LBL-30840

Lawrence Berkeley Laboratory
UNIVERSITY OF CALIFORNIA

Engineering Division

Measurements and Models of Wide Area TCP Conversations

V. Paxson
(M.S. Thesis)

May 1991

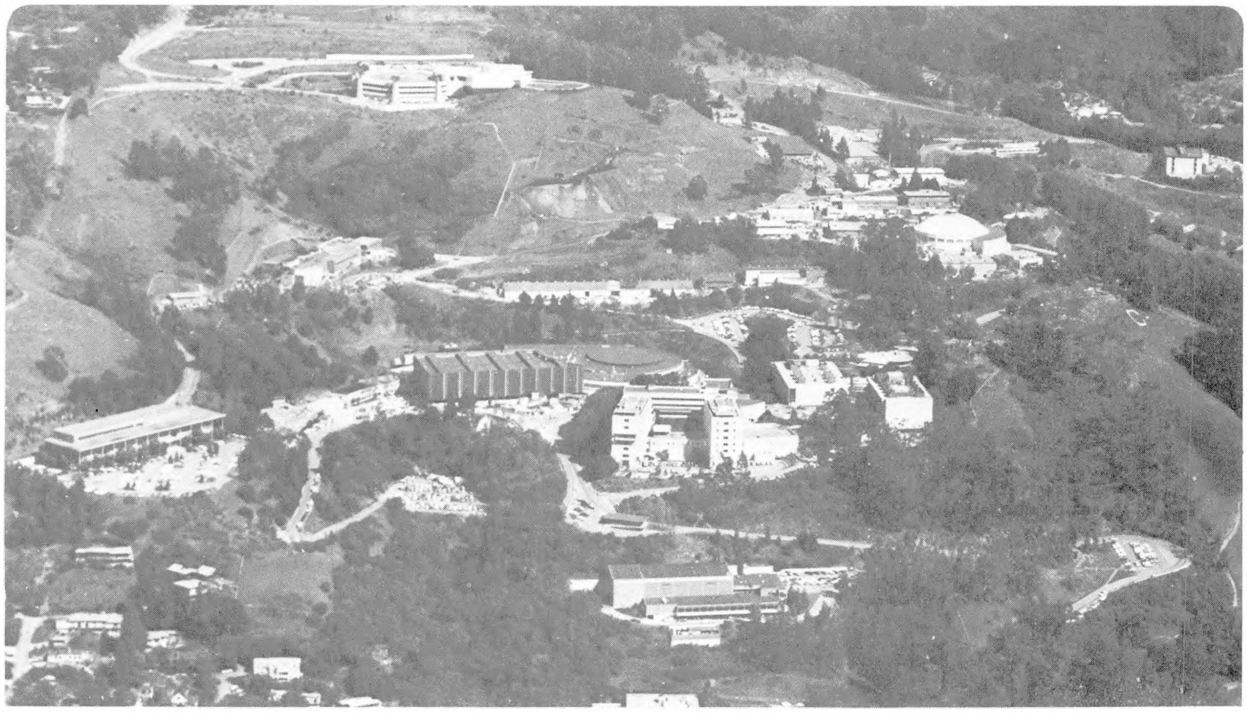

Prepared for the U.S. Department of Energy under Contract Number DE-AC03-76SF00098

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. Neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California and shall not be used for advertising or product endorsement purposes.

Lawrence Berkeley Laboratory is an equal opportunity employer.

LBL--30840

DE91 016424

Measurements and Models of Wide Area TCP Conversations

Vern Paxson

M.S. Thesis

Department of Electrical Engineering and Computer Sciences
University of California

and

Engineering Division
Lawrence Berkeley Laboratory
University of California
Berkeley, CA 94720

May 1991

This work was supported in part by a National Science Foundation Graduate Fellowship and in part by the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

Contents

List of Figures

List of Tables

Abstract

This paper describes measurements of all of the wide area network TCP conversations between the Lawrence Berkeley Laboratory (LBL) and the rest of the world for the months of November, 1990, and March, 1991. Some 500,000 conversations were recorded, encompassing 11 different major protocols. We look at aggregate characteristics of these conversations, both overall and by TCP protocol (e.g., smtp, ftp), computing the distributions of amount of data transferred, network bandwidth used, conversation lifetimes and conversation interarrival times. Temporal traffic variation is also investigated, showing the variation of number of active conversations and network bandwidth utilization over periods of 24 hours, 7 days and 30 days. Long term variation is also investigated by separately analyzing November and March data (which reveals a 10–20% increase in almost all aggregate traffic characteristics in just four months). We classify each conversation geographically and discover that the connectivity of the conversations was remarkably rich, including traffic to 48 of the 50 states in the U.S. and 23 foreign countries. Finally, we develop a number of models for describing conversations of the various protocols. From these models we can more readily assess how each protocol is used and how the use changes as network utilization grows.

1 Overview

This paper describes measurements of all of the wide area network TCP conversations between the Lawrence Berkeley Laboratory (LBL) and the rest of the world for the months of November, 1990, and March, 1991. LBL is a multidisciplinary basic research laboratory that is operated by the University of California under contract with the Department of Energy. The Laboratory has about 3,300 employees and 700 computers attached to the Internet. Also included in our analysis as part of LBL (though administratively the two are separate) is the Mathematical Sciences Research Institute (MSRI), an NSF-funded research institute with typically about 60 associated researchers. MSRI is part of LBL's Internet network.

First we discuss the structure of a TCP conversation and the tools we used for recognizing and recording the information necessary for analyzing these conversations. We next briefly discuss how the raw data we gathered was reduced to a tractable form, and give an overview of the characteristics of the November conversations. Our overview concentrates on the conversations' aggregate properties (such as total number of conversations and average conversation duration), histograms of these properties' distributions with some discussion of extreme values, and a look at how network use varies over the course of a day and the course of a week. We next look at the same aggregate properties for the March conversations and how network use grew between November and March.

We then discuss the geographical distribution of the network conversations, which we were able to derive from the subdomains of the hosts involved in the conversations. We find that the interconnectivity between LBL and the rest of the world is quite rich.

Our final efforts entail developing simple models for describing the behavior of the various protocols. We discuss the methodology we used for developing the models and introduce a metric for the "goodness of fit" of a model to data, followed by a discussion of the models themselves. Our goal is that these models be sufficiently descriptive to serve in the simulation of realistic wide-area network conversations.

2 The Structure of a TCP Conversation

TCP is the Transmission Control Protocol used for reliable communication between Internet hosts [Pos81]. Central to TCP is the notion of a *conversation*: a bidirectional connection between two Internet hosts over which data may be sent without losses or errors. TCP conversations are used with a variety of higher-level *protocols*. Each protocol is designed with a specific networking task in mind, and their use varies widely.

Host A begins a TCP conversation with a remote host B by first picking an initial *sequence number* and sending it to B in a TCP packet with the SYN ("Synchronize sequence numbers") bit set in the header. If B accepts the conversation, it replies to A with a separate sequence number in another SYN packet, acknowledging the sequence number given in A's original packet. The conversation is now established. When A or B wishes to terminate the conversation, it sends a FIN ("Finish") packet to its partner; the partner then replies with another FIN packet, acknowledging receipt of the first one. At this point, the conversation is terminated.[1]

As A and B 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. Thus, after the initial SYN, all sequence numbers are given as the sum of the initial sequence number plus the number of bytes transmitted so far. Therefore the difference between a FIN packet's sequence number and that of the host's earlier SYN packet is equal to the total number of bytes sent by that half of the conversation. Recording SYN and FIN packets along with their timestamps then suffices for extracting for each conversation the onset, duration, and number of bytes sent in each direction.

Also part of the conversation creation process is the selection by A of the TCP *port* numbers to be used by A and B. Two hosts can have multiple conversations as long as the pair $(\text{Port}_A, \text{Port}_B)$ is unique. Several port numbers are predefined for use by well-known protocols. Since the port number is included in all TCP packets, the SYN and FIN packets suffice for extracting the protocol used by each conversation.

3 Data-Taking and Reduction

The data was taken using a packet capture tool called *tcpdump* [JLM89]. One of the packet filters with which it works is the Berkeley Packet Filter (BPF) [MJ91]. When using BPF, *tcpdump* downloads into the operating system kernel a program written in a simple interpretive language. The kernel then runs the program on each network packet received and returns to the user those that the program accepts. Since the user program need only be scheduled when the kernel has buffered up a set of packets known to be interesting, BPF provides a highly efficient and flexible mechanism for packet capture.

All off-site communication at LBL funnels through a group of gateways that reside on a network separate from the rest of the Laboratory. To take the data discussed in this paper we used a Sun 3/50 residing on the gateway network to run *tcpdump*, which captured packets and saved their headers on a local disk. The filter captured all TCP packets in whose headers any of the FIN, SYN, or RST flags were set.[2] Two

[1] In addition, conversations may be abruptly terminated by sending an RST ("Reset") packet.

[2] Additionally, the filter excluded any traffic to or from the gateway network, to avoid including in the data any conversations induced by the measurement process itself.

large uninterrupted periods of data were taken.[3] One spanned the month of November, 1990, and the other the month of March, 1991.

From October 27th through December 1st, the packet filter received a total of 123,703,757 packets, or an average of about 3.5M per day (40 per second). No packets were dropped, by either the kernel or the Ethernet driver. Unfortunately, similar statistics are not available for the March data, except that no packets were dropped by the Ethernet driver. Given our experience with the high packet rates that *tcpdump* can sustain when using it with BPF, we would be surprised if the kernel had dropped any packets during March, either.

From the first mass of filter data we extracted a subset for the month of November, running from midnight, Thursday, November 1st, through midnight, Saturday, December 1st. This thirty day period included Thanksgiving, November 20th; that Thursday and the following Friday were LBL holidays, during which network use fell slightly below weekend levels.

Because we were interested in studying how network usage had changed between November and March, we extracted from the second data set a subset as similar as possible to the first in temporal size and phase. The "March" data thus runs from midnight, Thursday, February 28th through midnight, Saturday, March 30th. This thirty day period did not include a holiday. We discuss how we corrected for this difference in Section 5.

The file of raw packet headers for November consisted of 84MB of binary data; March consisted of 111MB. The data was reduced using a $\approx$ 300 line *awk* script. The reduced data is an ASCII file containing a one-line description of each conversation, giving its starting time, duration, total number of bytes transferred (in both directions), protocol, whether it was initiated locally (within LBL), whether it involved a U.C. Berkeley host[4], the IP addresses of the originator and the responder, and the number of bytes sent by the originator to the responder. From this latter value and the byte total we can derive the number of bytes sent by the responder to the originator.

For example, the first November conversation is:

```
657446427.53 26.57 540 ftp 1 1
      128.3.552.90 128.32.156.99 109
```

This line indicates that at 53/100's of a second[5] after 12:00.27AM on November 1st, 1990, an *ftp* conversation began that lasted 26.57 seconds and during which 540 bytes were exchanged. The **1** in the 5th column indicates that the conversation was initiated at LBL and the **1** in the 6th column indicates that the conversation involved U.C. Berkeley. The originating host's IP address was 128.3.552.90 and the responder's 128.32.156.99. The corresponding hostnames are isotop.lbl.gov and janus.berkeley.edu.

If we check the raw data file for that timestamp we get:

```
00:00:27.53 isotop.lbl.1414 > janus.Berkeley.ftp:
        S 1747136000:1747136000(0) win 4096
```

indicating that at that instant isotop.lbl.gov sent a SYN packet to the *ftp* server on janus.Berkeley.EDU, specifying an initial sequence number of 1747136000. By filtering on those two hosts and the TCP *ftp* protocol, we can run *tcpdump* on a slice of the raw data starting at the given time and lasting for another 26.57 seconds:

```
00:00:27.53 isotop.lbl.1414 > janus.Berkeley.ftp:
        S 1747136000:1747136000 ...
00:00:27.55 janus.Berkeley.ftp > isotop.lbl.1414:
        S 1339712000:1339712000 ... ack 1747136001
00:00:54.10 janus.Berkeley.ftp > isotop.lbl.1414:
        F 1339712431:1339712431 ... ack 1747136109
00:00:54.10 isotop.lbl.1414 > janus.Berkeley.ftp:
        F 1747136109:1747136109 ... ack 1339712432
```

The total number of bytes transferred during the conversation is equal to $(1747136109 - 1747136000) + (1339712431 - 1339712000) = 109 + 431 = 540$, as indicated in the reduced data.

4 A Look at the November Data

We begin with an overview of the basic properties of the November conversations, to give a feel for the different dimensions of the data to be explored. We first look at aggregate properties: number of conversations and their durations, sizes, and averages. We then present histograms of the distributions of these quantities to give a first impression of the variability of the data. We finish with a look at how conversation activity varies over the course of a day and the course of a week.

4.1 Aggregate Conversation Characteristics

The 84MB of raw November data reduced to 14.3MB of processed data. A total of 210,868 complete conversations occurred[6]. 128,137 (60.8%) were initiated within LBL, and 64,659 (30.7%) were with hosts on the U.C. Berkeley network[7]. The conversations involved the exchange of 5.6GB of user data, for a sustained data rate (over the entire month) of 2.2 KB/sec[8], not including TCP/IP overhead.

[3] To avoid overflowing the local disk, periodically a new instance of *tcpdump* was begun and then the old one terminated. The old data file would then be transferred to an LBL file server with larger capacity. Software was developed for splicing together a series of such files, taking care to weed out duplicates from the brief interval during which both the new and old instances of *tcpdump* were running.

[4] LBL and U.C. Berkeley hosts were identified by the network numbers of the IP addresses.

[5] Our timestamps were limited in accuracy to 10 msec.

[6] About 5% of the conversations were terminated abnormally (via RST packets instead of FIN packets); these conversations are not part of the 210,868, nor are they included in the subsequent analysis.

[7] As we shall see, the *nntp* protocol dominates the conversations, and contributes significantly to these percentages.

[8] 5,609MB / (86,400 secs/day × 30 days) = 2,163 B/sec

Protocol	Description
nntp	Network news
smtp	Electronic mail
ftp	File transfer; commands
ftp-data	File transfer; data
finger	Remote user lookup
telnet	Remote login
login	Remote Unix login
shell	Remote command execution
domain	Distributed nameserver
X11	X11 Window System
other	Unidentified

Table 1: TCP Protocol Names and Descriptions

Of the TCP protocols used, we could identify 17 by the port numbers used. The remainder we lumped together as *other*. 11 of these protocols (including *other*) were involved in 1,000 or more conversations. These are briefly described in Table 1. The remainder were each involved in 200 or fewer conversations, and we refrain from discussing them individually as our sample for them is small and their impact on the aggregate network traffic is low[9].

The bulk characteristics of each TCP protocol's November conversations are given in Table 2. Each row lists the protocol followed by the total number of completed conversations, the total number of megabytes exchanged in such conversations, the average data rate in bytes per second, the average conversation size in kilobytes, the average conversation duration in seconds, the percentage of conversations originating within LBL, and the percentage of conversations that were between LBL and U.C. Berkeley. Some notes on the numbers:

- The **all** entry gives the totals for all protocols, including the 7 minor protocols not listed in the table.
- The average overall conversation bandwidth of 127 B/sec, coupled with the total conversation bandwidth of 2.2 KB/sec over the entire month, gives an average of 17 active conversations at any particular moment.
- That 75% of all *nntp* (network news) conversations originate at LBL is due to the *nntp* routing algorithm (flooding) coupled with the fact that the main *nntp* server at LBL is very aggressive about forwarding news articles, doing so whenever it receives a new batch.
- TCP *domain* conversations occur between name servers. LBL name servers do not have U.C. Berkeley name servers as peers, which explains the nearly complete lack of *domain* conversations with U.C. Berkeley hosts.
- When exploring the *other* conversations we noticed that some of them are related to executing remote commands using *shell*, while others are apparently used to exchange (sometimes very large) files with supercomputer centers at Livermore and U.C. San Diego.
- Some *shell* conversations are related to use of the Unix *rcp* command (remote file copy), which explains the high bandwidth.

Figure 1 shows a histogram of the number of conversations whose size is in a particular range. The ranges begin with 1 byte, 2-3 bytes, 4-7 bytes, 8-15 bytes, and continue in that fashion, each being twice as large as the previous. From the graph we see that conversation sizes peak in two ranges: 48,744 of the conversations were between 129 and 256 bytes in size; 34,716 ranged from 1025 to 2048 bytes; and 25,046 from 2049 to 4096 bytes. The final histogram bin (not legible in the figure) shows 4 conversations of size between 33.5MB and 67MB. These were an *ftp-data* conversation of 54MB (to Los Alamos National Laboratory), and three *X11* conversations, of 44MB, 42MB, and 38MB, all to U.C. Berkeley. All in all, 77 conversations included 10MB or more of data: 41 of these involved *ftp-data*, 29 were *X11*, 4 were *other*, and 3 were *shell*.

Figure 2 shows a similar histogram for the number of conversations of a particular duration. Here the range is $\log_2$ seconds instead of bytes. The biggest peak is between 65 seconds and 256 seconds, with a total of 67,919 conversations in those two bins. 24,436 conversations lasted less than 1 second. 6 conversations occupied the top two bins, from 3 to 12 days. The longest of these was a *login* conversation between U.C. Berkeley and LBL that lasted 6 days, 18 hours, during which 3.4MB were exchanged. 73 conversations lasted 12 hours or longer. Of these, 39 were *login*, 20 were *telnet*, 5 were *X11*, 4 were *ftp*, 3 were *finger*[10], and 2 were *other*.

Figure 3 shows the same sort of histogram for the number of conversations that began a particular amount of time after the previous conversation. Again, the ranges are seconds, though now the first bin is from 1/16th second, the second 1/8th second, the third 1/4 second, and so forth. While conversation interarrival times peak in the 4 to 16 second range (a total of 66,428 conversations), a sizeable number of conversations—14,706—began less that 1/16th of a second after the previous conversation. It turns out that 87% of these are *nntp* conversations, a consequence of the aggressive flooding strategy used for news propagation. In 5 instances more than 512 seconds (8.5 minutes) elapsed between conversation starts. These all occurred between 1:30AM and 3AM, on three separate nights.

Figure 4 shows the same type of histogram for the number of conversations of a particular bandwidth. The bins are bytes

[9] These protocols and their total number of conversations were: *whois*, 196; *printer*, 26; *uucp-path*, 4; *time*, 1; *systat*, 1; *hostnames*, 1; *daytime*, 1.

[10] It is surprising that a *finger* conversation might take more than a few minutes. We conjecture that the long conversations may be due to the *finger* server waiting for a file server to recover in order to retrieve user information from it.

Protocol	Total Conv	Total MB	Avg. B/sec	Avg. KB	Avg. sec	% LBL orig.	% w/ UCB
all	210,868	5,609	127	26.6	209	61	31
nntp	91,426	991	113	10.8	96	75	37
smtp	49,046	168	268	3.4	13	41	21
ftp-data	29,061	2,261	4,713	77.8	17	50	18
finger	10,232	6	16	0.5	34	46	41
telnet	9,170	291	20	31.7	1,610	68	37
domain	5,969	33	1,211	5.5	5	80	0
ftp	5,329	7	3	1.3	420	49	29
login	4,951	167	12	33.7	2,735	62	71
other	1,964	198	288	100.6	349	15	23
X11	1,929	1,061	417	550.0	1,318	85	85
shell	1,561	426	4,811	272.7	57	75	57

Table 2: Bulk Characteristics of November TCP Conversations By Protocol

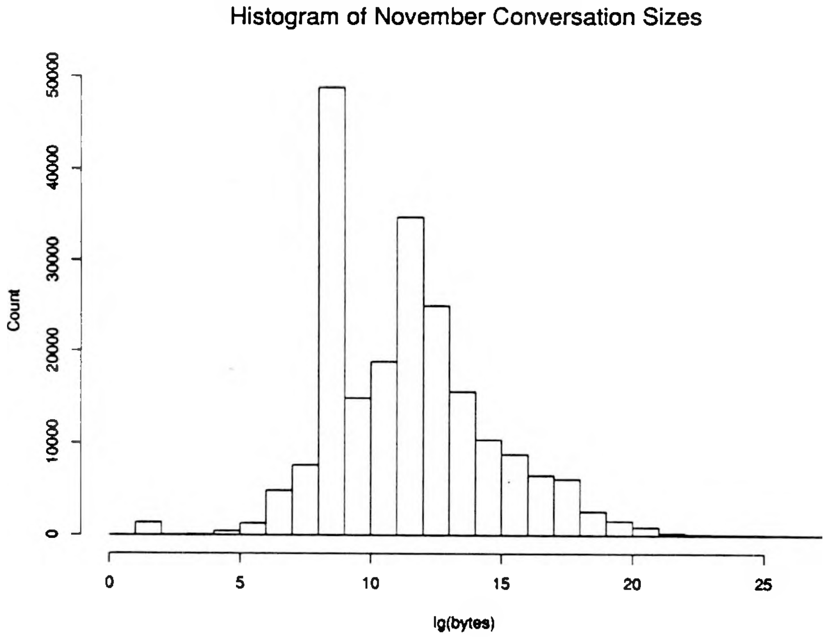

Figure 1: Logarithmic Distribution of Conversation Sizes

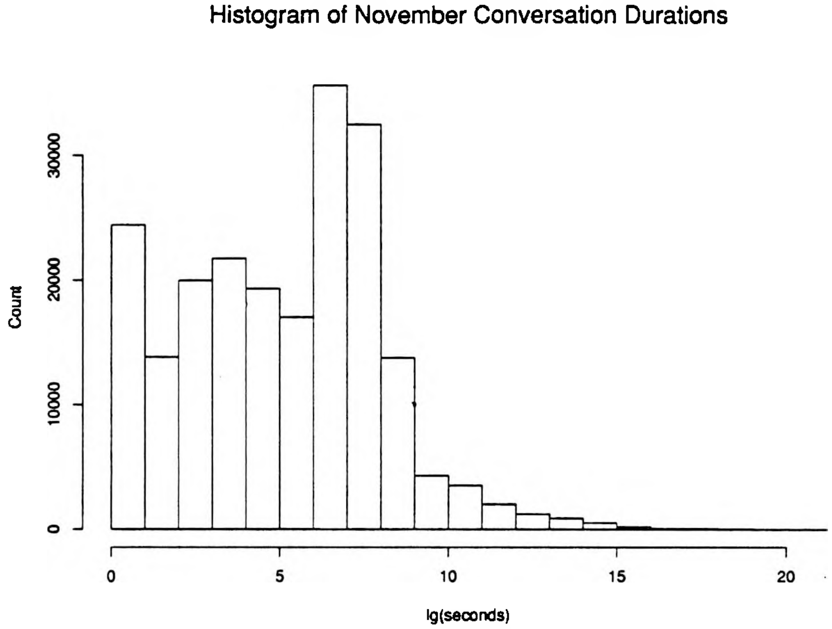

Figure 2: Logarithmic Distribution of Conversation Durations

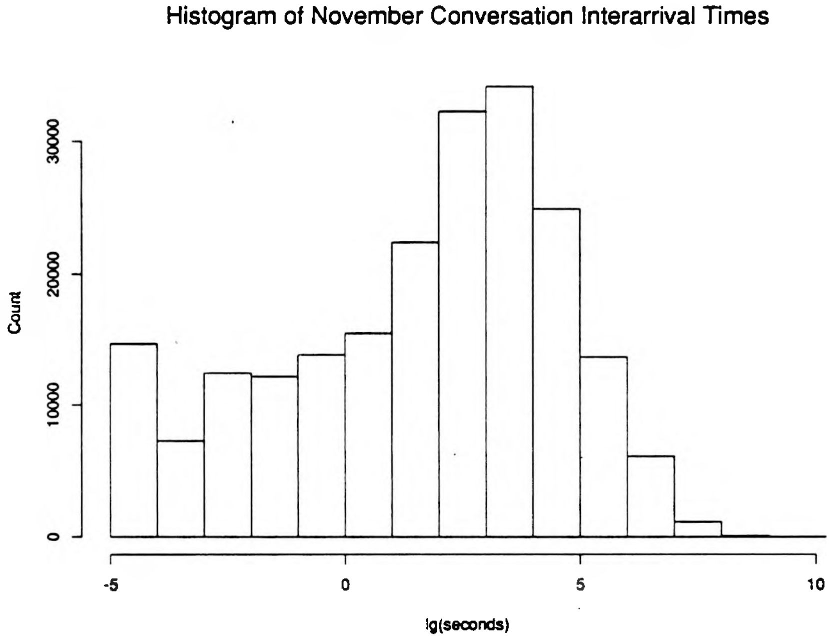

Figure 3: Logarithmic Distribution of Conversation Interarrival Times

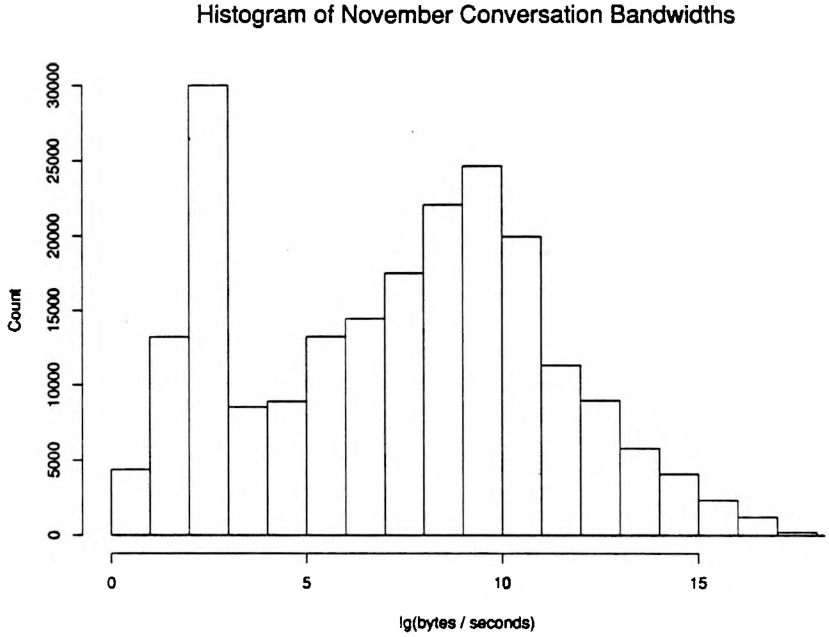

Figure 4: Logarithmic Distribution of Conversation Bandwidths

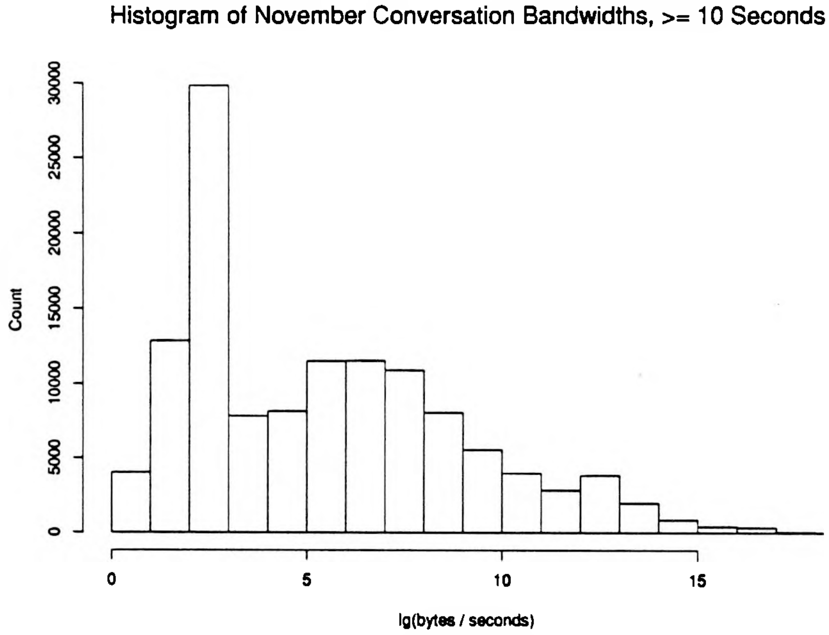

Figure 5: Distribution of Conversation Bandwidths, >= 10 Seconds

per second. The maximum bandwidth observed was 112.6 KB/sec (*ftp-data*). 544 conversations achieved a data rates of 50 KB/sec or greater. Of these, 303 were *ftp-data* and the other 241 were *shell* (most likely the *rcp* program).

Figure 5 shows the same histogram but restricted to conversations lasting 10 or more seconds[11]. The maximum sustained bandwidth drops to 81 KB/sec. 280 conversations sustained a bandwidth of 40 KB/sec or greater. Of these, there were 149 *ftp-data*, 125 *shell*, and 6 *other*.

While these histograms give an overview of different dimensions of the conversation data, they shed no light on the underlying causes. We examine the properties of the individual protocols that lead to these aggregates in section 7.

4.2 Dynamic Conversation Characteristics

The preceding histograms illustrate some of the static characteristics of the November conversations, but there is interesting temporal variation in the conversations, too. Figure 6 shows the pattern of active conversations for the week of November 11th (Sunday) through November 17th (Saturday), inclusive. Each point on the plot represents the number of conversations active at a particular time. The figure shows that conversation activity is greatest during daylight and the evening hours before midnight, attaining peaks from 40 to 60 active conversations during afternoons. Activity during the weekends (the leftmost and rightmost portions of the figure) is roughly half of that during the week. The appendix shows the same plots for the weeks of both the November and the March datasets. The effects of the Thanksgiving holiday (November 22) are readily apparent in those plots, extending from the previous Wednesday to the following Monday (on Monday the activity peak is broader and has 8 to 14 fewer active conversations than on other November Mondays); see Figures 29 and 30 in the appendix.

Figure 7 shows the active conversations for a single day, Wednesday, November 14th. The pattern is typical of work days. A rise begins around 8AM and reaches a peak almost exactly at noon. Activity declines to a nadir around 12:30PM and then reaches a new peak around 1:30PM. Afternoon activity stays high, reaching its greatest value at 3PM, and declines rapidly from 5PM to another nadir at 7:30PM. At this point evidently workers return from dinner and activity resumes, reaching a secondary peak at 11:15PM (the time of this peak varies from 9PM through midnight), and then declines to a low value until the following day's 8AM rise.

One would expect the bulk of the sustained conversations to be due to interactive conversations. This is somewhat borne out by Figure 8, which shows active conversations for interactive protocols (namely, *telnet*, *login*, *X11*, and *ftp*). The activity pattern is clearly similar to that for all conversations. One might expect that the non-interactive traffic would have very few simultaneous conversations, but Figure 9 shows that this is not the case. Here the peaks of activity reach 16 simultaneous conversations around noon, 15 at 4:20PM, and 11 at 11:15PM. This result is surprising in light of the fact that of the non-interactive conversations, only 5 were 30 minutes or longer and only 3 were an hour or longer (9.2, 4.4, and 2.4 hours). Thus the bulk of the simultaneous conversations are simply the sustained overlap of a number of short-lived conversations.

We suspected that the simultaneous non-interactive conversations were dominated by *nntp*, since the flooding propagation of new news results in LBL's newserver beginning new network conversations with all its peers whenever new news comes in. Furthermore, when done forwarding, it keeps the network connection open for 60 more seconds, hoping that it will receive more news to propagate. This forwarding strategy results in a large number of simultaneous *nntp* conversations, as seen in fig. 10. The striking decrease of conversations between 8PM and 10:30PM is a feature that is present during other days as well, though the time and width of the dip varies. Such dips might be caused by factors internal to LBL (such as the local *nntp* server running out of disk space and refusing new news) or external (perhaps a main newsfeed peer being down).

In Figure 11 we see the total bandwidth used by all active conversations on Wednesday, November 7th. Spikes indicate short, high-bandwidth conversations, while white rectangles such as that at the righthand edge of the plot indicate long-lived high-bandwidth conversations. Many of the spikes are due to conversations that do not involve a great amount of data but are so short that they nevertheless have very high data rates. We remove these by filtering out conversations whose duration is less than 10 seconds (for this day, the filtering removes 49% of the conversations). Figure 12 shows the results.

Clearly a very busy conversation or series of conversations is active during the period from about 6PM to 8PM, as there is a long, sustained bandwidth utilization then. Further investigation shows that there were three separate long-lived, high-bandwidth conversations, the first beginning at 6:00PM and lasting for 9 minutes, the next beginning 23 seconds after the first finished and lasting another 26 minutes, and finally a 64 second conversation beginning 2 minutes after the second ended. What's striking about these conversations is that they all shared nearly identical, very high bandwidths: about 24 KB/sec[12]. Further investigation showed that all three were *X11* conversations between the same U.C. Berkeley and LBL hosts, explaining the correlations between one conversation ending and another beginning. One of LBL's network specialists quickly identified the signature of a sustained, very high bandwidth *X11* conversation: a game of *xtrek*! An even more furious X11 conversation occurred in March: a 65 minute conversation between LBL and U.C. Berkeley during which

[11] 59% of the conversations lasted 10 or more seconds.

[12] 24 KB/sec sustained over 26 minutes is 37.4MB!

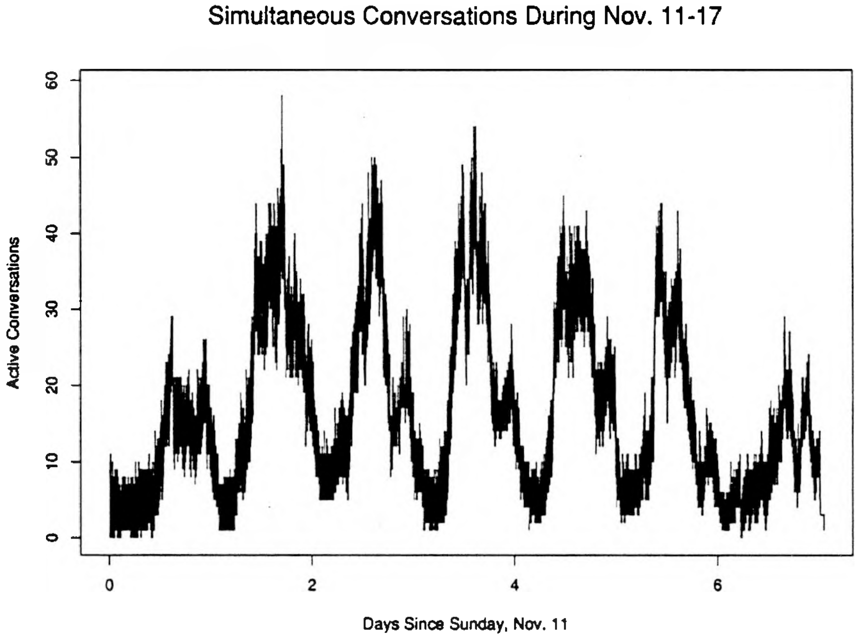

Figure 6: Active Conversations During the Week of November 11-17

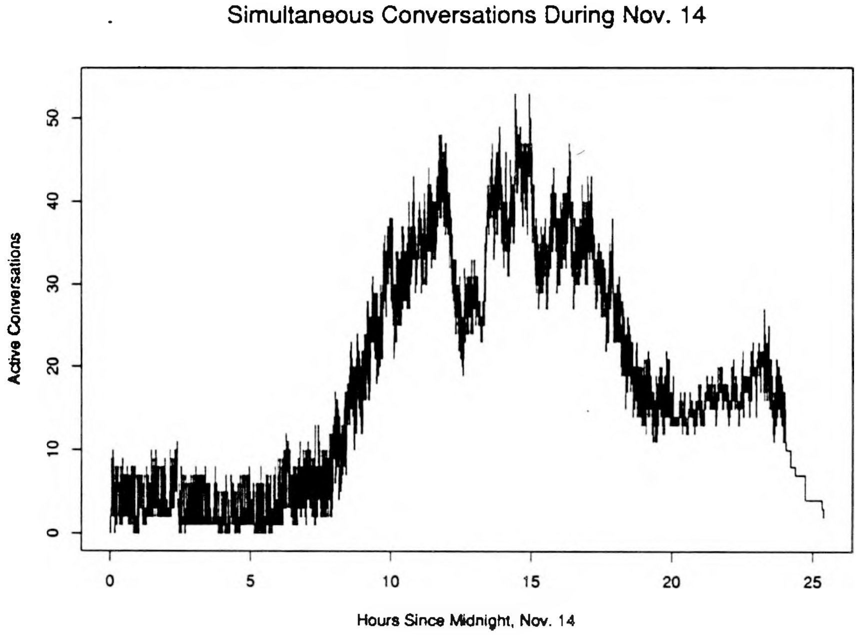

Figure 7: Active Conversations During Wed., Nov. 14, All Protocols

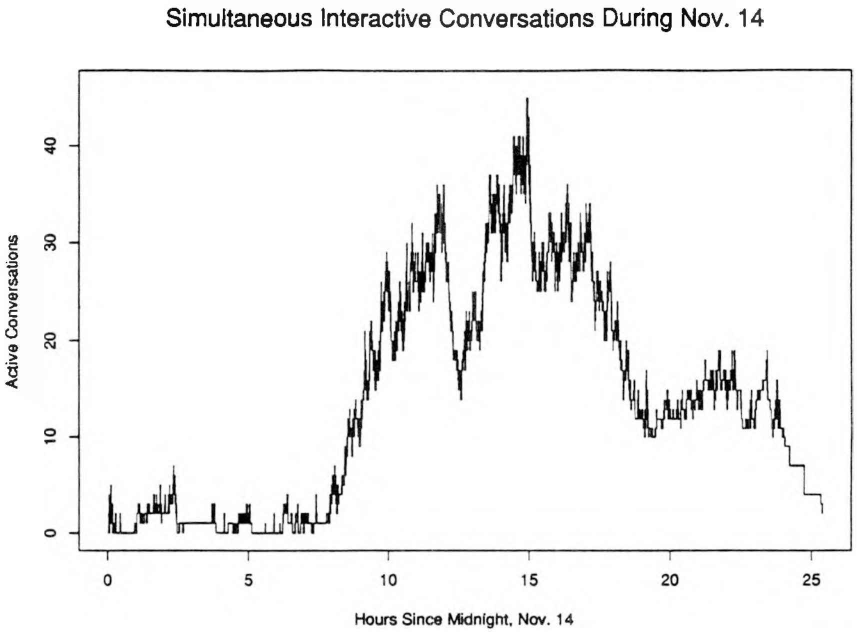

Figure 8: Active Interactive Conversations During Wed., Nov. 14

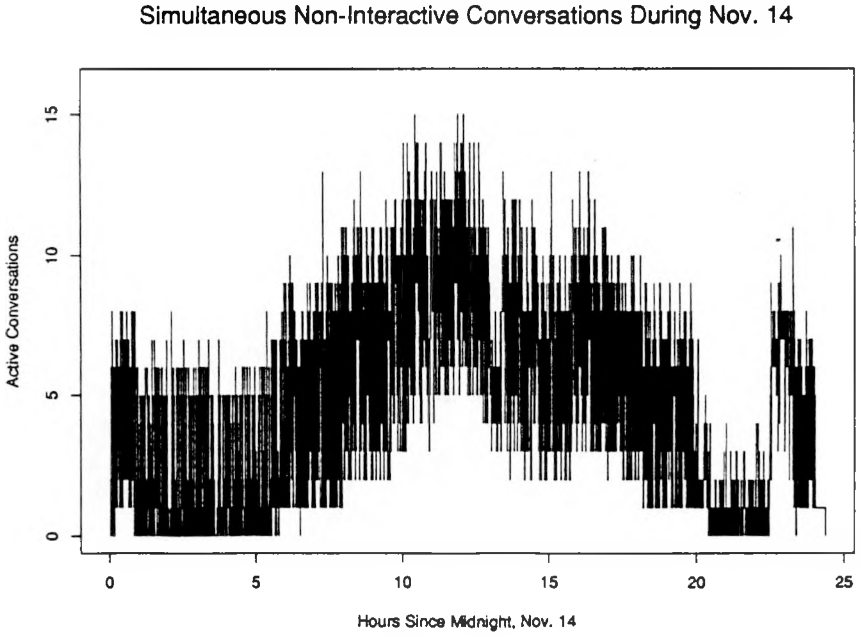

Figure 9: Active Non-Interactive Conversations During Wed., Nov. 14

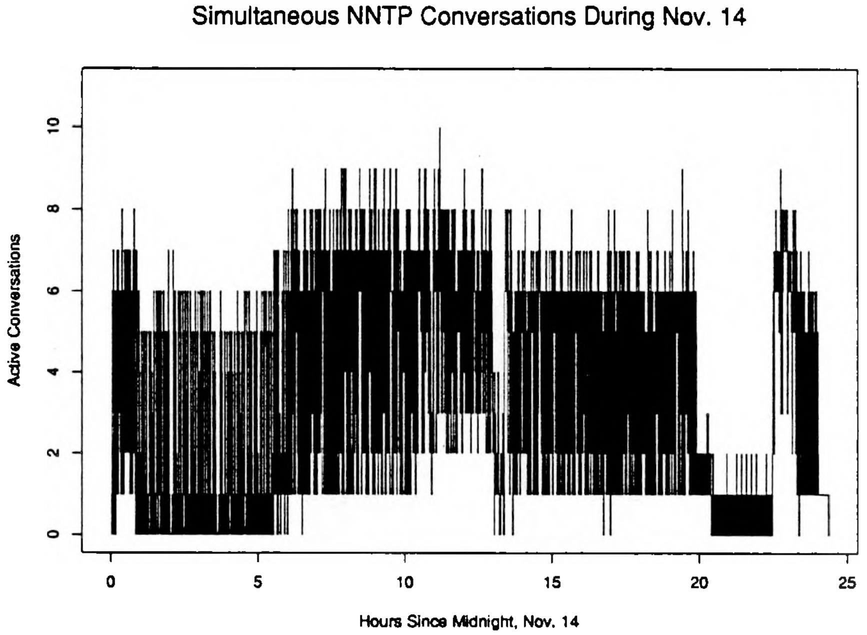

Figure 10: Active *nntp* Conversations During Wed., Nov. 14

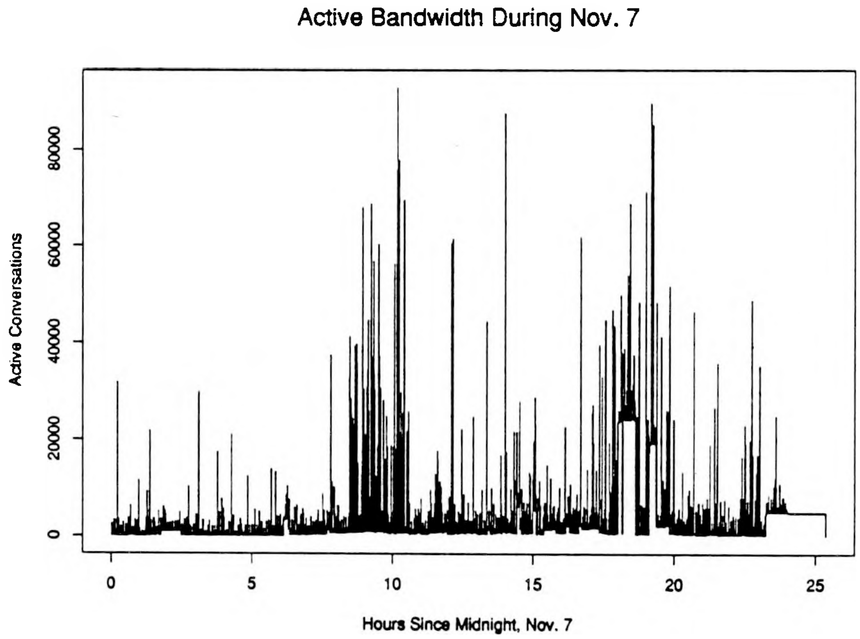

Figure 11: Active Bandwidth (B/Sec), Wed. Nov. 7, All Protocols

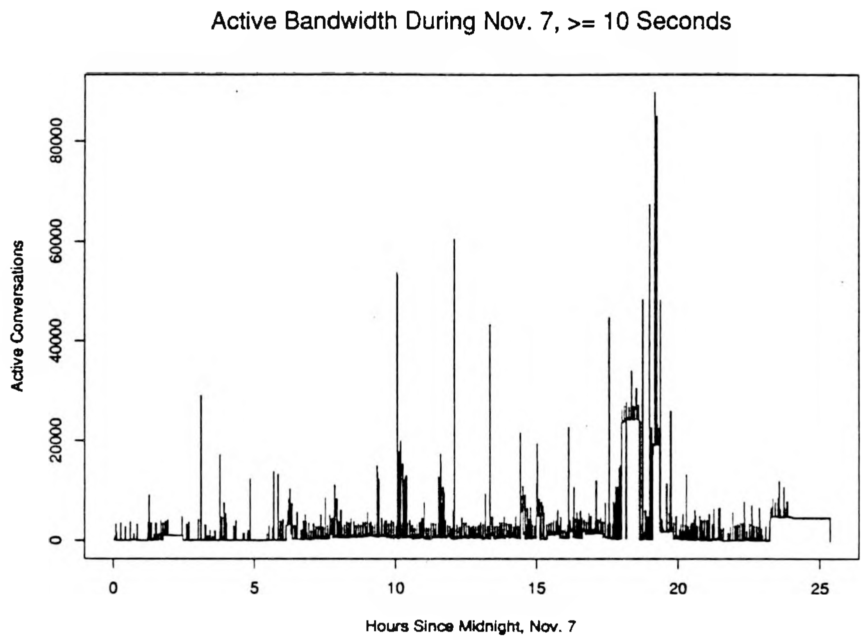

Figure 12: Active Bandwidth (B/Sec), Wed. Nov. 7, >= 10 Seconds

a whopping 144MB was exchanged, for a sustained bandwidth of 37 KB/sec. This was the single largest conversation observed during the two months.

5 The March Data

We now present an overview of the March data with an emphasis on what it indicates about the change (primarily growth) of network usage between November and March.

The 111MB of raw March data reduced to 19.4MB of processed data. A total of 286,868 complete conversations occurred, an increase of 36% over November (but see below). 16 distinct TCP protocols were identified, and again the remainder were lumped together as *other*. The same 11 protocols as before were involved in 800 or more conversations. The remainder except for *printer* were each involved in 325 or fewer conversations[13].

The bulk characteristics of the March conversations are given in Table 3. Perhaps the most striking change between November and March is the increase in *printer* conversations from 26 to 10,018! Of these, 9,973 (99.6%) were between the same LBL originating and U.C. Berkeley responding hosts. Of these, all but 18 entailed the LBL host sending just 7 bytes. These conversations recurred very nearly exactly one minute apart, almost continuously from March 21st through the end of the dataset, midnight March 30th.

From the system manager of the LBL host we learned the source of the conversation pattern: users had set up an account whose sole purpose was to show the printer queue on the U.C. Berkeley host every minute (a request that entails sending 7 bytes) and had left this task running continuously.

While we have found that the legitimate business of a single host can often considerably skew the overall picture of network utilization (this is most true with the *nntp* server), the *printer* anomaly seems more accidental than one of intent. It is not likely to be repeated (at least by the same parties) so it seems prudent to regard it as an abnormality and remove it from the remainder of our analysis. Therefore the **all*** entry in Table 3 reflects the totals for all the protocols except *printer*.

7,035MB over 30 days gives a bandwidth of 2.7 KB/sec sustained over the entire month. Coupled with the average overall conversation bandwidth of 124 B/sec this gives an average of 22 active conversations at any particular moment, a 27% increase over November.

Table 4 summarizes the changes between bulk characteristics of the November data and those for the March data. Values given as a percentage are the percentage change with respect to the November data. For example, there were 31.3% more total conversations in March than in November[14]. Values given within parentheses are changes in percentages. For example, in March the percentage of all conversations that originated within LBL dropped by 9 percentage points.

As mentioned earlier, before interpreting these changes we need to adjust for the fact that a two-day holiday occurred in November but there was no analogous Laboratory holiday

[13] These protocols and their total number of conversations were: *whois*, 325; *sunrpc*, 55; *hostnames*, 32; *uucp-path*, 5; *daytime*, 1. Except for *daytime* and *uucp-path*, all of these represent substantial increases over the corresponding November totals.

[14] The **all** entry corresponds to the differences between November's **all** entry and March's **all*** entry.

Protocol	Total Conv	Total MB	Avg. B/sec	Avg. KB	Avg. sec	% LBL orig.	% w/ UCB
all	286,868	7,038	124	24.5	198	54	40
all*	276,850	7,035	124	25.4	206	52	38
nntp	125,036	1,572	143	12.6	88	55	54
smtp	65,019	259	318	4.0	13	45	20
ftp-data	34,883	3,818	5,906	109.5	19	48	19
finger	15,249	8	11	0.6	51	58	44
telnet	12,063	397	20	32.9	1,667	65	37
printer	10,018	3	819	0.3	0.4	99.8	100
domain	7,266	36	5,949	4.9	0.8	63	0
ftp	7,121	8	3	1.2	352	50	24
login	5,725	211	11	36.9	3,228	55	72
shell	1,755	297	727	169.5	233	66	81
other	1,514	70	60	45.9	766	51	27
X11	801	332	343	415	1,210	49	51

Table 3: Bulk Characteristics of March TCP Conversations By Protocol

Protocol	Total Conv	Total MB	Avg. B/sec	Avg. KB	Avg. sec	% LBL orig.	% w/ UCB
all	+31%	+25%	-3%	-4%	-1%	(-9)	(+7)
nntp	+37%	+59%	+27%	+16%	-8%	(-20)	(+17)
smtp	+33%	+54%	+16%	+16%	0%	(+4)	(-1)
ftp-data	+20%	+69%	+26%	+41%	+12%	(-2)	(+1)
finger	+49%	+33%	-40%	-11%	+50%	(+12)	(+3)
telnet	+32%	+36%	0%	+4%	+4%	(-3)	(0)
domain	+22%	+9%	+460%	-10%	-84%	(-17)	(0)
ftp	+34%	+14%	+2%	-14%	-16%	(+1)	(-5)
login	+16%	+26%	-7%	+9%	+18%	(-7)	(+1)
shell	+12%	-30%	-85%	-38%	+309%	(-9)	(+24)
other	-23%	-65%	-79%	-54%	+119%	(+36)	(+4)
X11	-58%	-69%	-18%	-25%	-8%	(-36)	(-34)

Table 4: Changes in TCP Conversations Between November, 1990 and March, 1991

during March. We adjust our interpretations as follows. As mentioned above, from figures 29 and 30 in the appendix we can see that the effects of the Thanksgiving holiday extended noticeably from the Wednesday prior to Thanksgiving through the following Monday. During that interval, 34,824 conversations occurred. For similar intervals beginning November 7th and November 14th the number of conversations were 43,783 and 42,744, respectively. Therefore we can estimate that between 8,000 and 9,000 conversations were "lost" due to the holiday, about 4% of the total. Therefore we should assume that any aggregate value that grew by less than +4% between November and March is not a significant increase, and the other increases should be interpreted as being a little lower than they initially appear. Any decrease in an aggregate value is significant. Finally, values that are averages as opposed to aggregates are directly comparable without needing to adjust for the +4% difference.

There are a number of points of interest in Table 4:

- For all but *X11* and *other* protocols, the number of conversations increased significantly.

- Surprisingly, the use of *X11* fell precipitously. We contacted the system manager of November's most popular LBL *X11* host, involved in 323 conversations. In March this host participated in only 24 conversations, an astonishing drop. All of the conversations, both in November and March, were with hosts at U.C. Berkeley. By consulting the "last" logs of who had logged on when, the system manager determined that a U.C. Berkeley undergraduate had generated the bulk of the traffic. This person was engaged in a class project that ended in December, 1990. We do not know, however, if similar factors are behind the decrease in *X11* conversations involving other LBL hosts, though we note that in November 85% of *X11* conversations were with U.C. Berkeley, while only 49% during March.

- Most protocols experienced a growth in the amount of data transferred as well. For the three most common protocols—*nntp, smtp,* and *ftp-data*—the increase in bytes transferred grew faster than the number of conversations, indicating that these conversations are getting larger as well as more numerous.

- We conjecture that some of the changes in the *finger* conversations are in part due to the availability of a new *finger* weather-reporting service. 332 of the March conversations were with a single University of Washington host[15], which given a pseudo-username such as "weather-SFO" returns a weather forecast for the associated region (San Francisco Bay Area, in this case). While these conversations made up only 2.2% of the total, they account for 16% of the total bytes transferred. Their average bandwidth was 203 B/sec, far higher than the 16 B/sec for all *finger* conversations.

- The large changes in the *shell* and *other* protocols are difficult to analyze since both of these protocols are used for a wide variety of activities.

Plots showing the simultaneous active conversations for the four weeks of the March data are given in the appendix. The patterns are quite similar to November's, but the peaks are higher. November's peaks were 57, 58, 54, and 55 simultaneous conversations. March's are 67, 73, 68, and 68. The 23% increase is comparable to 31% increase in conversations.

6 Geography—How Wide is 'Wide'?

A natural question arising when studying wide area networks is what does "wide" mean? How are the remote peers in conversations distributed geographically? We're not surprised to find that between 30-40% of all LBL conversations are with U.C. Berkeley, given the very close ties between the two institutes, but what other geographic-related correlations might we find?

We set about identifying the location of each remote host as follows. First, we attempted to directly identify each host using the name server. This succeeded more than 90% of the time. For geographical information, though, subdomains usually suffice, and with some persistence we were able to identify the subdomains associated with all but 1 of the remaining unidentified hosts[16]. This was done by extracting the host's network number from its IP address and identifying the corresponding subdomain using the *whois* network service [HSF85].

7,277 distinct hosts took part in conversations during November or March[17].

The Internet domains involved in the conversations give an overview of the different types of hosts with which LBL communicated. Table 5 lists the counts of how many conversations involved a given domain in November and in March, as well as the percentage increase between the two. The *edu* domain is comprised of educational institutes, of which U.C. Berkeley is a member (the second line gives the counts for the *edu* domain less those conversations that were between LBL and U.C. Berkeley). The *gov* domain is comprised of institutes with ties to the federal government, of which LBL is a member. *mil* is the domain for military institutes; its

[15] `stormy.atmos.washington.edu`; this service has since been discontinued.

[16] The renegade unidentified IP address is `192.31.95.10`. It took part in 8 conversations.

[17] The November conversations spanned 4,262 distinct hosts, of which 407 were within LBL and 586 within U.C. Berkeley. The March data consisted of 5,407 distinct hosts, 457 within LBL and 627 within U.C. Berkeley. Thus the total number of distinct hosts grew by over 26%, though the distinct LBL hosts grew by 12% and the U.C. Berkeley hosts by 7%. As of this writing, LBL has about 680 hosts. U.C. Berkeley has an estimated 4,000-4,500 non-PC hosts [Fro91].

Domain	November	March	Increase
edu	148,325	208,193	+40%
edu except UCB	83,666	92,043	+10%
gov	30,291	36,748	+21%
mil	11,230	14,773	+32%
com	10,583	13,435	+27%
foreign	7,815	10,568	+35%
net	2,205	2,679	+21%
org	326	436	+34%
other	93	36	-61%

Table 5: Number of Conversations *vs.* Domain

Domain	Total	November	March
edu	597	234	263
gov	60	27	28
mil	28	10	9
com	1,744	112	124
net	54	14	17
org	149	19	21

Table 6: Number of Different Subdomains For Each Domain

Country	November	March
Canada	2,329	2,376
Austria	45	70
Belgium	0	3
Denmark	132	69
Finland	188	621
France	1,204	377
Germany	1,012	1,009
Greece	2	53
Iceland	157	188
Ireland	0	11
Italy	89	475
Netherlands	375	343
Spain	0	60
Sweden	412	280
Switzerland	264	1,035
United Kingdom	626	603
Europe	3,880	(+18%) 4,594
Australia	937	1,090
New Zealand	83	125
Australia	1,020	(+19%) 1,215
Japan	706	897
South Korea	116	600
Asia	822	(+82%) 1,497
Argentina	1	0
Mexico	18	997
Israel	0	1

Table 7: Number of International Conversations

counts would be much lower were it not for the fact that one of LBL's *nntp* peers resides in this domain. Without the peer, the counts would be 2,203 and 2,032, respectively, for an increase of -8%. *com* is the domain for commercial entities; IBM, for example, is in this domain. The domain labeled "foreign" consists of international traffic; we discuss it in more depth shortly. The *net* domain consists primarily of institutes doing work related to computer networking. The *org* domain is for non-commercial organizations such as users' groups. Finally, the domains listed under "other" are *bldrdoc* (the National Institute of Standards and Technology) and *arpa*.

It is interesting to see the degree of "penetration" into a domain by LBL hosts, in the sense of how many different subdomains in a given domain engaged in conversations with LBL. Table 6 lists for each domain the total number of subdomains that are registered for that domain in the *whois* database. Subdomains usually correspond to a single entity. For example, all the U.C. Berkeley hosts are in the subdomain *berkeley.edu*. From the table we see that in either given month, LBL computers engaged in conversations with nearly half of all of the educational and governmental institutions that are accessible via the Internet. This is a strikingly rich degree of connectivity, and in the *edu* domain it is growing.

We next look at LBL's international network traffic. Nearly 4% of all wide area conversations were with hosts in foreign countries, a surprisingly high amount. There are 48 foreign countries connected to the Internet; in March LBL connected to 24 of them (21 in November). Table 7 lists the different countries with which LBL hosts had conversations, along with a count. The countries are grouped into geographical regions: Canada, Europe, Australia, Asia, Central and South America, and the Middle East. While the change in the number of conversations fluctuates a great deal for individual European countries, it shows considerable overall growth in Europe, Australia, and Asia. Some care must be exercised in interpreting the totals, however, particularly the smaller ones. For example, 365 of the 600 conversations with South Korea involved repeated failed attempts of a South Korean host to engage in an *smtp* conversation with LBL's principle mail server. 915 of the conversations with Mexico are *smtp* conversations that occurred between 7:45AM on March 19th and 10:15AM on March 20th.

Finally, we studied the geographical distribution of network traffic within the United States. To do this, we first constructed a list of all of the 534 subdomains that were accessed during the two months. The 293 *edu* subdomains were easy to translate into geographical locations using a reference that gives the location of all colleges and universities in the United States [Web79]. The remaining subdomains were identified using information provided by *whois*. If the *whois* data provided a physical address for the subdomain, that was used. If only a telephone number was listed, we assumed that the subdomain is geographically diffuse. There were 14 such subdomains: most of these were in the *com* domain (such as `ibm.com`) or the *mil* domain (e.g., `army.mil`).

We then produced a mapping between subdomains and states, with the 14 diffuse subdomains (along with the unidentified IP address, `192.31.95.10`) being mapped to "unknown". We furthermore divided the California subdomains into two regions, the San Francisco Bay Area for subdomains north of Monterey (inclusive), and Los Angeles for those to the south. Due to the low resolution of the atlas we had handy [Ran84], we were unable to pinpoint 14 of the 121 California subdomains, so we split their conversations evenly between San Francisco and Los Angeles[18]. Finally, we split conversations to subdomains in Washington D.C. evenly between Maryland and Virginia.

The incomparable *S* data analysis environment[BCW88], which we used heavily for our analysis, provides among many other built-in data sets a map of the United States and enough geographical information regarding states and cities to make good use of it. Figure 13 shows how the network conversations were distributed throughout the world. In addition to the continental United States locations, "AK" and "HI" represent Alaska and Hawaii, "S. Am." South America, "M.E." the Middle East, "P.R." Puerto Rico, and "?" the unknown locations. Also included are Asia, Australia, and Europe, where those regions are comprised of the countries indicated in Table 7.

In the figure the number of conversations with a particular area is proportional to the area of the circle drawn around that area, so a circle twice as wide represents four times as many conversations. November conversation counts are indicated with dotted circles and March counts with solid circles. We see then, for example, that Texas experienced a fair amount of growth in conversations between November and March, while if we shift two states over to the east, Mississippi lost most of its traffic over the same period of time (see the discussion of *domain* traffic in Section 7.9). California's domination of the traffic is no surprise, nor is that of the Bay Area over the Los Angeles region. The growth in the Bay Area is fueled primarily by the 80% increase in the number of conversations between LBL and U.C. Berkeley. This in turn is largely due to the additional of another U.C. Berkeley *nntp* peer. The prevalence of traffic to Utah is no surprise, either, since one of the LBL's other main *nntp* peers is at the University of Utah. 98% of the Utah traffic is due to *nntp*, as was 67% of the November Los Angeles traffic and 80% in March.

Figure 14 shows the same plot with the California data removed. Both in November and in March LBL hosts engaged in conversations with 48 of the 50 U.S. states. Only South Dakota and West Virginia had no conversations. We find this degree of connectivity surprising; we had expected perhaps 30 out of the 50 states to have been contacted.

We look at how different protocols contribute to the geographical distribution of conversations in the next section.

[18] These subdomains participated in just 0.6% of all the California conversations.

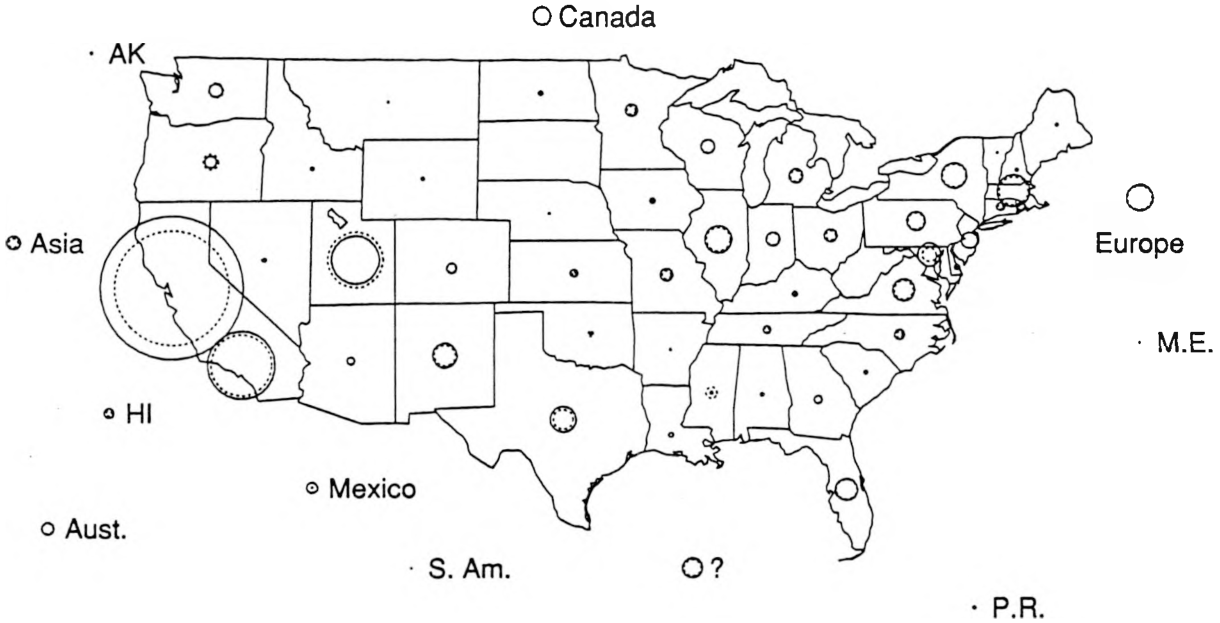

Figure 13: Geographical Distribution of Wide Area Conversations

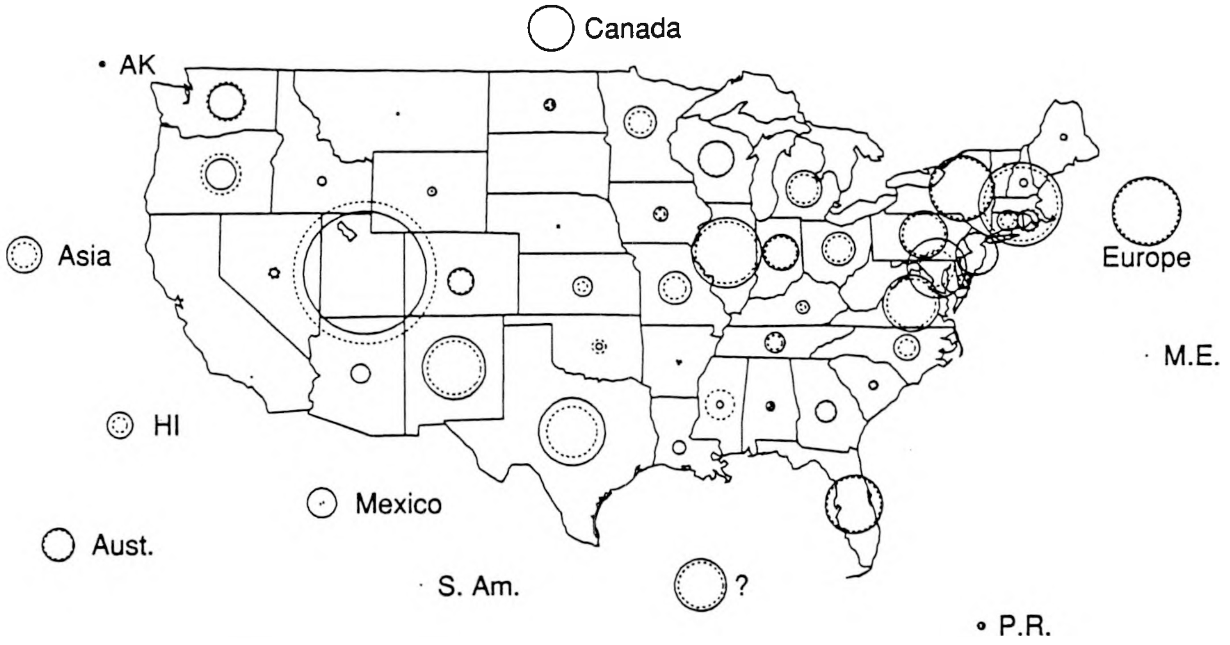

Figure 14: Geographical Distribution for Regions Other Than California

7 Protocol Models

The final topic to which we turn our attention is that of devising models to describe the behavior of different protocols. Our goal was to develop models detailed enough for use in generating realistic simulations of wide area network conversations. We did so for each of the major protocols other than *shell* and *other*, which are too noisy to permit any simple form of model.

The models are summarized in Table 15.

7.1 Constructing Conversation Models

Since data flow in network conversations is asymmetric, it is important to distinguish in our models between the number of bytes sent by the originator and the number sent by the responder. For interactive conversations we were also interested in the conversation durations. Finally, for each protocol we wanted to model the interarrival period between the beginning of one conversation and the beginning of the next.

Unfortunately the traffic data is plagued with spikes, some of which are due to aberrant behavior (e.g., repeated failed attempts to complete a protocol transaction) and some due to the basic nature of the protocol (e.g., *ftp-data* being used to periodically transfer a large data file between a pair of hosts). In our models we identify such spikes and remove them before proceeding. Many protocols also have "failure modes" that may not result in spikes but instead yield many short conversations. We attempted to quantify the proportion of conversations that are failures and remove them, too. Finally, several protocols have some overhead inherent in establishing the conversation. When we were able to identify such overhead we put an approximate value on it (for example, for *smtp* we estimate the originator sends 250 bytes of overhead in addition to the length of the mail message) and added it as a constant offset to the model.

Our models are simple in that for the most part they consist of a single Gaussian fitted to $\log_2$ of the data (after making the above-mentioned adjustments). We found initial estimates of the mean and standard deviation of the Gaussian by making quantile-quantile plots of the data versus the quantiles for a normal distribution [CCKT83], and then using a robust regression to pick out the slope and intercept of the best line fitting the resulting plot. In general we would do this process once for the November data and then repeat it for the March data. If there was much disagreement between the two initial models we sought to identify the simplest change to the models that would bring them into closer agreement.

The key question with any model is: "How good is it?" If our data naturally fell into a fairly small number of fixed regions, we would use a χ^2 test to determine the probability that the data did indeed come from an hypothesized distribution. For continuous data, the analogous test is the Kolmogorov-Smirnov (K-S) test [PFTV86]. This test compares the cumulative distribution of the data *vs.* that of a model as follows: Let N be the number of data points and D be the maximum distance between the cumulative distribution of the data at any ordinate and that of the model at the same ordinate. Define $Q_{KS}(\lambda)$ as follows:

$$Q_{KS}(\lambda) = 2\sum_{j=1}^{\infty}(-1)^{j-1}e^{-2j^2\lambda^2}$$

Then if the data does indeed come from the model's distribution, the probability that we would have observed a value for the data greater than D is:

$$\text{Probability(observation} > \text{D)} = Q_{KS}(\sqrt{N}D)$$

We use the symbol $\Phi = \sqrt{N}D$. From the K-S test we then derive a metric for measuring whether one model is more predictive than another: the model with the least corresponding Φ is the one that is most predictive. We therefore optimized our models by attempting to minimize both Φ_{November} and Φ_{March}.

Ideally we would like our models to have associated Φ's that indicate large probabilities that the model is correct. For example, $\Phi = .5$ gives a correctness probability of 96%; $\Phi = .75$ is 63% probable; $\Phi = 1$ is 27% probable; and then the probabilities deteriorate rapidly, with $\Phi = 1.5$ only 2.2% probable, and $\Phi = 2$ a scant 0.06% probable. It is important to bear in mind that a highly probable value of Φ suggests that the model is "exact"; the data did indeed come from precisely the distribution given by the model. Unfortunately most of our models have $\Phi > 2$, and thus are clearly not exact. We believe, however, that this is due to the large degree to which the data is polluted by spikes and clumps, and that most of our models provide good approximations to the underlying phenomena.

7.2 SMTP

smtp is the Simple Mail Transfer Protocol [Pos82], used for sending electronic mail. The typical conversation involves an originator connecting to a responder; identifying itself; receiving an acknowledgment that the responder is willing to receive mail; identifying the recipient of the mail message; receiving an acknowledgment that the recipient is acceptable; sending the mail message; and receiving an indication that the message was properly received.

Figure 15 shows the distribution of *smtp* conversations. Not surprisingly, they are quite diverse. Of all the regions involved in any sort of network conversation, all but Arkansas, Belgium, Ireland, and Montana participated in an *smtp* conversation. The growth in Bay Area conversations is mostly due to a 28% increase in conversations between LBL and U.C. Berkeley. Los Angeles, Massachusetts, Virginia, New York, and Illinois are the leading other mail destinations. As mentioned above, the conversations with Mexico are aberrant; 95% occurred over a 28 hour period.

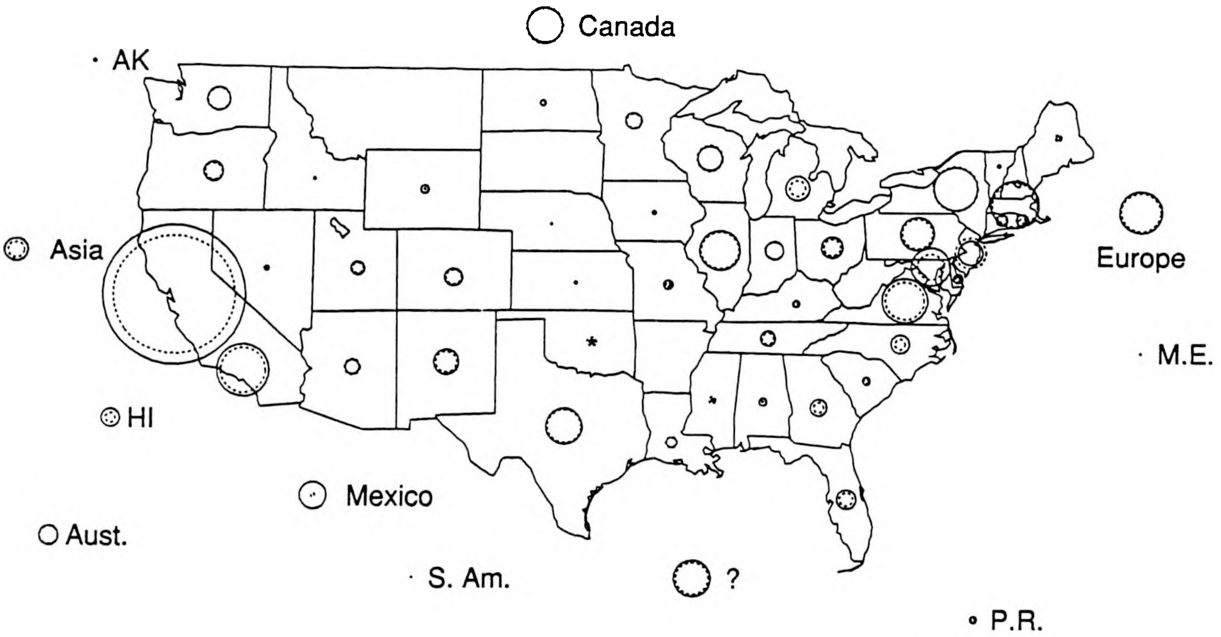

Figure 15: Geographical Distribution of *smtp* Conversations

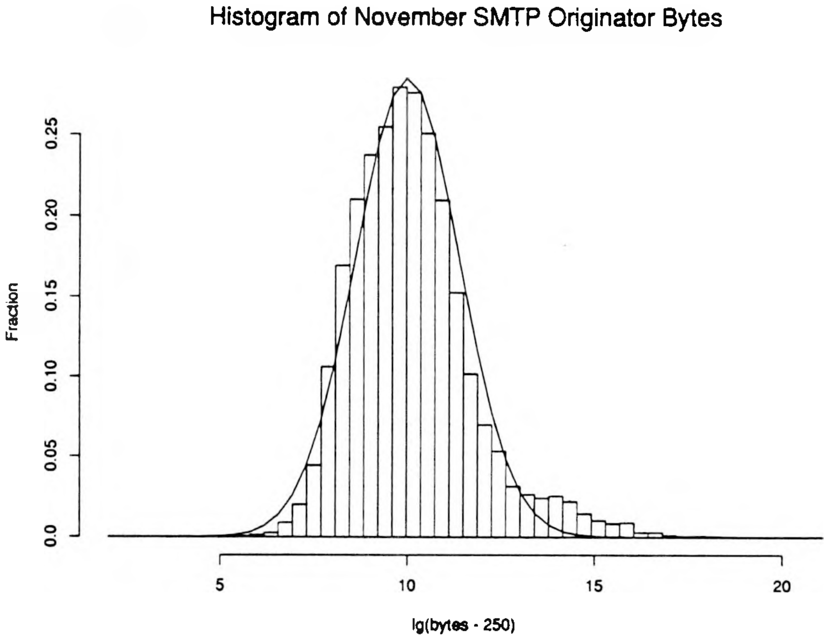

Figure 16: Histogram of log_2 of November *smtp* Originator Bytes

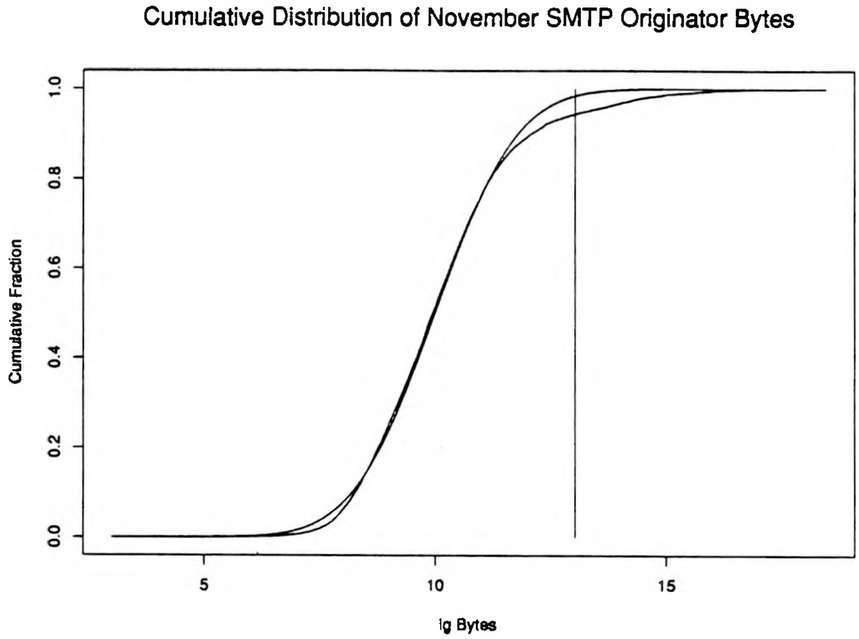

Figure 17: Cumulative Distribution of November *smtp* Originator Bytes

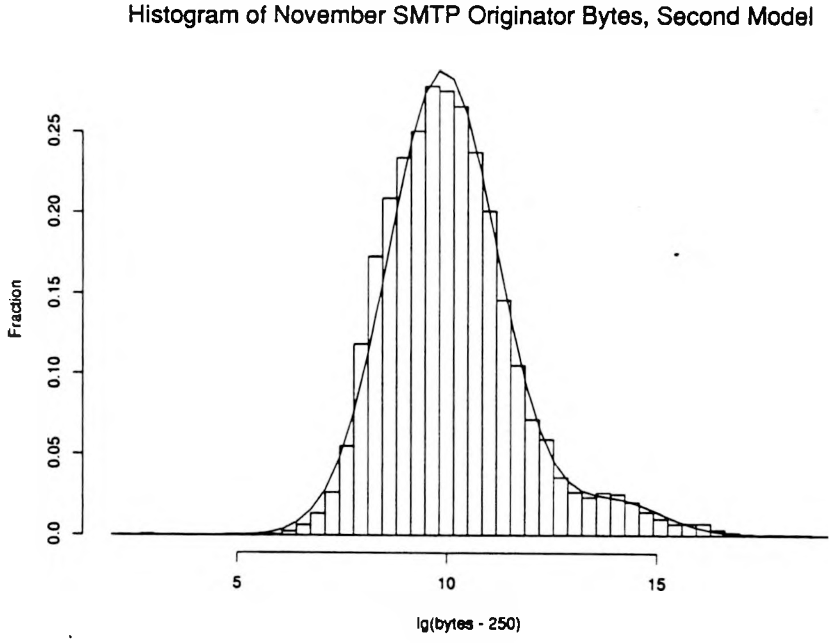

Figure 18: November *smtp* Originator Bytes Fitted to Second Model

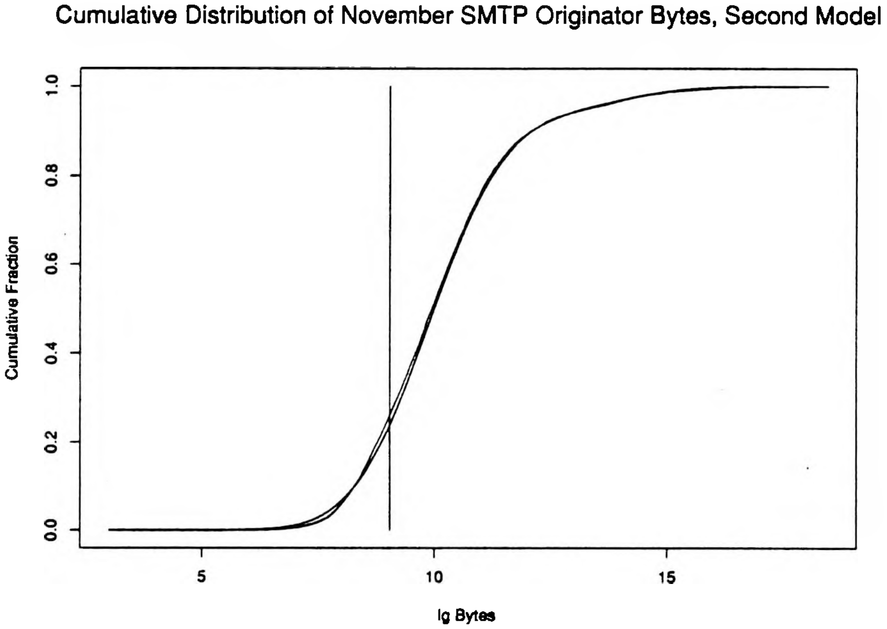

Figure 19: November *smtp* Originator Bytes *vs.* Second Model

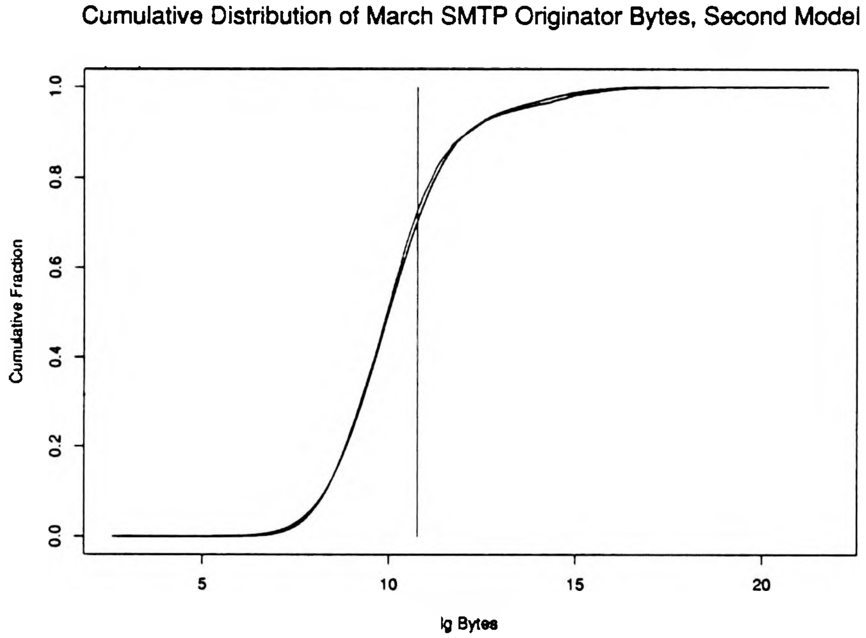

Figure 20: March *smtp* Originator Bytes *vs.* Second Model

Originator

1B failures	$\approx 1\%$
2-250B failures	$\approx 1\%$
250B + Gaussian($\bar{x} = 955$B, $\sigma = \times 2.46$)	$\approx 92.1\%$
250B + Gaussian($\bar{x} \approx 20$KB, $\sigma = \times 2.22$)	$\approx 5.9\%$
$\bar{x}_1$, $\log_2$ of 94% of successes	9.9
$\bar{x}_2$, $\log_2$ of 6% of successes	≈ 14.2
σ_1, $\log_2$ of 94% of successes	1.3
σ_2, $\log_2$ of 6% of successes	1.15
Φ_{November}	2.6
Φ_{March}	3.1

Responder

$<= 100$B	$\approx 1\%$
$>= 500$B	$\approx 1\%$
$\bar{x} = 343$B, $\sigma = 42$	$\approx 98\%$
Φ_{November}	8.0
Φ_{March}	7.2

Interarrival

$\bar{x}$ interarrival, November	19.5 seconds
σ interarrival, November	×4.32 seconds
Φ_{November}	1.01
$\bar{x}$ interarrival, March	15.5 seconds
σ interarrival, March	×4.23 seconds
Φ_{March}	1.16

Table 8: Model for *smtp* Conversations

We first look at modeling the November *smtp* conversations, beginning with the bytes sent by the originator, which includes the body of the email message. 2.2% of all the conversations involved the originator sending 250 or fewer bytes. We deemed these failed delivery attempts. Of these, 38% involved just 1 byte being sent and 13% 7 bytes, with other small spikes present, as well.

Only two conversations of between 200 and 250 bytes occurred, but 14 occurred between 250 and 300 bytes. From these numbers we conjectured that mail messages have roughly 250 bytes of overhead in addition to the message body. We subtracted 250 bytes from the remaining data. Figure 16 shows a histogram of $\log_2$ of the resultant data, along with the fitted Gaussian ($\bar{x} = 10$, $\sigma = 1.4$). Clearly the fit is unacceptably poor in the upper tail. Figure 17 shows the corresponding cumulative distribution plot[19]. The vertical line marks the point of maximal discrepancy between the model and the data, and the corresponding K-S Φ value is 5.35.

The model needs refinement to account for the large discrepancy in the upper regions. We postulated that the distribution might instead be the sum of two separate Gaussians, and with time developed the following model: 94% of the traffic is distributed with $\bar{x} = 9.9$ and $\sigma = 1.3$. The remaining 6% is distributed with $\bar{x} = 14$ and $\sigma = 1.15$. This fit is shown in Figure 18, and the corresponding cumulative plot in Figure 19. The model has $\Phi = 2.57$; not excellent, but certainly much better than a simple Gaussian.

We now turn our attention to the March data, where we find that 1.7% of the conversations were less than 250 bytes long, 52% of those were exactly 1 byte, and another 9% were 7 bytes. After again removing values less than 250 from the data and subtracting an overhead of 250 bytes, Figure 20 shows the March data fitted against the November model. The fit is not quite so good—Φ has risen to 3.1—but still considerably better than a single Gaussian.

If we restrict our K-S measurement to the largest 10% of the March data, though, Φ drops to 1.25 with a probability of about 9%. Interestingly, if we then increase the mean of the second Gaussian from 14 to 14.4 then Φ drops to 0.23 and the fit is perfect (probability 100%). The original mean of 14, though, provides a perfect fit to the top 10% of the November data. We interpret this difference below.

For *smtp* responses, we find that about 1% of the November responses were <= 100B, and about 1% were >= 500B. The rest all fell between 100 and 500 bytes. The data is difficult to fit further since it is reft with spikes, so our model is crude.

The model is summarized in Table 8. The percentages in the second column give the proportion of the traffic with the given characteristic. When converting from a logarithmic model to a linear model, standard deviations become factors instead of additive values. For example, the mean of the bulk of the data lies at $2^{9.9} \approx 955$ bytes, and a data value is within one standard deviation of the mean if it's within a factor of $2^{1.3} \approx 2.46$ of this value, i.e., from 388 to 2,349. The mean of the second (6%) group of data is deliberately somewhat vague since clearly it changed between November and March. Finally, we show two different models for the conversation interarrival time, one for the November data and one for the March data. Both models are very good fits to the data from which they were constructed, but neither model fits the other month's data (Φ_{March} for March fitted to the November model is 4.96; the corresponding Φ_{November} is 4.86).

We interpret the model as follows. The majority of mail messages are composed by the user at the keyboard just prior to sending the mail. They run about 1KB in size, a reasonable amount to type at one sitting. But about 6% of the time email is used to transfer files instead of messages, and these tend to be substantially larger. It is noteworthy that the 94% model fits March nearly as well as it fits November; this implies that the model is somewhat invariant under growth in network utilization, which we would expect since the amount of characters a person is willing to type into an email message should not change much over time. On the other hand, the files being transferred with *smtp* grew in size. We will see below that the same was true of *ftp-data* transfers, suggesting that file sizes increase along with growth in network utilization. Finally, the difference in interarrival times indicates that network usage is increasing by growing during daytime hours and not by spreading out towards greater off-hours network usage, since the mean "density" of conversations is climbing. We will see this trend repeated in almost every other model of interarrival times.

7.3 FTP and FTP-DATA

ftp is the File Transfer Protocol [Bhu72], used for sending files between hosts. *ftp* is the interactive half of the file transfer: the user sends their username and password to the remote host for verification, and then issues commands to list the files available on the remote host and to `get` a file from the remote host or `put` a file to it. The actual file transfer is done using a separate *ftp-data* connection that is initiated by the remote host. Commands exist to transfer multiple files at one time, too. Each file in such a group is sent using a separate *ftp-data* connection.

Figure 21 shows the distribution of *ftp* conversations. The connectivity is as rich as *smtp*'s: of all the regions involved in any sort of network conversation, all but Argentina, Israel, Puerto Rico, and Wyoming participated in an *ftp* conversation.

A fairly good model for *ftp* conversations is summarized in Table 9. The final section of the table indicates how many *ftp-data* companion conversations are expected for an *ftp* conversation. The last entry states that if the *ftp* conversation spawns multiple *ftp-data* conversations then on average it will spawn 4.9 of them. This value was computed from the mean of the log of the number of *ftp-data* conversations that occurred dur-

[19] Due to graphical limitations, the plot was made using every fifth data value instead of all the data values.

Originator

<= 30B failures	$\approx$ 1%
30B + Gaussian($\bar{x}$ = 119B, σ = ×2.83)	$\approx$ 99%
Φ_{November}	3.3
Φ_{March}	3.8

Responder

$\bar{x}$ = 484B, σ = ×2.27	100%
Φ_{November}	3.5
Φ_{March}	4.7

Corresponding FTP-DATA conversations

0 corresponding conversations	$\approx$ 20%
1 corresponding conversation	$\approx$ 29%
> 1 corresponding conversation	$\approx$ 51%
if > 1, avg. how many	4.9

Interarrival

$\bar{x}$ interarrival, November	172 seconds
σ interarrival, November	×4.43 seconds
Φ_{November}	0.84
$\bar{x}$ interarrival, March	137 seconds
σ interarrival, March	×4.20 seconds
Φ_{March}	0.76

Table 9: Model for *ftp* Conversations

originator is sender	$\approx$ 80%
originator is receiver	$\approx$ 20%
sender: Gaussian($\bar{x} \approx$ 2,400B, $\sigma \approx$ ×20)	100%
receiver: 1B	100%
Φ_{November}	3.09
Φ_{March}	3.44
$\bar{x}$ interarrival, November	8.9 seconds
σ interarrival, November	×6.15 seconds
Φ_{November}	1.65
$\bar{x}$ interarrival, March	8.5 seconds
σ interarrival, March	×7.16 seconds
Φ_{March}	1.93

Table 10: Model for *ftp-data* Conversations

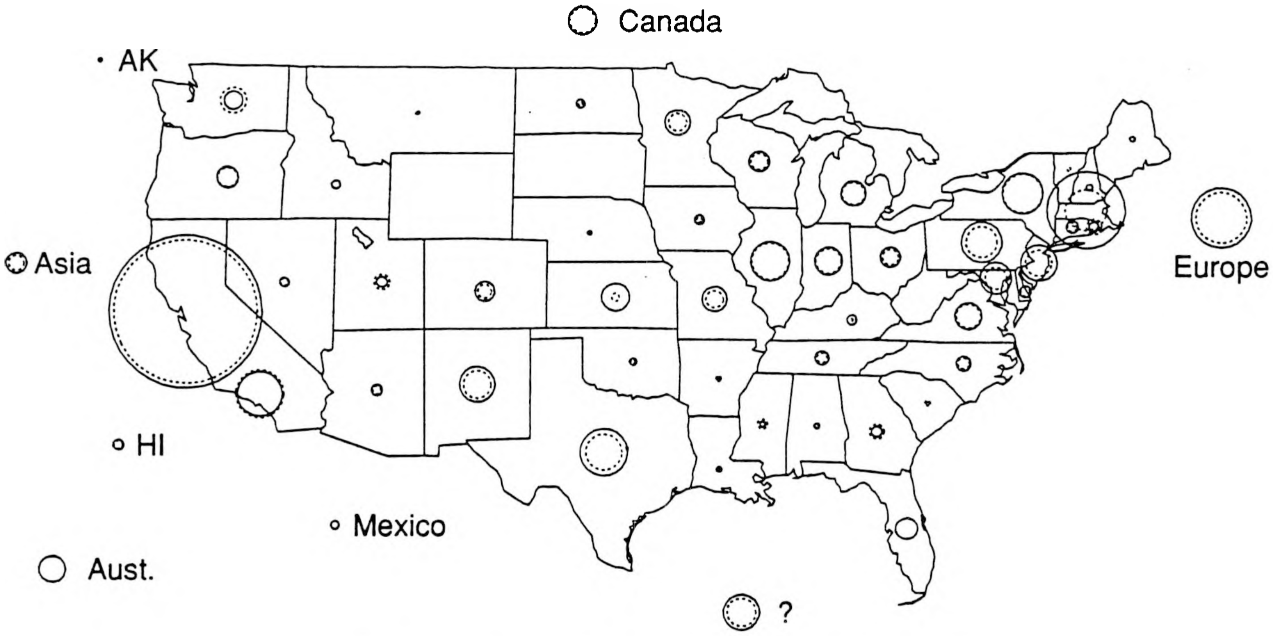

Figure 21: Geographical Distribution of *ftp* Conversations

ing each *ftp* conversation in November and March. The mean of the counts themselves (not their logarithms) is ≈ 9, due to some amazing outliers (one *ftp* conversation spawned 1,006 *ftp-data* conversations!).

ftp-data conversations are virtually mono-directional. Either the originator (for a `get` command) or the responder (`put` command) sends only one byte[20] while the other sends the file being transferred. In November 24% of the conversations were `get`'s; in March, 16%. The data is riddled with spikes, including 61 of 8,406,027 bytes and 5 of 27,341,210 bytes. These are presumably the same data set being repeatedly shipped back and forth between hosts.[21]

Table 10 presents a fairly good model for *ftp-data*. The model's main weakness lies in underestimating the frequency of large transfers, probably due to the effects of spikes. A better fit can be made to the November data by lowering the $\tilde{x}$ to 2,195B and σ to 3.95; Φ_{November} then becomes 2.67; and if we restrict ourselves to the first 80% of the data, $\Phi_{\text{November}} =$ 1.06, with likelihood 21%. From this we conclude that the average size of the files being transferred grew significantly between November and March, by roughly 10%.

7.4 TELNET

telnet is an Internet protocol for remote login [PR83]. A conversation is begun by the originator and responder engaging in a logging-in sequence (verification of username and password), along with possibly the negotiation of options related to terminal characteristics. Provided that this initial exchange is successful, the originator then sends whatever keystrokes are typed by the user, and the responder replies with the corresponding output generated by the remote host.

Figure 22 shows the geographical distribution of *telnet* conversations. The coverage is surprisingly broad for an interactive protocol where the low bandwidths one might expect with longer distance connections would tend to make interactive use painful.

Table 11 presents a good model for *telnet* conversations. Failures are probably either access denied (password incorrect or no user account on remote machine) or logins initiated to execute just one command. The large responder overhead is no doubt due to the cost of sending the login prompt, system identification banner, and the output of any commands the user executes automatically upon logging in. The slightly high Φ_{November} can be lowered to 1.40 by dropping $\tilde{x}$ from 6654 to 6000. Thus the average amount of traffic generated by a *telnet* responder grew by about 10% between November and March, without a corresponding growth in *telnet* originator traffic. We do not know how to interpret this change, and are unsure of its significance.

7.5 LOGIN

login is a remote login protocol for Unix machines. It is quite similar to *telnet* regarding the roles of the originator and the responder, as well as to the initial overhead of establishing a connection.

Figure 23 shows the geographical distribution of *login* conversations. The coverage is considerably less than that for *tel-*

[20] We observed 2 conversations in which no bytes were sent.

[21] The spike of 61 8.4MB conversations occurred in March, but a spike of 6 such conversations of identical size was also present in November.

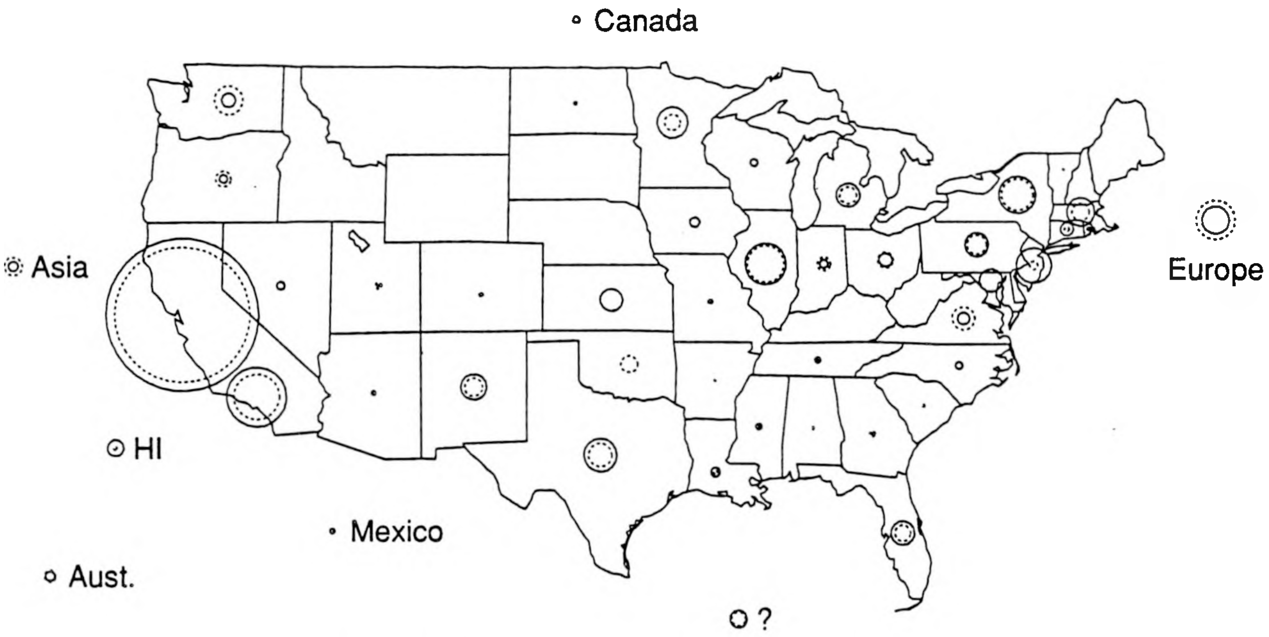

Figure 22: Geographical Distribution of *telnet* Conversations

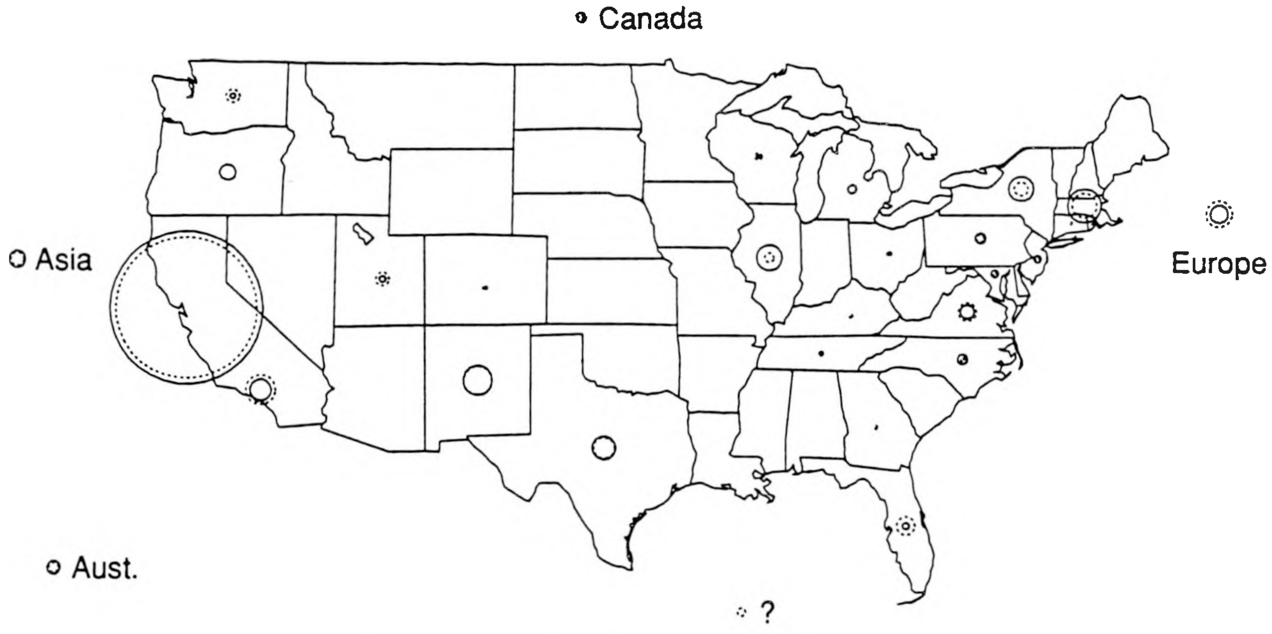

Figure 23: Geographical Distribution of *login* Conversations

Originator

<= 55B failures	$\approx$ 16.5%
55B + Gaussian($\bar{x}$ = 194B, σ = $\times$7.0)	$\approx$ 83.5%
Φ_{November}	1.67
Φ_{March}	1.85

Responder

<= 600B failures	$\approx$ 16.1%
600B + Gaussian($\bar{x}$ = 6654B, σ = $\times$6.3)	$\approx$ 83.9%
Φ_{November}	2.10
Φ_{March}	1.56

Duration

<= 30 sec	$\approx$ 15%
30 sec + Gaussian($\bar{x}$ = 324 sec, σ = $\times$7.7)	$\approx$ 85%
Φ_{November}	1.35
Φ_{March}	1.17

Interarrival

$\bar{x}$ interarrival, November	94 seconds
σ interarrival, November	$\times$4.25 seconds
Φ_{November}	2.24
$\bar{x}$ interarrival, March	69 seconds
σ interarrival, March	$\times$4.17 seconds
Φ_{March}	2.88

Table 11: Model for *telnet* Conversations

Originator

<= 60B failures	$\approx$ 18%
60B + Gaussian($\bar{x}$ = 162B, σ = ×8.0)	$\approx$ 82%
Φ_{November}	1.81
Φ_{March}	2.02

Responder

<= 600B failures	$\approx$ 12.4%
600B + Gaussian($\bar{x}$ = 5997B, σ = ×7.5)	$\approx$ 87.6%
Φ_{November}	2.42
Φ_{March}	2.60

Duration

<= 30 sec	$\approx$ 14.5%
30 sec + Gaussian($\bar{x}$ = 401 sec, σ = ×8.8)	$\approx$ 85.5%
Φ_{November}	1.66
Φ_{March}	1.59

Interarrival

$\bar{x}$ interarrival, November	172 seconds
σ interarrival, November	×4.22 seconds
Φ_{November}	1.90
$\bar{x}$ interarrival, March	164 seconds
σ interarrival, March	×4.11 seconds
Φ_{March}	1.17

Table 12: Model for *login* Conversations

net conversations. During both November and March there were only half as many *login* conversations as *telnet*, which is probably because *login* in general is between Unix hosts and *telnet* need not be. We conjecture that the lower geographical coverage of *login* is simply due to there being fewer total conversations, and so areas visited rarely by *telnet* conversations might be missed entirely by *login*.

Table 12 presents a fairly good model for *login* conversations. Forming the model was somewhat problematic, though, because the November and March datasets had qualitative differences. The "overhead" cutoff of 600B removed 14.3% of the November data but only 10.5% of the March data. Furthermore, increasing $\bar{x}$ to 178 for the March originator data lowers Φ_{March} to 0.79, which is 56% probable—i.e., the model is excellent. Decreasing $\bar{x}$ to 147 for the November originator data lowers Φ_{November} to 1.07, 20% probable (and if the smallest one-fifth of the data is excluded, Φ_{November} decreases to 0.65, 79% probable).

Increasing $\bar{x}$ to 6650 and σ to ×7.7 for the March responder data lowers Φ_{March} to 1.27, 8% probable; lowering $\bar{x}$ to 5400 and σ to ×7.2 lowers Φ_{November} to 1.06, 21% probable[22].

Thus excellent models exist for both the November and March data, but they are different. March shows an overall 20% increase in bytes sent, both by originator and responder.

Comparing the *login* model with that for *telnet* we see that *login* conversations tend to consist of about 15% fewer user keystrokes and 10% fewer response bytes, and to run about 24% longer. We conjecture that the former difference may be due to the terseness of the Unix operating system in both user input required and feedback provided by utilities. The latter difference may be rooted in the fact that more than two thirds of all *login* conversations were with U.C. Berkeley while only just over one third of the *telnet* conversations were. The fast network response between the two sites may encourage users to leave their conversations open longer than they would otherwise.

7.6 FINGER

finger is the Internet protocol for accessing remote user information [Zim90]. A *finger* conversation consists of the originator sending the name of the remote user for whom the information is requested, and the responder sending back a description of the remote user. The conversation then terminates.

Figure 24 shows the geographical distribution of *finger* conversations. Clearly there are a lot of curious people on the Internet. The large number of conversations with Massachusetts, almost the same number both months, will be discussed shortly. The large growth in traffic to Washington state has been discussed previously; it is due to the appearance of a *finger*-based weather-reporting service that's proven quite popular.

Messages sent by *finger* originators incur 3 bytes of overhead in addition to the name being looked up, if any. 13% of the November originators and 23% of the March ones sent just 3 bytes. This is not a "failure" but rather a request to see a list of all users presently logged into the remote system. Other than that, the originator tends to send between 6 and 12 bytes total (84% of the November conversations, 73% of March), with the majority lying between 9 and 11 bytes.

Both the November and the March data are polluted by a large number of repeated *finger* conversations, possibly for cracking purposes. In November, three MIT hosts[23] made 1,380 *finger* conversations with the same LBL host. 98% of the conversations followed a previous *finger* conversation by two minutes or less. 11 bytes were always sent, and 81 bytes received (this is the number of bytes in the LBL host's "no such user" response). Sure enough, in the March data four M.I.T. hosts (two in common with November) made 905 *finger* conversations with the same LBL host. Again 98% followed a previous conversation by under two minutes.

Other spikes exist in the *finger* data, usually associated with two particular hosts as well. It also became clear that some LBL users have scripts they run to "finger" several different U.C. Berkeley hosts in quick succession (no doubt looking for a friend or colleague), which they run fairly frequently. Not surprisingly, the responder data is also highly skewed by spikes, and we do not attempt to model it except to compute a mean and standard deviation of $\log_2$ of the byte count: $\bar{x} \approx$ 256B, $\sigma \approx$ ×6.

Finally, a crude model for interarrival is $\bar{x} \approx$ 45 seconds, $\sigma \approx$ ×4.28.

7.7 NNTP

nntp is the Network News Transfer Protocol [KL86], used for propagating network news messages across the Internet (and reading news locally as well). Typically, once an originator connects to a responder it offers the responder a series of what it believes are new news articles. The responder replies to each in turn indicating whether it wants the article, and if so, the article is transferred. The responder might also initially indicate that it is unable to receive news at the present time.

nntp servers have a set of peers with whom they transact the great majority of their conversations. When a server receives a new news article from one of its peers, it immediately offers it to all its other peers.

Figure 25 shows the geographical distribution of *nntp* conversations. LBL's *nntp* peers are at U.C. Berkeley, Lawrence Livermore National Laboratory (which is in the San Francisco area), San Diego, and Utah. The conversations with Massachusetts all involved MSRI's *nntp* server. The lone

[22] One remaining puzzle, though, is why the percentage of responder failures is different from that of originators. This discrepancy suggests that the failure modes are more complex than the model accounts for.

[23] e40-008-9.mit.edu, e40-008-8.mit.edu, and hawaii.mit.edu.

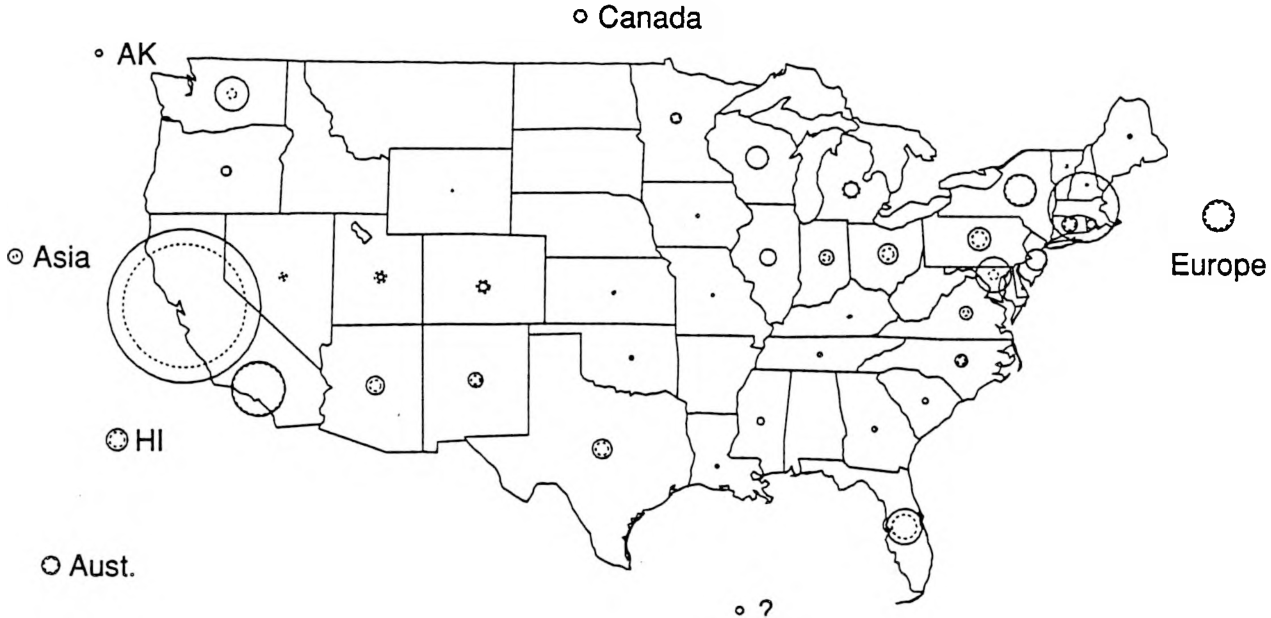

Figure 24: Geographical Distribution of *finger* Conversations

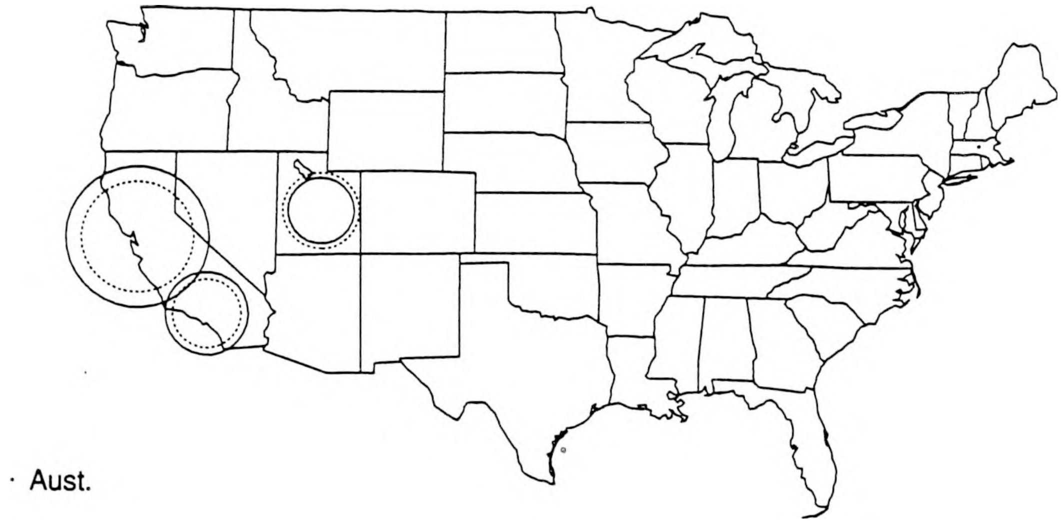

Figure 25: Geographical Distribution of *nntp* Conversations

November

7B failures	$\approx 38\%$
Gaussian($\bar{x} = 68\text{B}, \sigma = \times 1.8$)	$\approx 15.5\%$
Gaussian($\bar{x} = 6080\text{B}, \sigma = \times 5.3$)	$\approx 46.5\%$
Φ_{November}	2.7
Φ_{March}	14

March

7B failures	$\approx 25\%$
Gaussian($\bar{x} = 68\text{B}, \sigma = \times 1.8$)	$\approx 9.7\%$
Gaussian($\bar{x} = 6080\text{B}, \sigma = \times 5.3$)	$\approx 65.3\%$
Φ_{November}	8.0
Φ_{March}	2.8

Table 13: Model for *nntp* Conversations

conversation between Australia and LBL is unexplained; it originated in Australia.

The *nntp* data is complicated and noisy. The responder data is so riddled with spikes (for both months about 50% of the data appears in one of twenty or so spikes) that we simply model the responses as having $\bar{x} \approx 237$ and $\sigma \approx \times 2.16$. This model is quite bad: it gives $\Phi_{\text{November}} = 22$ and $\Phi_{\text{March}} = 18$.

The originator data is a little easier to model, though the differences between November and March are significant. The model is given in Table 13. While $\bar{x}$ and σ are the same for both months, the associated percentages have changed dramatically. This is due to the addition between November and March of another U.C. Berkeley *nntp* peer (and a particularly well-connected one at that), and dramatizes the volatility of *nntp* conversation patterns.

nntp interarrival times are likewise noisy. We simply model them as correlated conversations (due to propagation via flooding), with $\bar{x} \approx 1/10$th of a second and $\sigma \approx \times 4$; and uncorrelated conversations (arrival of new news), with $\bar{x} \approx 25$ seconds and $\sigma \approx \times 3$.

7.8 X11

X11 is the network protocol used by the X11 Window System [SG86]. In an *X11* conversation, the originator is the X11 *client*, which sends X11 graphics directives to the remote X11 server. The responder is the X11 *server*, the entity that actually paints pixels on the user's screen. The server primarily sends back input events (such as key presses and mouse motion) and status values for previous graphics directives.

Figure 26 shows the geographical distribution of *X11* conversations. The large amount of traffic with Texas is due to graphical interface work being done at LBL for a project at the Superconducting Super Collider Laboratory in Dallas. Interestingly, this is one of the only regions where *X11* traffic actually grew between November and March. The Asia and Europe connections are somewhat surprising—one wouldn't expect enough bandwidth to make such conversations tolerable. The connections with Europe were all to the same host in Switzerland, averaging 600 B/sec. This is actually higher than the average X11 bandwidth of about 400 B/sec, so presumably response was acceptable.

The *X11* data is full of spikes and clumps, particularly the responder. Table 14 shows a crude model for these conversations. The only rather good part of the model is the fit for the duration.

7.9 DOMAIN

The final protocol we modeled was *domain*, used for exchanging hostname information among the Internet's distributed name servers. Name servers maintain caches of hostnames and their corresponding IP addresses. Entries in the cache have a "time to live" after which the entry is no longer valid and must be refreshed from one of the name server's peers. Name servers also periodically poll their peers to verify whether their cache is still valid. The majority of the time it is, but when the peer informs the name server that the local cache is invalid this is followed by the peer sending to the name server a new table of hostname data. Thus a typical *domain* conversation entails the originator sending a short request to the responder, and then usually receiving a short reassurance that all is well, but occasionally receiving a potentially quite large table.

Figure 27 shows the geographical distribution of *domain* conversations. LBL's main name server peers are in New

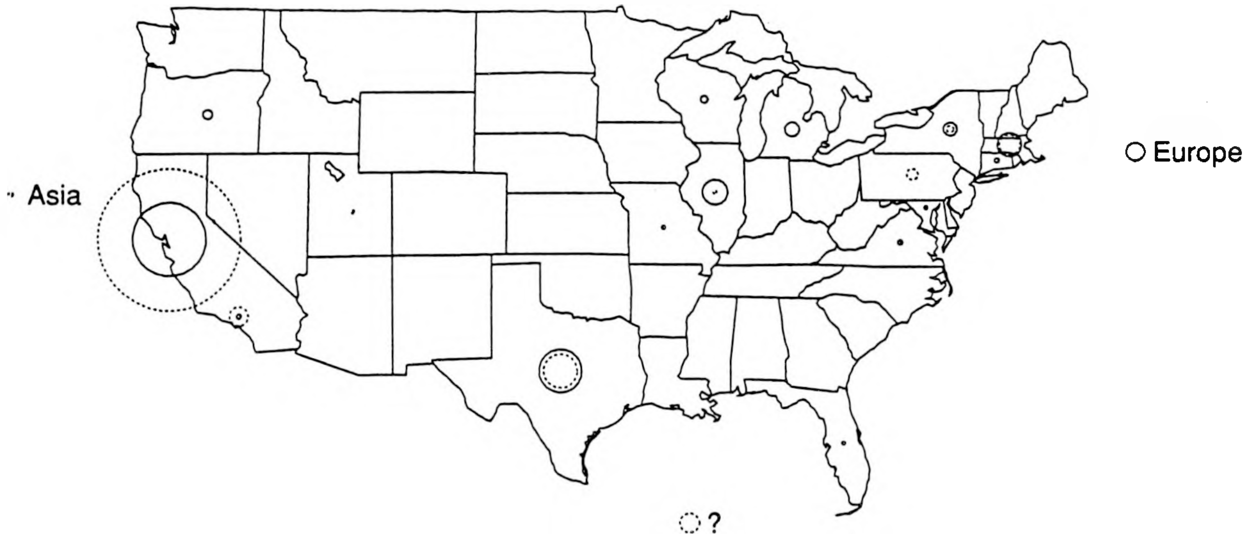

Figure 26: Geographical Distribution of *X11* Conversations

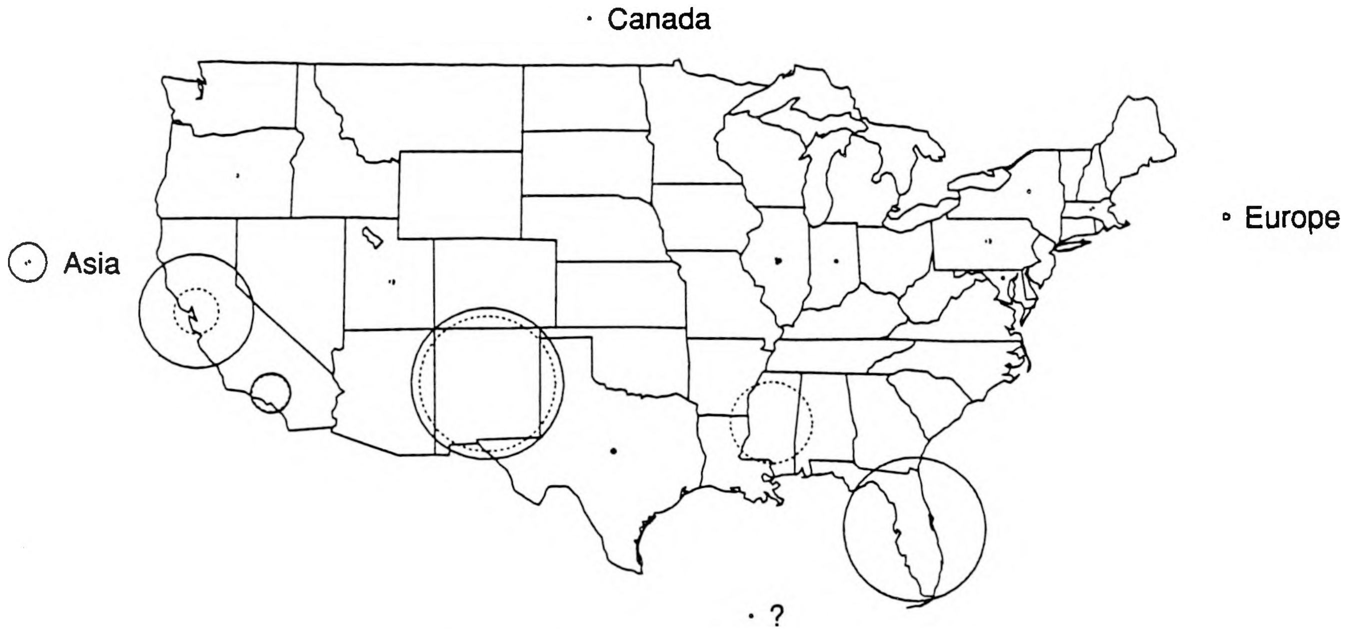

Figure 27: Geographical Distribution of *domain* Conversations

Originator

Gaussian($\bar{x}$ = 11.5KB, σ = ×14)	100%
Φ_{November}	4.2
Φ_{March}	2.6

Responder

Gaussian($\bar{x}$ = 5.8KB, σ = ×7.5)	100%
Φ_{November}	8.4
Φ_{March}	6.6

Duration

<= 10 sec	≈ 28%
10 sec + Gaussian($\bar{x}$ = 97 sec, σ = ×9.2)	≈ 72%
Φ_{November}	2.28
Φ_{March}	1.74

Interarrival

$\bar{x}$ interarrival, November	111 seconds
σ interarrival, November	×10.5 seconds
Φ_{November}	2.73
$\bar{x}$ interarrival, March	199 seconds
σ interarrival, March	×11.8 seconds
Φ_{March}	1.15

Table 14: Model for *X11* Conversations

Mexico, Florida, and Los Angeles. The large number of conversations with Asia (Japan, in particular) show an 11% success ratio, and hence indicate legitimate name server interaction. The even larger number of conversations with Mississippi (which only occurred during November) were all failures, and indicate an erroneous configuration on the Mississippi end (all conversations originated in Mississippi).

domain originators send very short messages: all were between 20 and 87 bytes, with 60-70% being spikes of 28, 41, or 44 bytes.

domain responses are very clumped, too, with the vast majority (85% in November and 96% in March) lying between 70 and 120 bytes, but with numerous clumps (different between November and March) in the 100-300KB range, as well as smaller clumps between 120B and 85KB. We do not attempt to further model these clumps, other than to note that about 3% of the responses were in the 100-300KB range, and only one response was more than 309KB in size (it was 589KB). Finally, the *domain* interarrival times fall primarily into two regions: $\approx$ 28% between 1/16th of a second and 1 second, and $\approx$ 40% between 256 seconds and 1,024 seconds. The larger region corresponds with a roughly ten minute polling period, perhaps the phase difference between several larger polling periods. The smaller region is due to conversations between the same LBL name server and different remote servers becoming synchronized due to name server restarts.

7.10 Summary

We modeled nine different protocols, four interactive and five non-interactive. Our modeling efforts were repeatedly hampered by variability in the data in the form of spikes and clumps, but a general picture emerged of a constant proportion of each protocol's conversations being *failures* and the rest following a pattern of a fixed-size overhead plus a logarithmic Gaussian or possibly the sum of two such Gaussians. Applying the Kolmogorov-Smirnov test to our models revealed that nearly all of them are incomplete, though some are very predictive for a significant fraction of the conversations. Salient features of six of these models are summarized in Table 15. For each type of conversation, "Failures" refers to the prevalence of very short conversations, and "Modality" to the number of modes excluding failures.

Protocols	*telnet, login*
Modality	Single
Failures	15%
Avg. duration of remainder	5-6 minutes
Avg. bytes sent	175
Avg. ratio of bytes sent to received	1:35
Model fit for bytes transferred	Good
Model fit for duration	Very good

Protocol	*smtp*
Modality	Bimodal
Failures	2%
Avg. email message size	1KB; 92%
Avg. email file transfer size	20KB; 6%
Model fit	Good

Protocols	*ftp, ftp-data*
Modality	Single
Avg. file size	2.4KB
Change in avg. file size in 4 months	+10%
Avg. number transfers if > 1	4.9
Likelihood no file is transferred	20%
Model fit	Fair

Protocol	*nntp*
Modality	Bimodal
Failures	$\approx$ 30%
Little or no news transferred	$\approx$ 12%
News transferred	$\approx$ 58%
Avg. size of transfer	6KB
Compromise model fit to either month	Poor
November model fit to November data	Good
March model fit to March data	Good

Table 15: Summary of Models

8 Acknowledgments

This work was supported in part by an NSF Graduate Fellowship, for which the author is grateful.

Van Jacobson made this project possible by providing the necessary tools, inspiration, and first-rate strategic and tactical advice. His guidance has been a great help. He also served as an invaluable resource when attempting to interpret network-related phenomena.

Steve McCanne's help with using *tcpdump* and BPF proved invaluable. The speed with which he incorporated new features and bug fixes into these tools was a delight, and greatly aided the rapid development of the data reduction software. He was also instrumental in helping me master the necessary PostScript magic for correctly inserting *S* plots into LaTeX figures.

Craig Leres provided much-needed help, ranging the spectrum from answering networking questions to installing hardware to finding additional disk space. He served as chief *nntp* and *domain* analyst.

John Baker and Peter Fuller of LBL were very helpful in analyzing patterns of behavior among their particular LBL hosts.

I wish to thank my research advisor, Prof. Domenico Ferrari, for his excellent guidance at key junctures during the

course of this project, and Ed Theil, the head of Computer Systems Engineering at LBL, for arranging for the necessary support at LBL. I also wish to thank the authors of the superlative *S* tool, without which the data analysis would have been much, much more painful, and Lindsay Schachinger and Johan Bengtsson for especially helpful insights into the data modeling process.

Finally, I especially wish to thank my wife, Lindsay, for her continual patience and support during the pursuit of this project.

All of these efforts are very much appreciated.

References

[BCW88] Richard A. Becker, John M. Chambers, and Allan R. Wilks. *The New S Language*. Wadsworth & Brooks, Pacific Grove, CA, 1988.

[Bhu72] A. K. Bhushan. File transfer protocol (ftp) status and further comments. RFC 414, Network Information Center, SRI International, Menlo Park, CA, 1972.

[CCKT83] John M. Chambers, William S. Cleveland, Beat Kleiner, and Paul A. Tukey. *Graphical Methods for Data Analysis*. Wadsworth International Group, 1983.

[Fro91] Cliff Frost. Personal Communication, 1991.

[HSF85] K. Harrenstien, M. K. Stahl, and E. J. Feinler. NICNAME/WHOIS. RFC 954, Network Information Center, SRI International, Menlo Park, CA, 1985.

[JLM89] Van Jacobson, Craig Leres, and Steve McCanne. *TCPDUMP(1)*. Available via anonymous ftp to `ftp.ee.lbl.gov`, June 1989.

[KL86] B. Kantor and P. Lapsley. Network news transfer protocol. RFC 977, Network Information Center, SRI International, Menlo Park, CA, 1986.

[MJ91] Steven McCanne and Van Jacobson. An efficient, extensible, and portable network monitor. Forthcoming, 1991.

[PFTV86] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes*. Cambridge University Press, 1986.

[Pos81] J. B. Postel. Transmission control protocol. RFC 793, Network Information Center, SRI International, Menlo Park, CA, 1981.

[Pos82] J. B. Postel. Simple mail transfer protocol. RFC 821, Network Information Center, SRI International, Menlo Park, CA, 1982.

[PR83] J. B. Postel and J. K. Reynolds. Telnet protocol specification. RFC 854, Network Information Center, SRI International, Menlo Park, CA, 1983.

[Ran84] *The New Rand McNally College World Atlas*. Rand McNally & Company, 1984.

[SG86] R. Scheifler and J. Gettys. The X window system. *ACM Transactions on Graphics*, (63), 1986.

[Web79] *Webster's New Collegiate Dictionary*. G. & C. Merriam Company, Springfield, MA, 1979.

[Zim90] D. P. Zimmerman. Finger user information protocol. RFC 1194, Network Information Center, SRI International, Menlo Park, CA, 1990.

9 Appendix

Figure 28 through 34 show the active conversations for all protocols for the weeks of November 4-10, 18-24, 25-30, and March 3-9, 10-16, 17-23, and 24-29, respectively. See Figure 6 for the active conversations for the week of November 11-17.

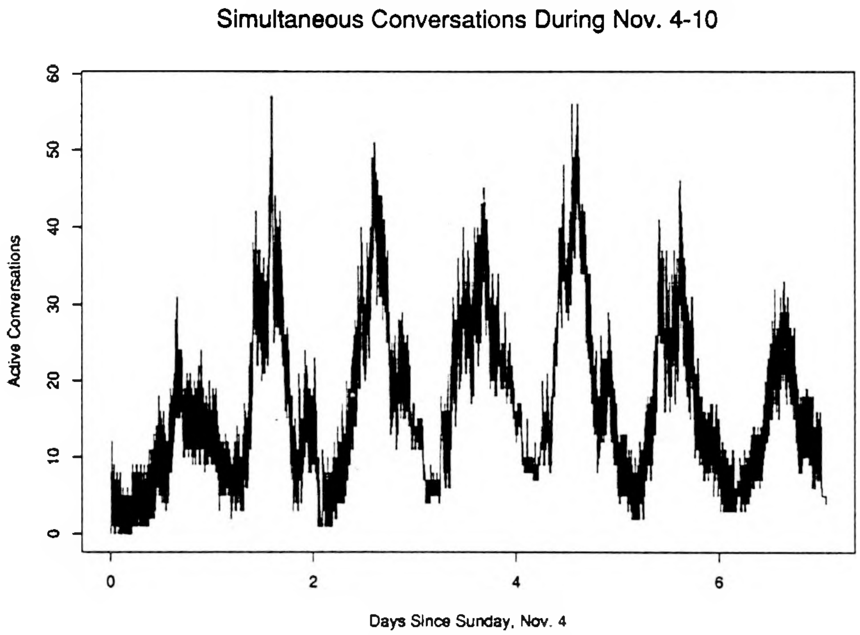

Figure 28: Active Conversations During the Week of November 4-10

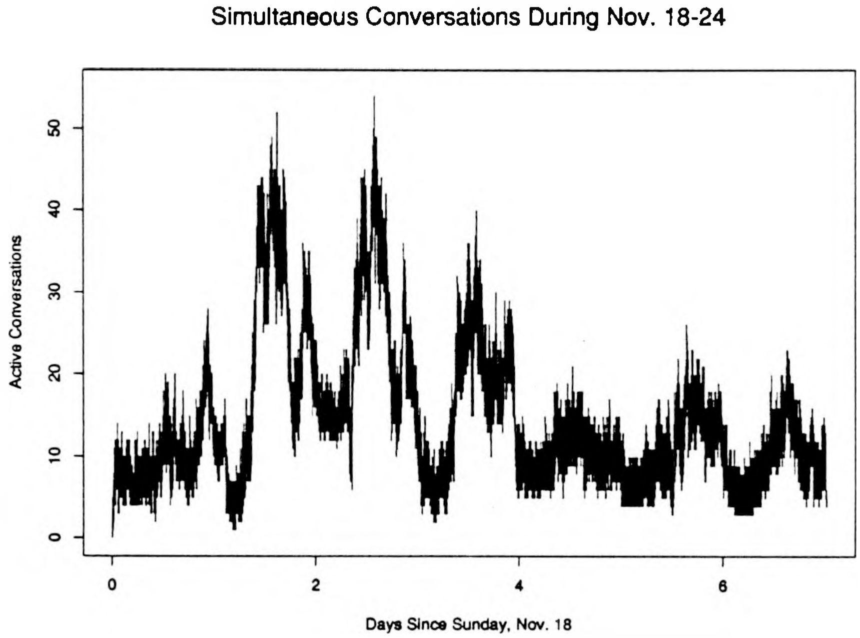

Figure 29: Active Conversations During the Week of November 18-24

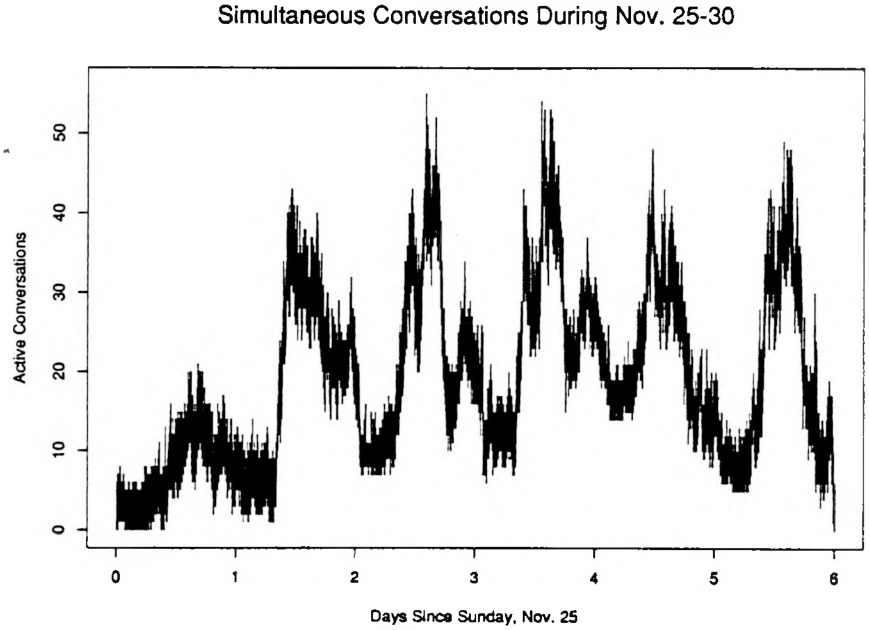

Figure 30: Active Conversations During the Week of November 25-30

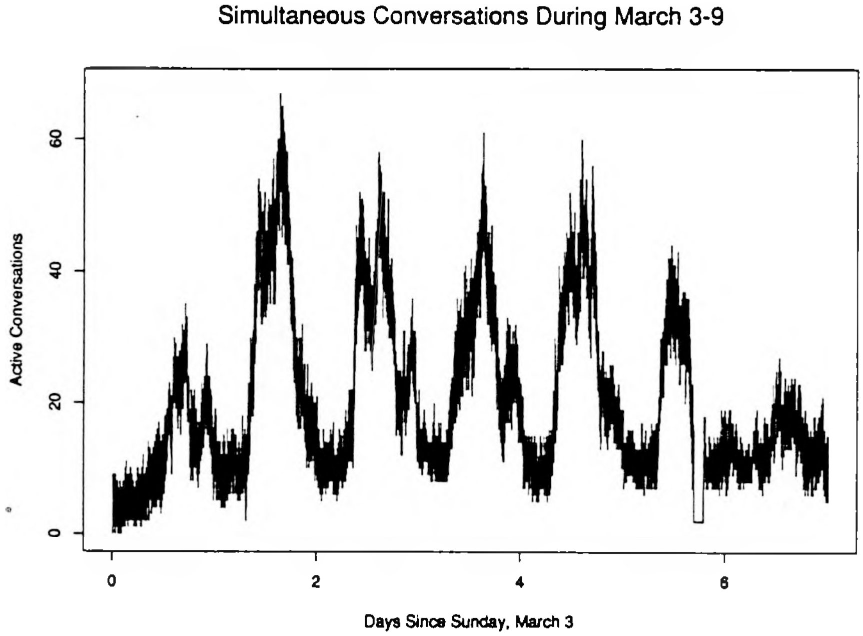

Figure 31: Active Conversations During the Week of March 3-9

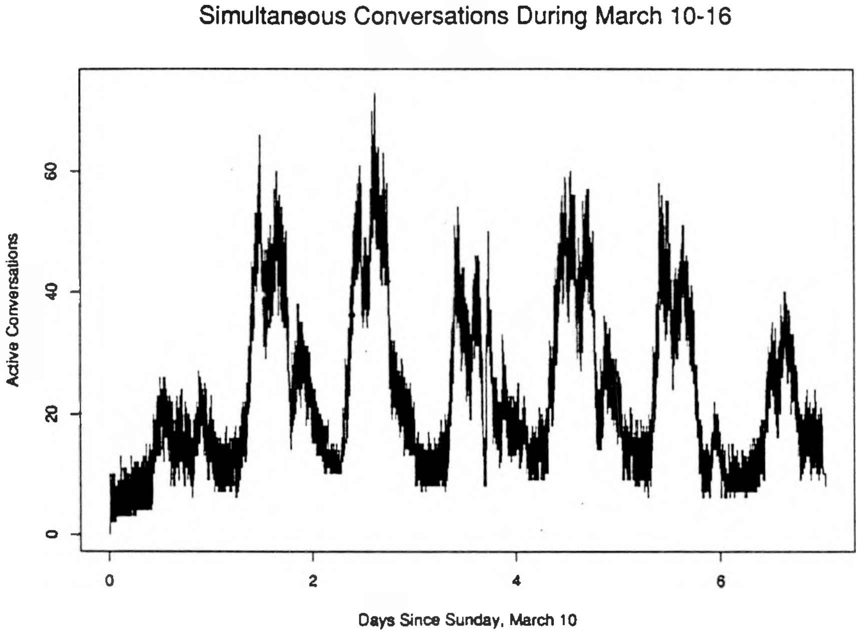

Figure 32: Active Conversations During the Week of March 10-16

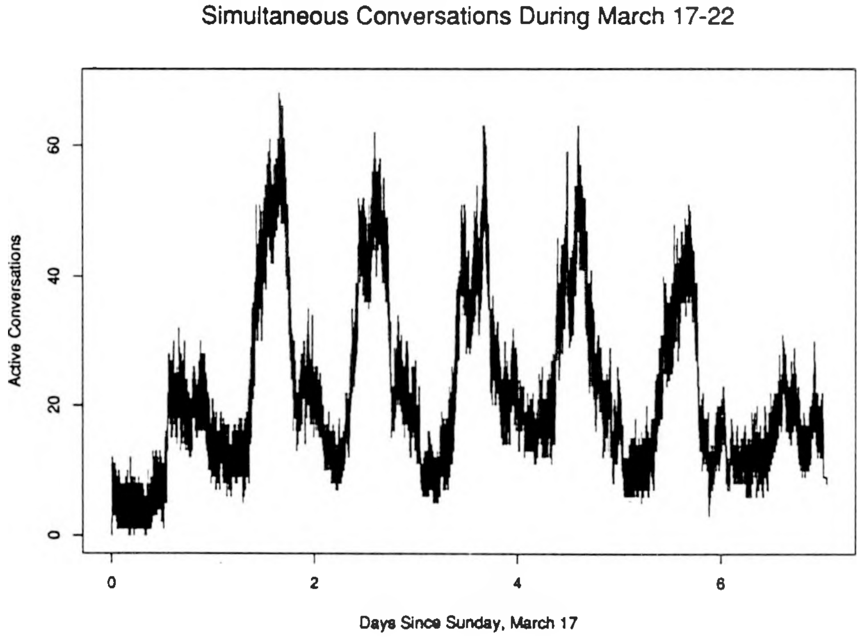

Figure 33: Active Conversations During the Week of March 17-23

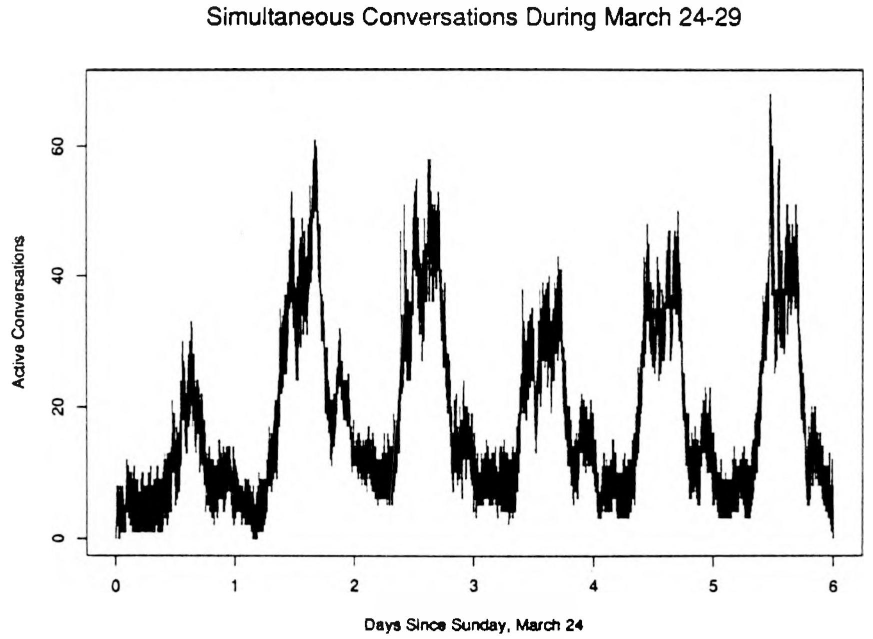

Figure 34: Active Conversations During the Week of March 24-29