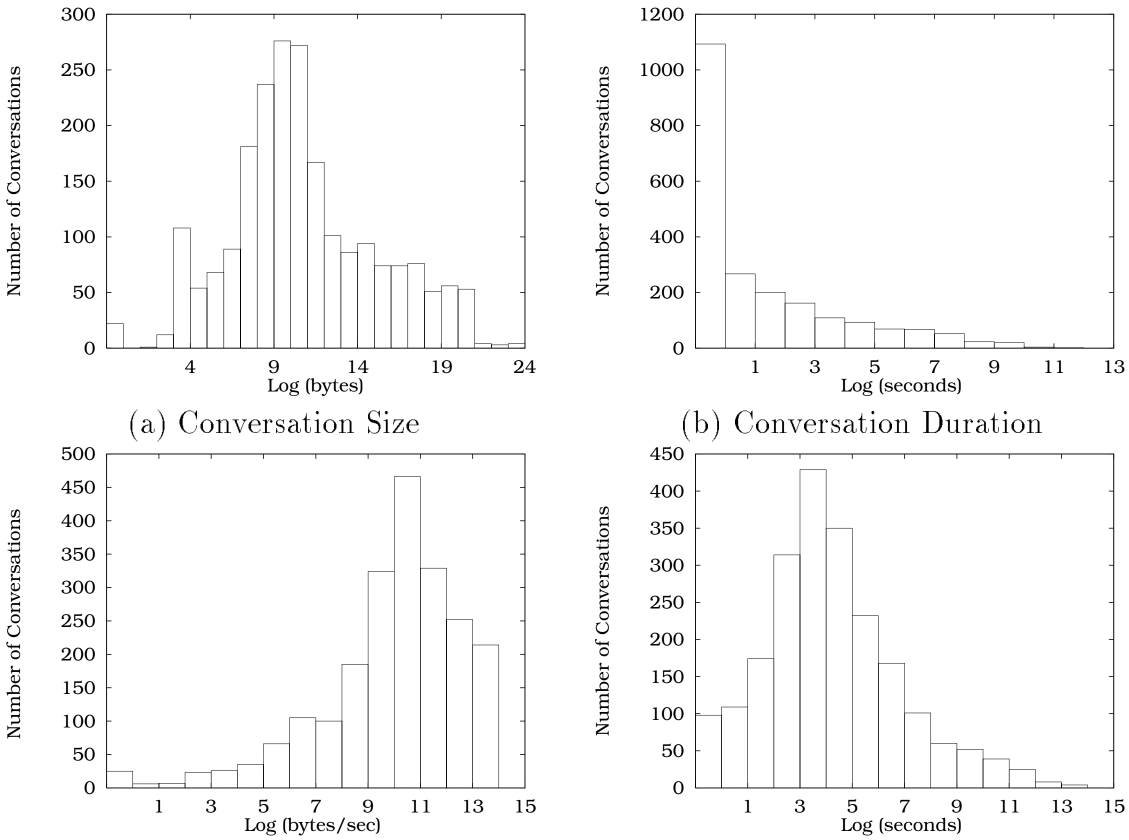

(a) Conversation Size (b) Conversation Duration

(c) Conversation Bandwidth (d) Conversation Interarrival Time

Figure A.8: Characteristics of FTP-Data Conversations

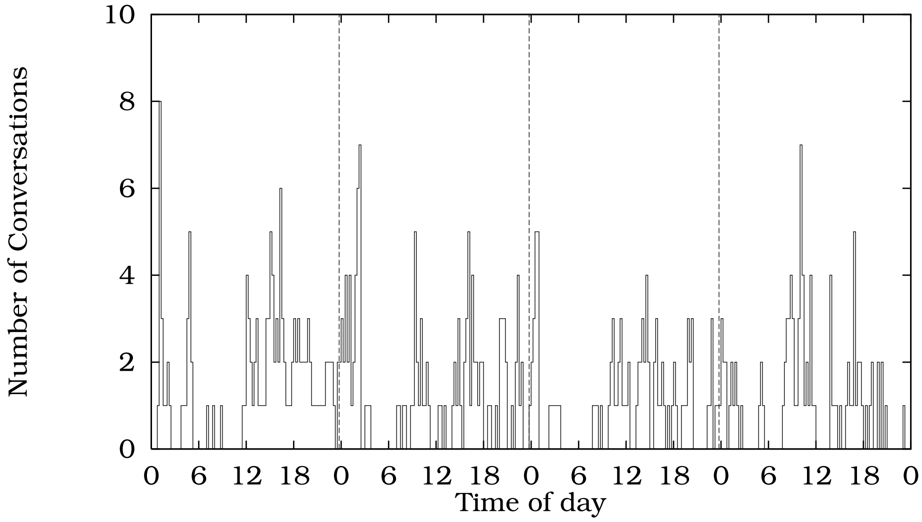

Figure A.9: Number of Active FTP (FTP-Control) Conversations

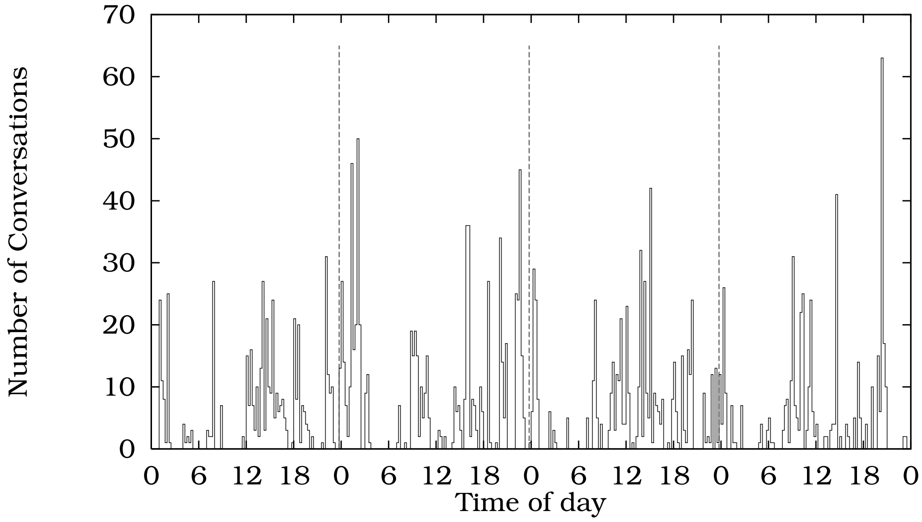

Figure A.10: Number of Active FTP-Data Conversations

A.4.3 SMTP Traffic

Table A.14: Statistics of SMTP Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(MB)	13.258	0.000	0.013	21.375	92.97
Duration(sec)	1683.247	0.163	24.014	2.379	79.91
Bandwidth(KB/sec)	203.788	0.000	0.584	6.240	72.21
Interarrival time(sec)	2860.085	0.001	105.921	1.628	70.82

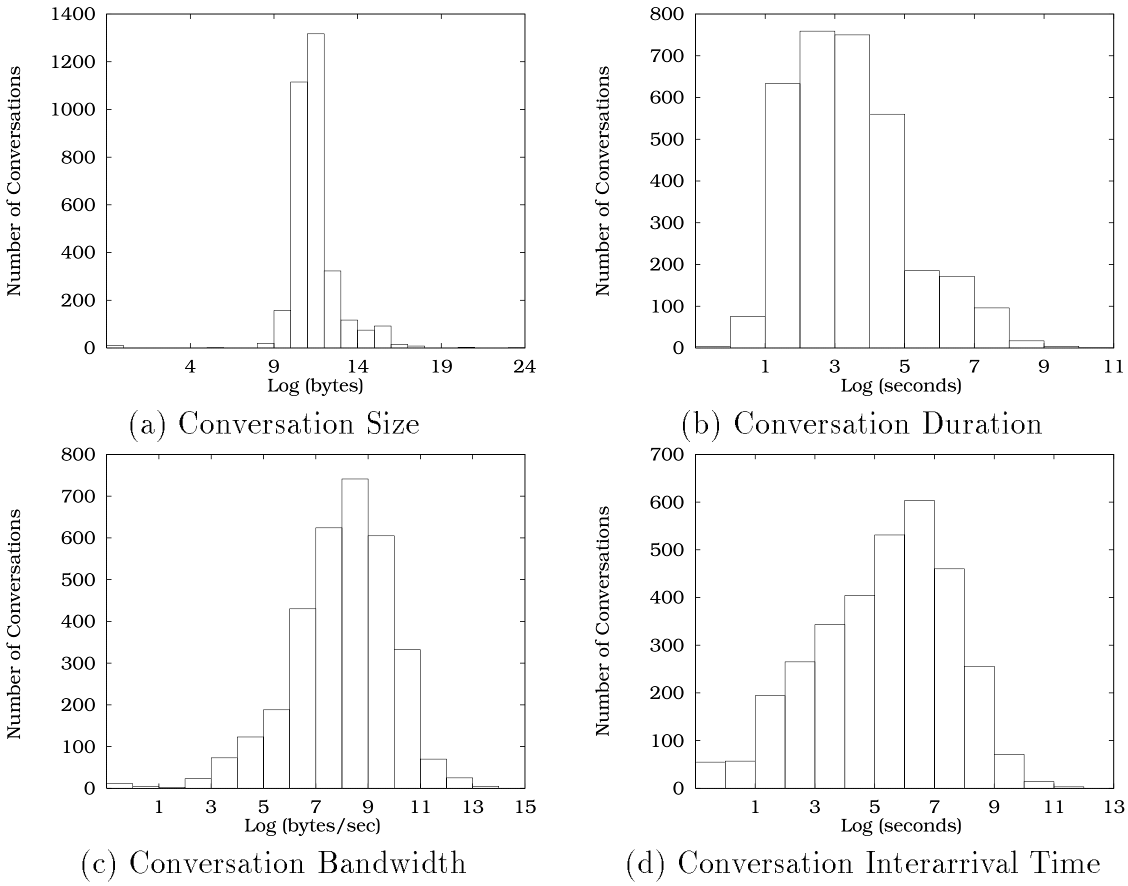

Figure A.11: Characteristics of SMTP Traffic

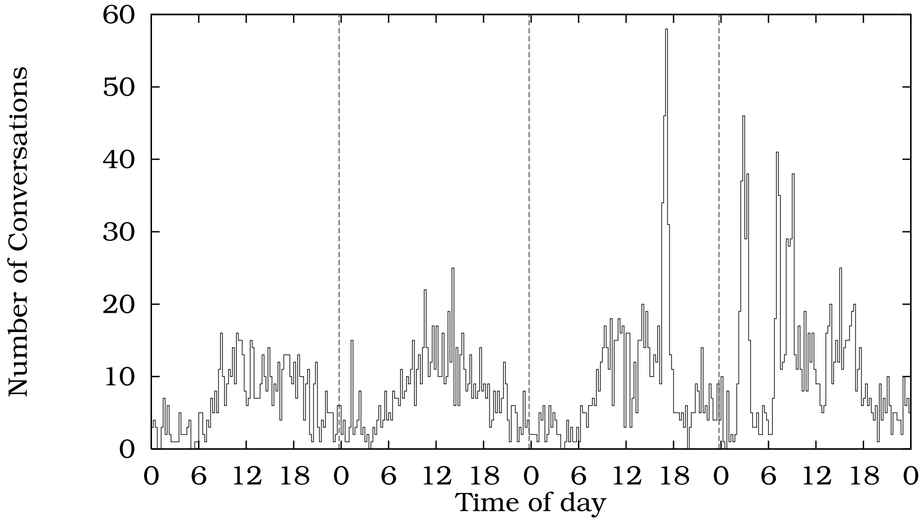

Figure A.12: Number of Active SMTP Conversations

A.4.4 Telnet Traffic

Table A.15: Statistics of Telnet Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(KB)	1136.283	0.047	37.383	2.539	75.20
Duration(sec)	5970.572	1.763	647.378	1.456	69.20
Bandwidth(KB/sec)	0.486	0.000	0.060	1.099	68.00
Interarrival time(sec)	22143.955	1.085	1362.288	1.853	74.00

A.4.5 Gopher Traffic

Table A.16: Statistics of Gopher Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(KB)	165.712	0.010	6.920	3.269	86.92
Duration(sec)	245.733	0.308	5.460	2.734	75.54
Bandwidth(KB/sec)	9.873	0.000	1.206	1.311	70.94
Interarrival time(sec)	47328.520	2.318	702.100	5.690	92.98

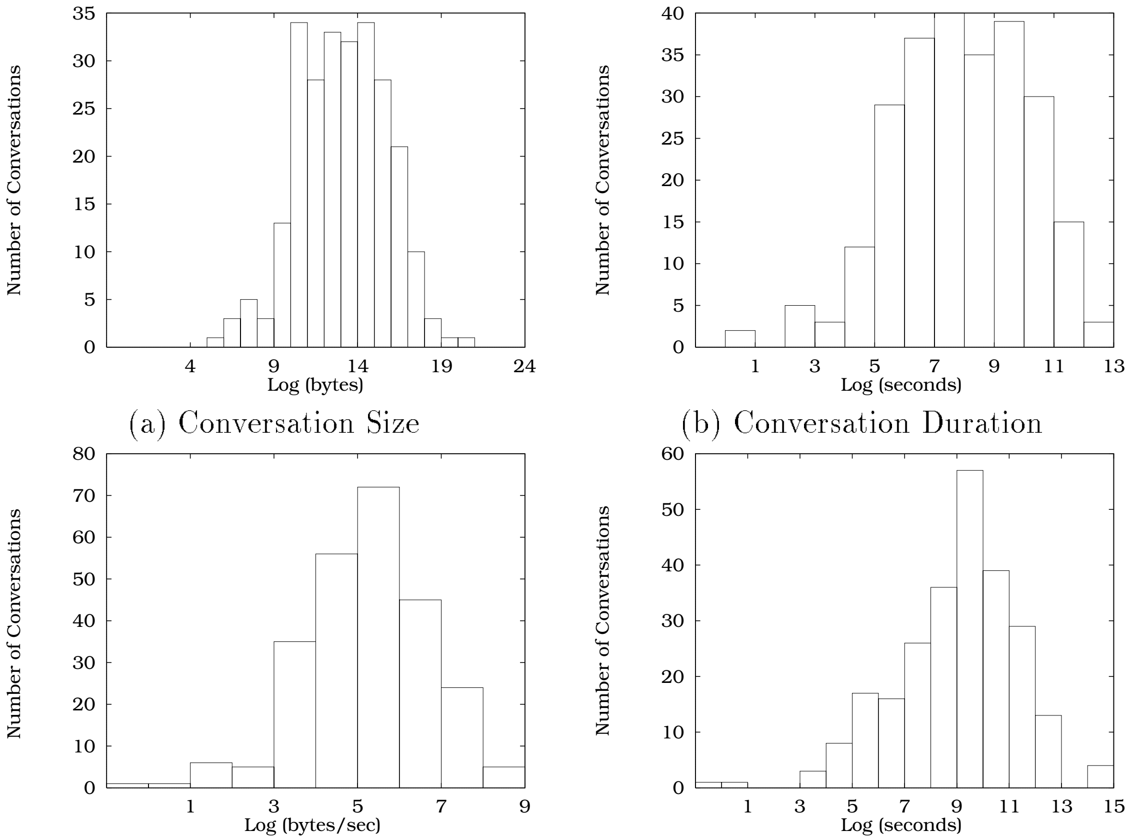

(a) Conversation Size (b) Conversation Duration

(c) Conversation Bandwidth (d) Conversation Interarrival Time

Figure A.13: Characteristics of Telnet Traffic

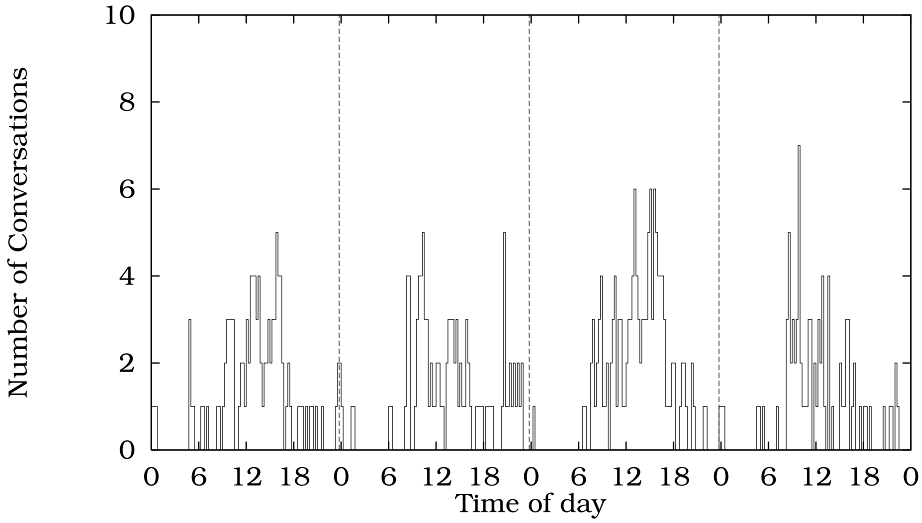

Figure A.14: Number of Active Telnet Conversations

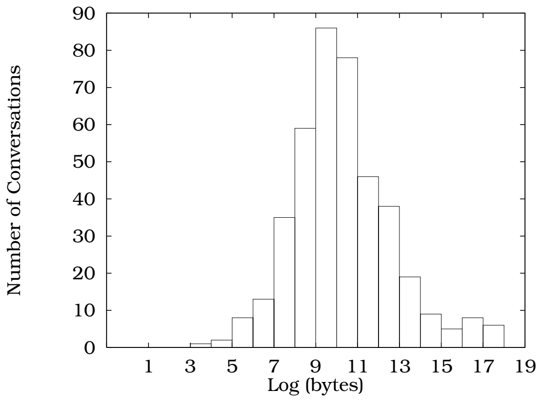

(a) Conversation Size

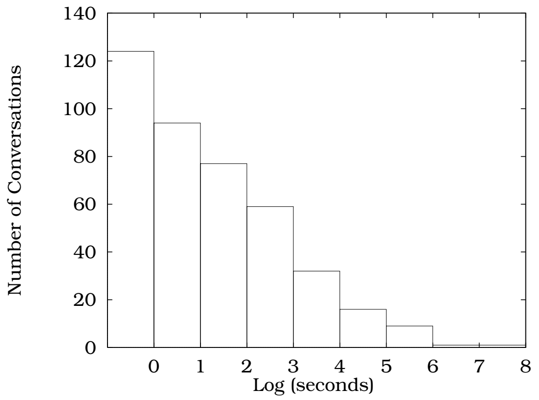

(b) Conversation Duration

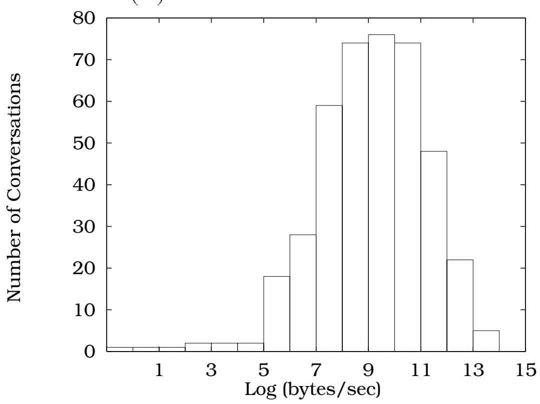

(c) Conversation Bandwidth

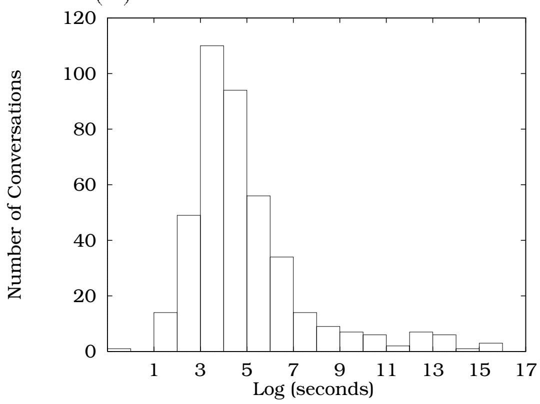

(d) Conversation Interarrival Time

Figure A.15: Characteristics of Gopher Traffic

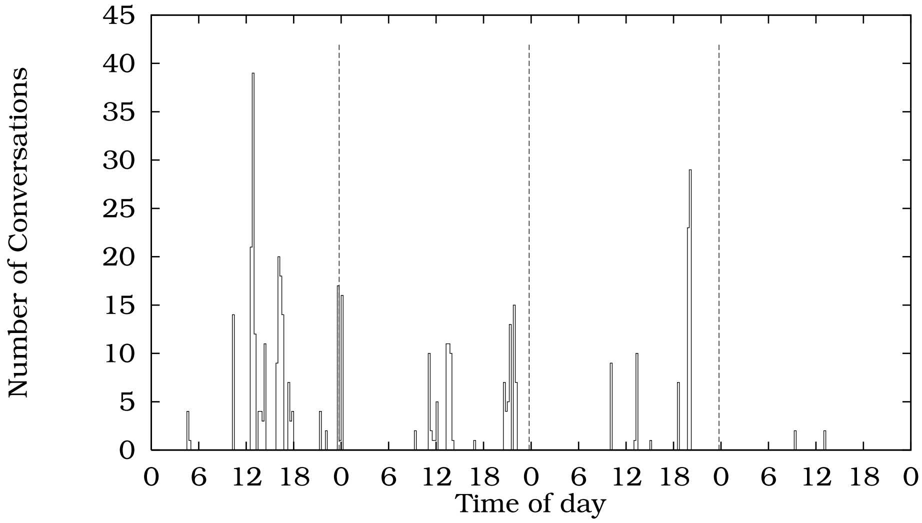

Figure A.16: Number of Active Gopher Conversations

A.4.6 Finger Traffic

Table A.17: Statistics of Finger Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(KB)	8.411	0.000	0.511	1.462	72.99
Duration(sec)	127.752	0.270	6.789	1.568	67.18
Bandwidth(KB/sec)	11.833	0.000	0.347	3.654	82.91
Interarrival time(sec)	24713.088	0.001	590.343	2.592	73.85

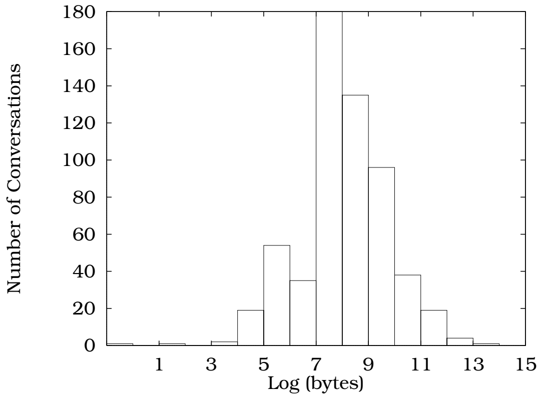

(a) Conversation Size

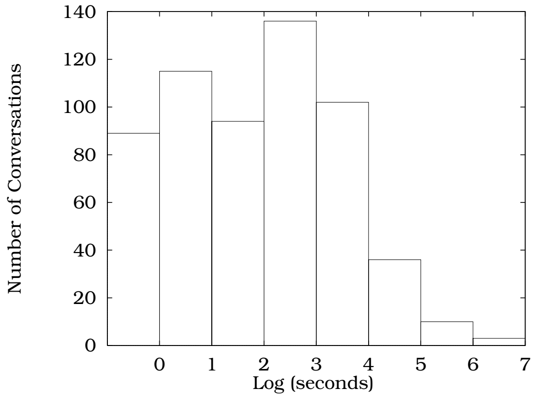

(b) Conversation Duration

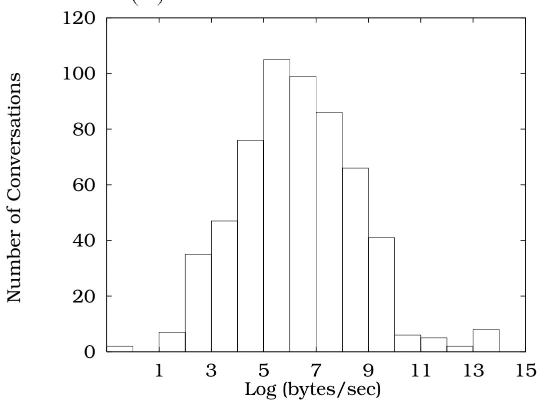

(c) Conversation Bandwidth

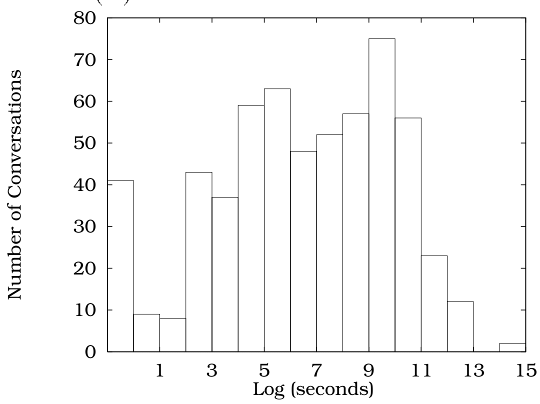

(d) Conversation Interarrival Time

Figure A.17: Characteristics of Finger Traffic

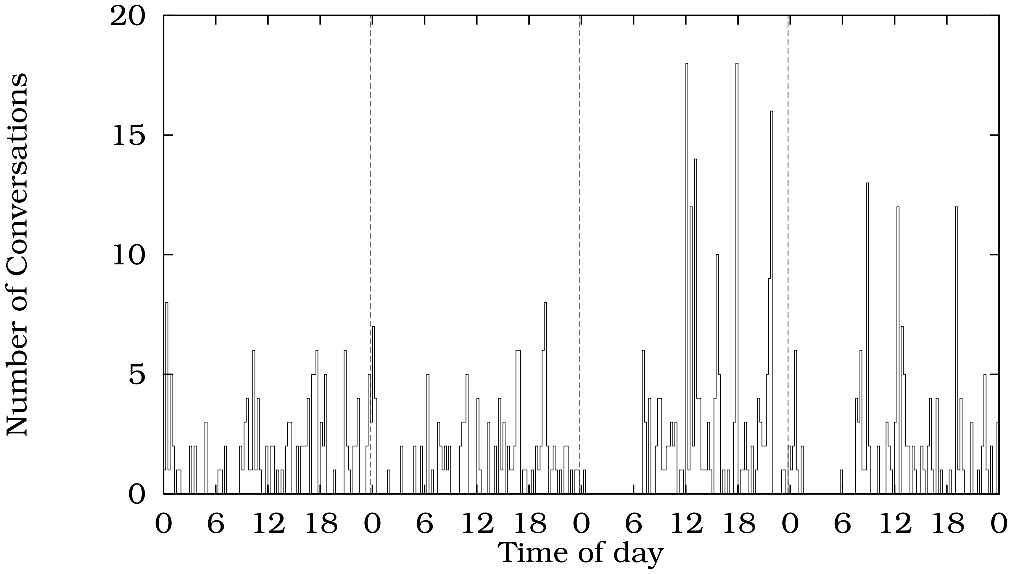

Figure A.18: Number of Active Finger Conversations

A.4.7 Login Traffic

Table A.18: Statistics of Login Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(KB)	920.371	0.306	83.131	2.281	81.36
Duration(sec)	38115.543	7.950	3763.964	2.043	76.27
Bandwidth(KB/sec)	0.655	0.001	0.068	1.413	61.02
Interarrival time(sec)	39454.500	65.541	5208.162	1.469	64.41

A.4.8 Ident Traffic

Table A.19: Statistics of Ident Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(Bytes)	51.000	9.000	45.256	0.084	63.18
Duration(sec)	19.174	0.244	1.781	1.374	75.58
Bandwidth(KB/sec)	0.197	0.002	0.050	0.730	63.57
Interarrival time(sec)	27820.496	0.122	1202.424	2.273	73.64

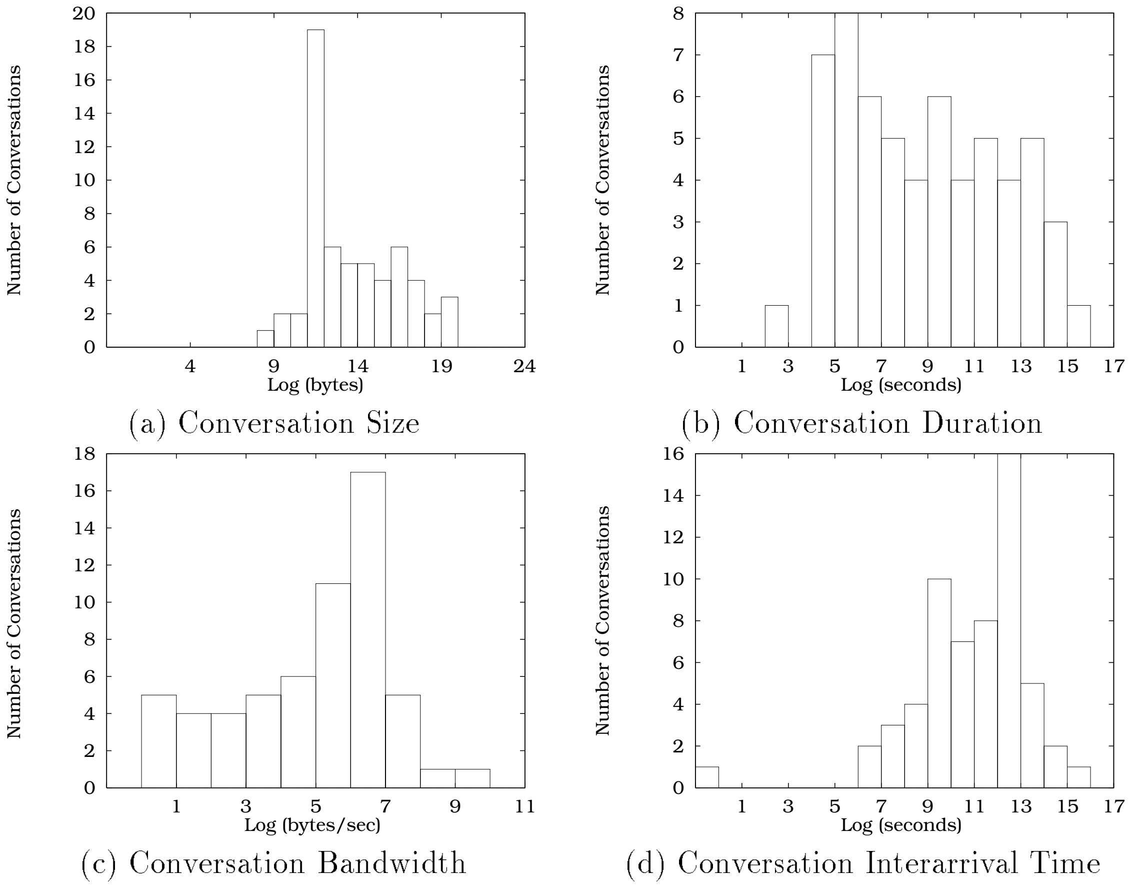

Figure A.19: Characteristics of Login Traffic

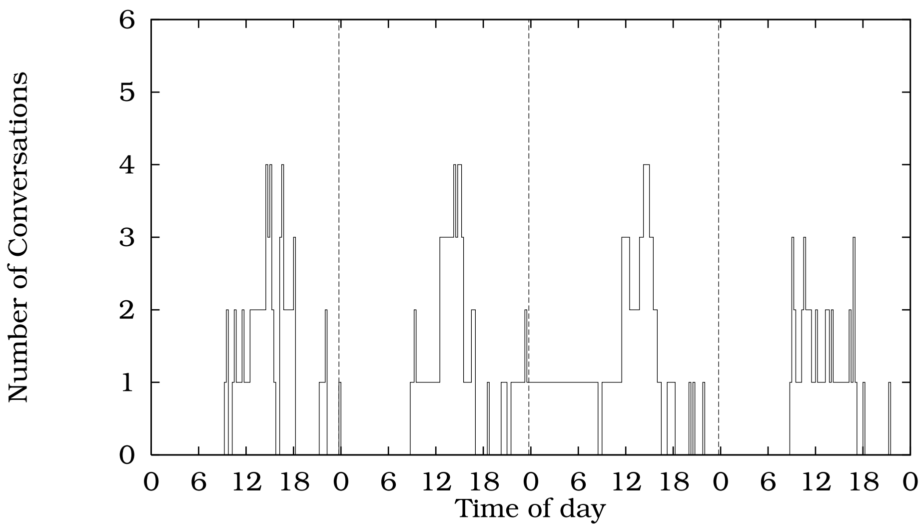

Figure A.20: Number of Active Login Conversations

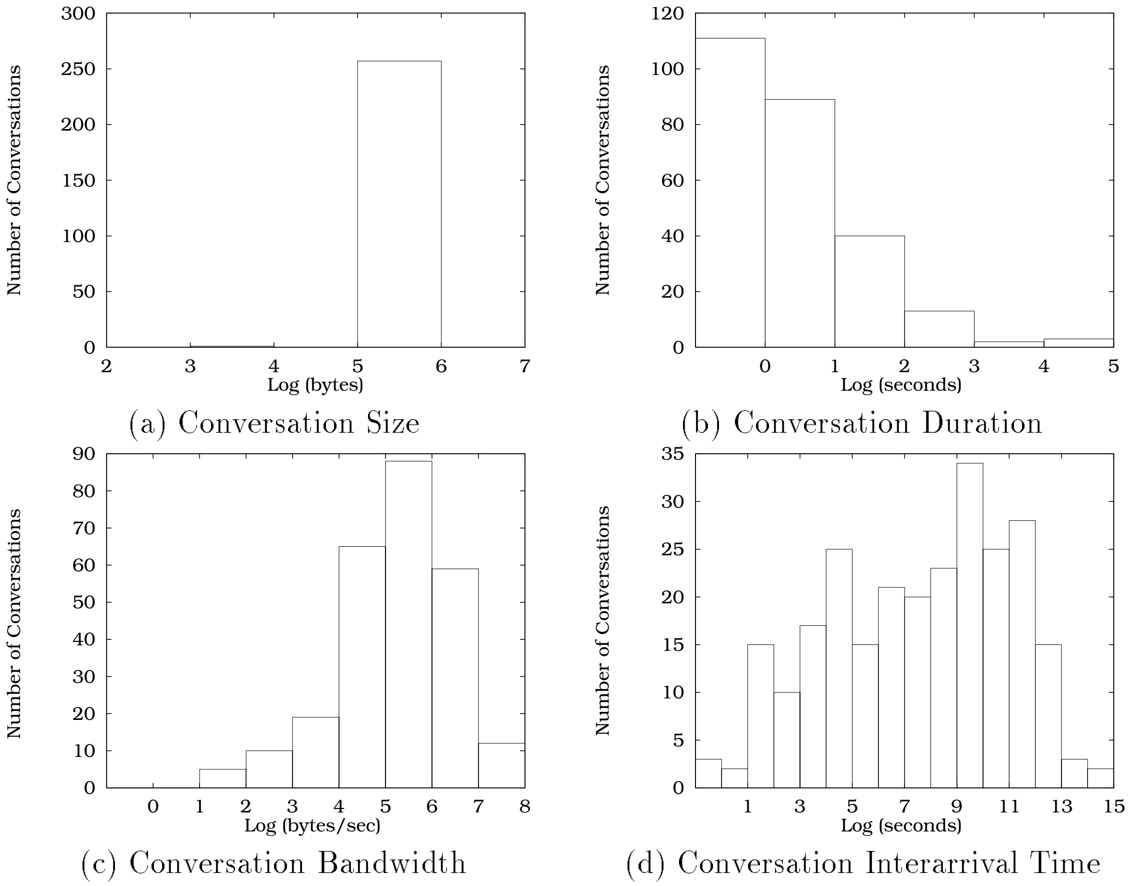

Figure A.21: Characteristics of Ident Traffic

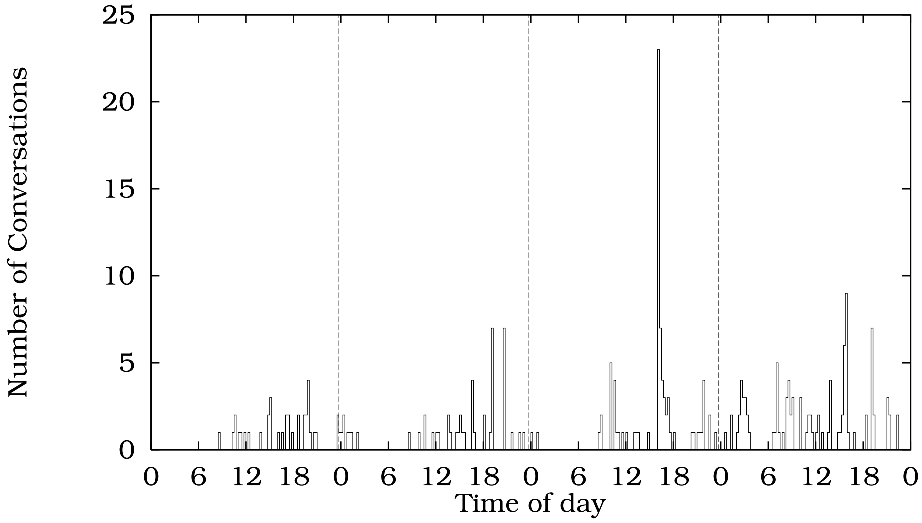

Figure A.22: Number of Active Ident Conversations

A.4.9 WWW Traffic

Table A.20: Statistics of WWW Conversations

	Max	Min	Average	Cov	% smaller than average
Conversation size(KB)	386.355	0.000	14.582	1.846	71.05
Duration(sec)	1656.601	0.260	6.383	7.471	76.07
Bandwidth(KB/sec)	24.303	0.000	3.359	0.923	62.85
Interarrival time(sec)	43726.172	0.089	248.525	8.377	93.59

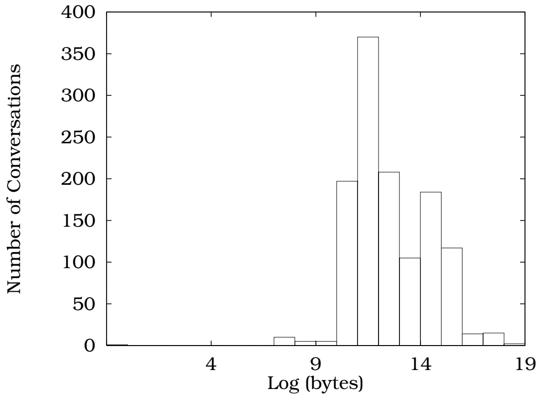

(a) Conversation Size

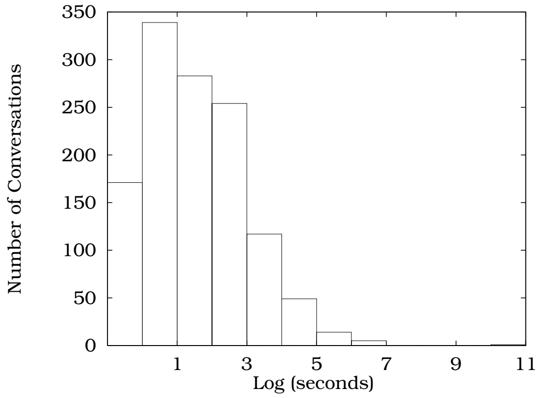

(b) Conversation Duration

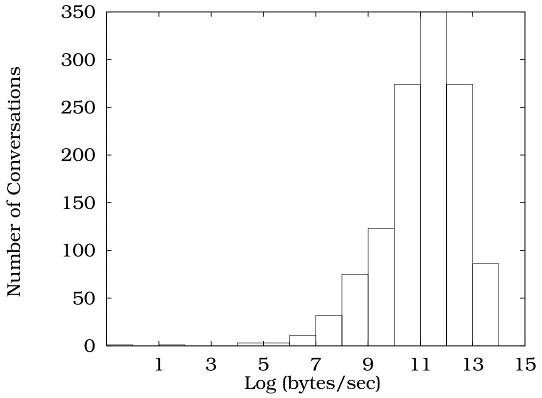

(c) Conversation Bandwidth

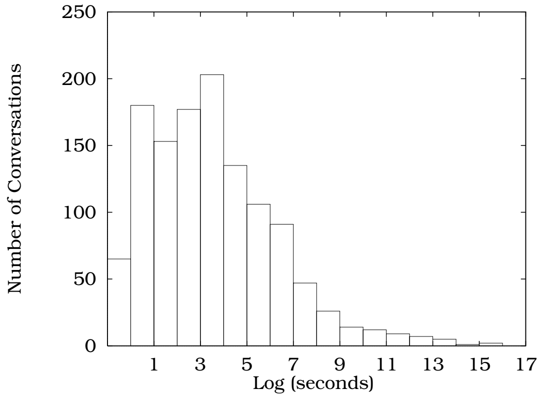

(d) Conversation Interarrival Time

Figure A.23: Characteristics of WWW Traffic

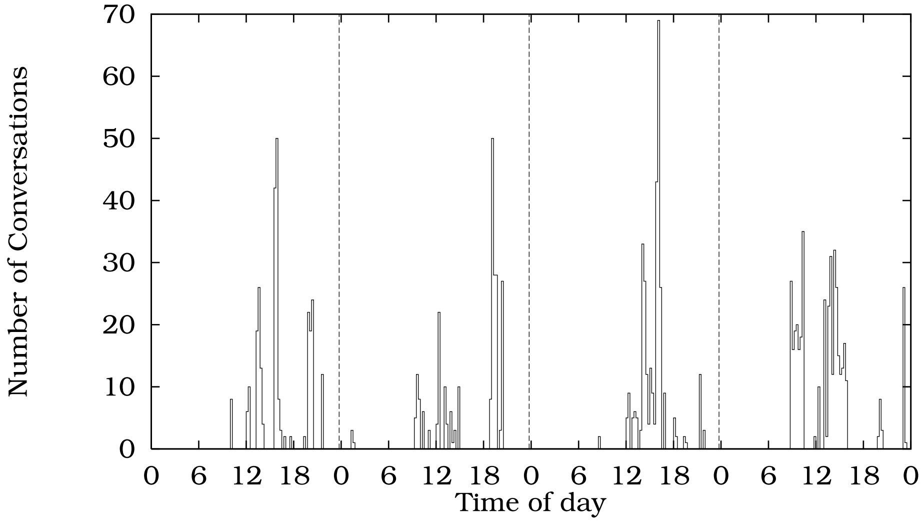

Figure A.24: Number of Active WWW Conversations

A.4.10 Direction of Data Flow

Table A.21: Data Flow From Source To Destination and Vice Versa

Protocol	Mbytes Src to Dst	Mbytes Dst to Src	Ratio S-D:D-S	Kbytes/con Src to Dst	Kbytes/con Dst to Src
Smtp	29.291	14.370	2:1	8.996	4.413
FTP-Data	236.475	0.000		109.327	0.000
WWW	1.364	16.615	1:12	1.106	13.475
Finger	0.003	0.295	1:98	0.006	0.505
Gopher	0.013	2.845	1:219	0.031	6.889
FTP	0.111	0.473	1:4	0.376	1.605
Ident	0.002	0.009	1:5	0.010	0.036
Telnet	0.108	9.237	1:86	0.433	36.949
Login	0.031	4.873	1:157	0.533	82.597

A.5 Summary of Results

Measurement results of wide area Internet conversations involving hosts on the CS Ethernet have been presented. The following results have been identified:

- A total of 98 application protocols were identified. Nine applications (SMTP, FTP-Data, WWW, Finger, Gopher, FTP, Ident, Telnet, Login) were responsible for 96% of the wide area TCP conversations.

- *Skdad*, *skorpio*, and *sparkle* were the hosts on the CS Ethernet that used the Internet the most frequently. Networks *umich.edu*, *uiuc.edu*, and *uu.net* were used the most frequently by hosts on the CS Ethernet.

- 37 highest level Internet domains were identified. 24.7% of all the conversations involved hosts within Canada, 38.26% involved hosts in the EDU domain, and 14.59% involved hosts in the COM domain.

- Conversations with a small number of local area networks dominated since 80% of the conversations involved only 25% of the external LANs. In fact only 10 external LANs (out of a total of 454) were responsible for 22.79% of all the conversations.

- The average setup latency for the conversations initiated by hosts on the CS Ethernet was 0.99 seconds and the average setup latency for the conversations initiated by hosts on the Internet was 0.01 seconds.

- FTP-Data conversations transfer large amounts of data which made up 71% of all wide area data bytes to and from the CS Ethernet.

- Although the browsing tool *Mosaic* for the WWW application was only installed on hosts in this department in December 1993, this application is already responsible for the third most wide area conversations on the CS Ethernet.

- The average overall conversation size was 38 Kbytes, the average duration was 112.86 seconds, the average bandwidth was 1.594 Kbytes/sec, and the average interarrival time was 38.8 seconds. Table A.22 shows the summary results of the conversation statistics of the 9 most frequently used applications.

Table A.22: Statistics of the Nine Most Frequently Used Applications

Protocol	Ave. Size (Kbytes)	Ave. Duration (Seconds)	Ave. Bandwidth (KBytes/sec)	Ave. Interarrival Time (Seconds)
SMTP	13	24	0.584	106
FTP-Data	109	28	3.278	157
WWW	15	6	3.359	249
Finger	1	7	0.347	590
Gopher	7	5	1.206	702
FTP	2	444	0.017	1155
Ident	0.05	2	0.050	1202
Telnet	37	647	0.060	1362
Login	83	3764	0.068	5208

- Longer conversations did not necessarily create the most traffic.
- The data flow of FTP-Data traffic is unidirectional while the data flow of SMTP, WWW, Finger, Gopher, FTP, Ident, Telnet, and Login are bidirectional (but asymmetric).