

Measurements of Wide Area Internet Traffic

Ramón Cáceres

Computer Science Division
University of California
Berkeley CA 94720
ramon@berkeley.edu

ABSTRACT

Measurement and analysis of current behavior are valuable techniques for the study of computer networks. In addition to providing insight into the operation and usage patterns of present networks, the results can be used to create realistic models of their traffic sources. Such models are a key component of the analytic and simulation studies often undertaken in the design of future networks. This paper presents measurements of wide area Internet traffic gathered at the junction between a large industrial research laboratory and the rest of the Internet. Using bar graphs and histograms, it shows the statistics obtained for packet counts, byte counts, and packet length frequencies, broken down by major transport protocols and network services.

1. Introduction

Wide area networks are currently the subject of considerable research activity. The main thrust behind this activity is the deployment of high speed long haul communications technology, which introduces many interesting problems to the design of a network. Much of the research into these problems takes the form of analytical and simulation studies. Both approaches require a suitable traffic model to approximate the behavior of the traffic sources in the network. Experience shows that intuitive notions of how traffic sources behave often lead to unrealistic traffic models, and that a more quantitative approach must be used.

A tried and proven method for developing models of computer systems is to extract the necessary information from existing systems. In the case of high speed wide area networks, it can be argued that future traffic patterns will differ substantially from current ones. For example, these networks are expected to carry real time video and audio traffic that exhibits very different characteristics from more traditional data traffic. Nevertheless, traditional forms of traffic will continue to be carried for quite some time. In particular, protocol families such as the Internet protocol suite will remain in use for the foreseeable future.

It follows that analyzing the traffic patterns in the current Internet will be helpful in future wide area network research. The insight obtained from this analysis can be coupled with what is known about new traffic sources to form realistic traffic models. For example, in the case of deterministic real time traffic such as video, reasonable models can often be formed without experience with a real network. In any case, measurements of present Internet behavior are useful in their own right, since current traffic patterns should prevail for a considerable transition period even as the network evolves.

This paper presents a set of measurements of wide area Internet traffic taken at AT&T Bell Laboratories in Murray Hill, New Jersey. They consist of statistics gathered by tracing network traffic flowing between the Bell Labs corporate network and the rest of the Internet. Section 2 describes the environment in which the traces were obtained, first the network configuration and then the tracing apparatus itself. Section 3 presents the results of the measurements in the form of graphs. They show packet counts, byte counts, and length histograms, including statistics for individual transport protocols and network services.

2. Trace Environment

2.1. Network Configuration

The network configuration at the time the traces were obtained is shown in Figure 1. The intent of the experiment was to gather statistics for wide area Internet traffic. For this reason, the traces were obtained from the Ethernet that linked the Internet gateway machine for the Bell Labs corporate network to the Internet router shown in the figure. The router was connected to Columbia University in New York by means of a 56 Kilobit per second transmission line. The effect was to connect Bell Labs to the JVNC (John Von Neumann Center) regional network of the NSFnet (National Science Foundation network). The regional network in turn connected to the NSFnet backbone, and thus to the rest of the Internet.

2.2. Hardware and Software Configuration

As a result of this network configuration, the Ethernet of Figure 1 carried all wide area Internet traffic coming in and out of Murray Hill. This traffic was inspected by tapping the Ethernet from a DEC uVAX II equipped with a DEQNA Ethernet interface. The interface was placed in promiscuous mode, allowing the collection of all the traffic on that Ethernet without disturbing the network or any of its hosts.

Unfortunately, there was no guarantee that the Ethernet interface could collect every packet that traveled on the cable, especially if high Ethernet loads occurred. The fact that all traffic on that 10 Megabit per second Ethernet was throttled by the aforementioned 56 Kilobit per second line decreased the chances of high loads, but it did not completely eliminate the possibility of packet loss. Beyond this hardware uncertainty, the tracing software insured that no further packet loss occurred, as described

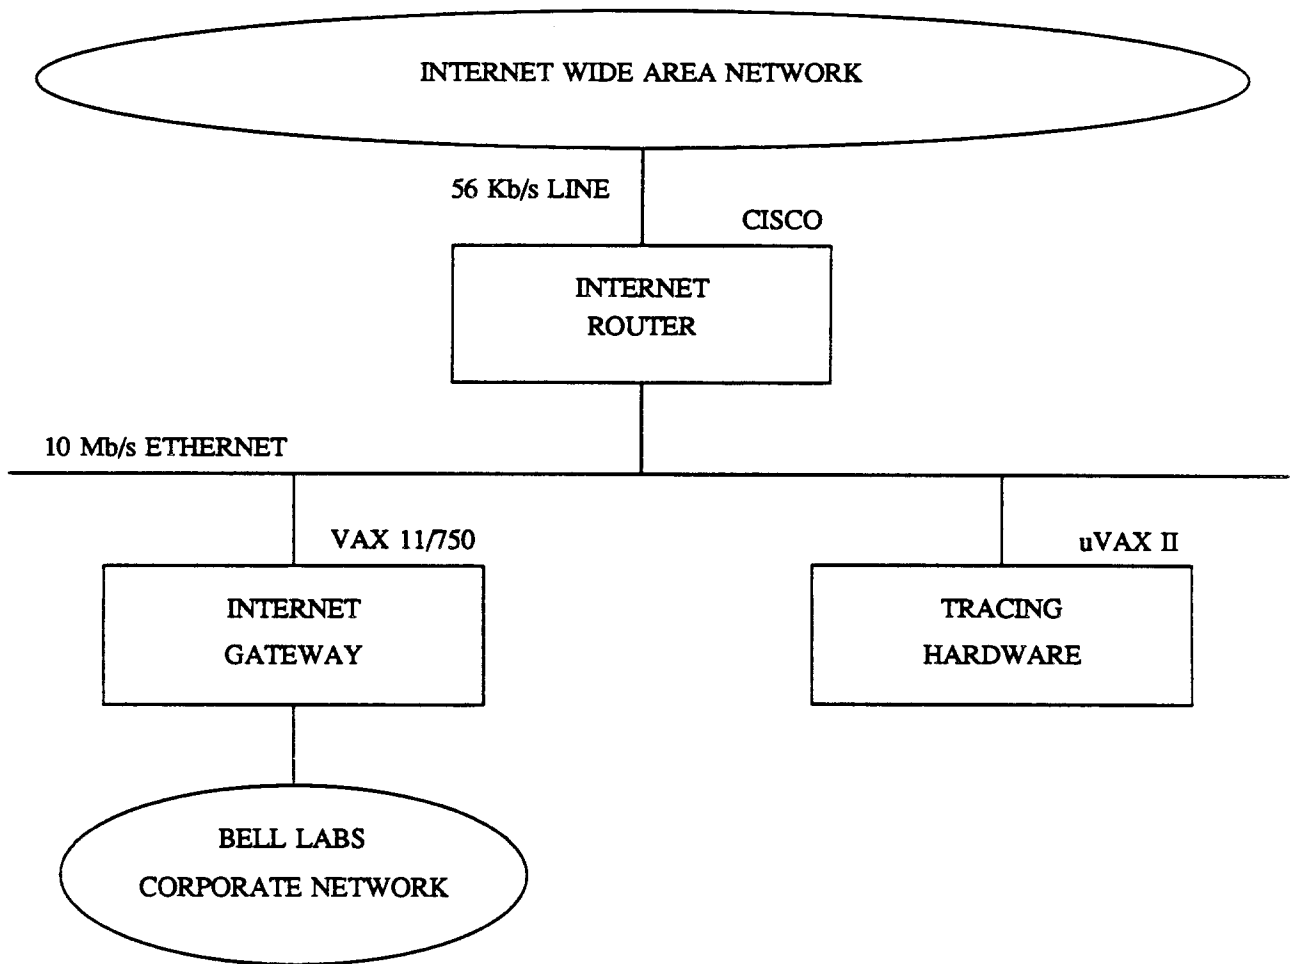


Figure 1 - Network Configuration

below.

The software configuration used in the uVAX II is shown in Figure 2. The system was running Version 9 UNIX and the standard streams device driver for the DEQNA interface. A new streams module was inserted in the kernel between the Ethernet driver and user space. The function of this module was simple: it provided additional buffering in an attempt to insure that no Ethernet data was dropped between hardware read interrupts and user buffer space, it verified whether in fact any data was lost, and it placed a timestamp on every Ethernet packet that flowed through it. No data loss was detected during any of the tracing runs reported here.

The main bulk of the tracing software resided in a user space program. This program placed the Ethernet interface in promiscuous mode and instructed the device driver to give it only Ethernet packets containing IP datagrams. This first level of filtering eliminated all non-IP traffic, such as packets transmitted by ARP (Address Resolution Protocol). The program then looped, reading packets and collecting traffic statistics. Further filtering was accomplished by discarding any Ethernet packets that were not directed to or originating from the Internet router. At the end of the trace, the program wrote the statistics to a set of disk files.

Statistics were gathered in memory as the packets came in, and the raw trace data was immediately discarded. This avoided the I/O overhead of dumping data to disk or tape, but meant that all relevant statistics had to be collected on the first pass as there was no original data left to post-process. Packet length and interarrival time histories were considered sufficient for the purposes of this study.

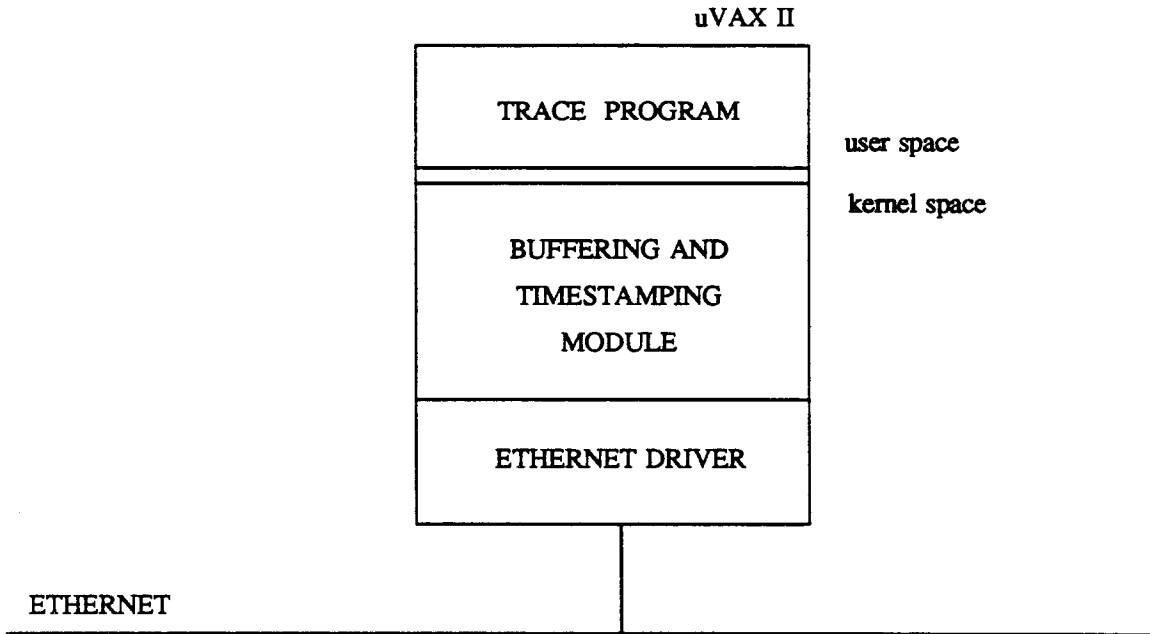


Figure 2 - Software Configuration

Packet count and length statistics were collected on a per network service basis for both TCP (Transmission Control Protocol) and UDP (User Datagram Protocol). Network services were differentiated by a heuristic that computed a hash function on the smaller of the source and destination port numbers used by TCP and UDP. This hash scheme relied on the convention that major Internet services rendezvous with their clients by means of well-known port numbers, which are typically the smaller of the source and destination port numbers used in a conversation. Major TCP network services include FTP (File Transfer Protocol), SMTP (Simple Mail Transfer Protocol), TELNET (remote terminal protocol), and LOGIN (Berkeley remote terminal protocol). Major UDP network services include DOMAIN (name domain service) and ROUTE (routing service).

Packet interarrival time statistics were gathered for the complete traffic stream, without differentiating between TCP, UDP, or individual network services. Arranging for the software to separate the statistics would have been a simple matter, as was done for packet lengths. However, this effort was left for future work because the hardware clock resolution available on the microVAX II was not good enough to extract meaningful interarrival time results. The remainder of this document restricts itself to presenting packet length results.

3. Trace Results

Several traces were collected starting at different days of the week and different times of the day. Whereas interarrival time statistics varied considerably among these traces, the distributions for packet counts, byte counts, and length frequencies remained relatively constant throughout these runs. In the interests of brevity, the following presents results from a single representative trace started on Thursday, July 13, 1989. Roughly speaking, the trace ran for 24 hours between 8:00 AM Thursday and 8:00 AM Friday, and inspected 500,000 IP packets.

3.1. Packet Counts

The first quantity of interest is the number of packets transmitted by each transport protocol and network service. Figure 3 graphs this packet count for TCP and UDP. As shown, TCP accounts for approximately 80% of the packets.

Next we graph the breakdown of packet counts for the individual network services that use TCP and UDP. Figure 4 shows these numbers for TCP. SMTP and FTP account for the majority of the




Figure 3 - Overall Packet Count

packets, with SMTP alone responsible for more than 50% of all TCP packets.

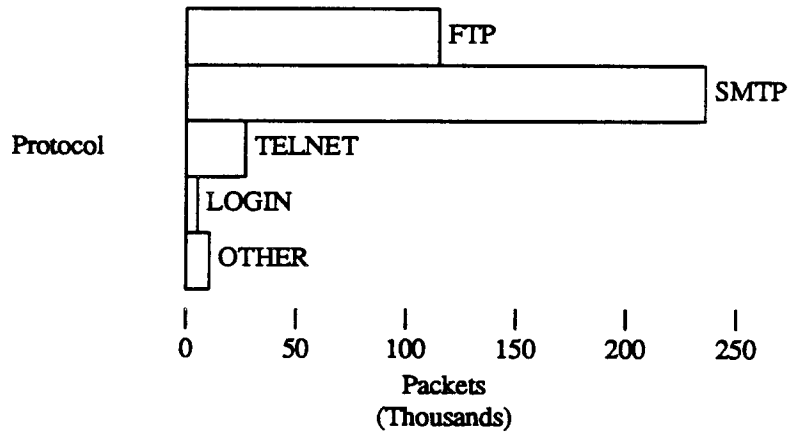


Figure 4 - TCP Packet Count

Figure 5 shows that for UDP, an overwhelming majority of the packets are transmitted by DOMAIN, the Internet name domain service. It is noteworthy that the number of DOMAIN packets is roughly equal to the number of FTP packets for the same time period. This lends weight to the common perception that current Internet name server implementations are creating an unreasonable amount of traffic. A name service such as DOMAIN, intended to support the basic operation of the network, should not be responsible for such a significant percentage of the network load. The packet count shown for ROUTE, the Internet routing service, is much more consistent with this philosophy.




Figure 5 - UDP Packet Count

3.2. Byte Counts

In addition to packet counts, another interesting statistic is the number of bytes transmitted by each protocol and service. The figures reported here correspond to the number of data bytes passed to TCP and UDP by higher layers. That is, they are transport data lengths not including the lengths of the IP, TCP, or UDP headers. The data actually given to the network by IP includes an IP header plus either a TCP or a UDP header. In the absence of IP fragmentation, IP options, and TCP options, these header lengths are: 20 bytes for IP, 20 bytes for TCP, and 8 bytes for UDP.




Figure 6 - Overall Byte Count

Figure 6 graphs the number of bytes associated with TCP and UDP. TCP accounts for more than 90% of the bytes, as compared to 80% of the packets in Figure 3. Figure 7 shows the breakdown of TCP bytes by network service. FTP and SMTP are again the clear leaders, but in terms of bytes FTP slightly overtakes SMTP. In contrast, FTP accounts for less than half the number of packets of SMTP, as per Figure 4. The last graph of this type, Figure 8, shows the UDP breakdown per network service. This byte count graph is very similar to the UDP packet count graph of Figure 5.




Figure 7 - TCP Byte Count

3.3. Length Frequencies

The final statistic reported here is the frequency of the different transport data lengths observed during the trace run. This measure is very useful since, together with interarrival time statistics, it allows the creation of detailed traffic models for the individual protocols and services. These models can then be used for analytical and simulation studies of network behavior.

The first such length histogram is graphed in Figure 9 for all the traffic observed, both TCP and UDP. It has a number of interesting characteristics. Most noticeable are the multiple peaks at particular data lengths. These are better explained by the histograms for the individual protocols and services that will appear below.




Figure 8 - UDP Byte Count




Figure 9 - Overall Length Frequencies

One other feature bears mentioning here: the abrupt absence of data points above 536 bytes, with a lone exception at 857 bytes. This property of wide area Internet traffic is explained by the prevailing policy regarding IP fragmentation of higher level data. The high performance cost of software fragmentation and reassembly dictates that such activity be avoided whenever possible. Since most current IP implementations reside in host software, they adjust their maximum transmission unit parameter so that no fragmentation takes place. For TCP traffic destined for the Internet wide area backbone, as was true of most of the traffic involved in these traces, this maximum transmission unit translates into 536 transport bytes.

Figure 10 shows the length histogram for TCP traffic. The very large peak at 0 bytes corresponds to TCP acknowledgments, which carry only TCP and IP headers with no transport data. TCP is a reliable stream protocol, and these acknowledgments are used for window flow control and for insuring reliable data delivery. They are extremely common, accounting for almost 50% of all TCP packets, and indicate that the ability of TCP to piggy-back acknowledgements on reverse-flowing data is not being used. The absence of piggy-backed acknowledgements is an interesting phenomenon, since state of the art TCP implementations make an effort to use them whenever possible. The most likely explanation is




Figure 10 - TCP Length Frequencies

that most conversations between wide area network services and their clients are intrinsically one-way.




Figure 11 - FTP Length Frequencies

The length frequencies exhibited by FTP are shown in Figure 11. Since FTP is a bulk file transfer protocol, it should produce a lot of large packets. This is certainly the case, as evidenced by the large peak at 512 bytes (recall that 536 bytes is the maximum length of TCP transport data for wide area

Internet traffic, so 512 bytes is indeed large in this context). On the other hand, the peak at 250 bytes is harder to explain; it may be caused by a built-in buffer size in some FTP or TCP implementations. The peak at 0 bytes is again caused by TCP acknowledgements.




Figure 12 - SMTP Length Frequencies

Figure 12 graphs the length histogram for SMTP traffic. As usual, the peak at 0 bytes is caused by TCP acknowledgements. Other significant peaks are found in the 1-100 byte range, and in particular there is a peak at the 1-10 byte range. Aside from data points created by occasional very short mail messages, these frequencies are explained by the handshaking that takes place at the beginning and end of SMTP conversations. The contents of SMTP handshaking messages are usually short ASCII strings similar to "HELLO hostname." The peaks at 512 and 536 bytes are easily explained by maximum-sized transmissions due to large mail transfers, similar to the large file transfers carried out by FTP. Once again, however, the peak at 250 bytes is harder to explain. Aside from an unlikely propensity of mail messages to contain exactly 250 bytes, a possible explanation is a fixed-size buffer built into current SMTP or TCP implementations. A full explanation must await further investigation.

Figures 13 and 14 show the length frequencies for TELNET and LOGIN traffic, respectively. These graphs agree with intuitive notions of remote terminal traffic. Aside from the TCP acknowledgement peak at 0 bytes, they show their biggest peaks in the 1-10 byte range. This is naturally explained by the single-character ASCII terminal traffic that results from human typing and its ensuing echo, along with a few double-character carriage return and line feed combinations. The rest of the traffic is caused by the larger transfers that occur when a host sends back program output to the remote terminal. The mysterious 250-byte peak is again noticeable in the LOGIN graph, as well as the more easily explained 512-byte peak.

Figure 15 graphs the length histogram for all the UDP traffic. Unlike TCP, UDP is a datagram protocol and does not use acknowledgements. Therefore, it does not produce any traffic with 0 bytes of transport data. All the peaks in the UDP histogram are caused by network service data, and are explained by other graphs below.

Figure 16 graphs the length histogram for DOMAIN traffic. A large portion of this traffic lies in the 20-40 byte range, with a very large peak at 29 bytes and a smaller one in the 30-40 byte range. These numbers can be explained by the nature of name server communication, which is composed




Figure 13 - TELNET Length Frequencies




Figure 14 - LOGIN Length Frequencies

mainly of simple database queries and responses.

Finally, Figure 17 graphs the length frequencies for ROUTE traffic. As shown, the routing traffic fell into two very well-defined categories. In fact, only 4-byte and 24-byte data lengths were observed.




Figure 15 - UDP Length Frequencies




Figure 16 - DOMAIN Length Frequencies

4. Conclusion

This paper presents statistics of wide area Internet traffic by means of bar graphs and histograms. It shows the measurements obtained for packet counts, byte counts, and length frequencies, without elaborate interpretation or analysis. Packet interarrival times, although gathered during tracing, are not




Figure 17 - ROUTE Length Frequencies

presented due to the inadequate resolution of the clock used to perform the measurements.

Raw data of the type presented here is both interesting and useful. Already these figures have been used to study the transmission line efficiency that results from using the ATM (Asynchronous Transfer Mode) networking standard to transport long haul Internet traffic. Basing such calculations on real wide area network data is more valid than basing them on any intuitive perception of network behavior.

In future work, the same experiments should be repeated with better clock hardware in order to obtain useful interarrival times. In addition, the experiments should be run at different network sites for comparison and validation purposes. Of particular interest are the junction points between the current NSFnet regional networks and the NSFnet wide area backbone. At these points, extensive multiplexing of different traffic types occurs. It is important to offer good performance to each type of traffic without wasting network resources. Traces from these points will yield exclusively long-haul traffic with the interarrival time distribution appropriate for such a junction, enabling an accurate study of the cost-performance tradeoffs present there.

The measurement methodology described in this paper is a valuable approach to network research. The data thus obtained from a current network in operation can be used directly to realistically model network behavior. Creating such models is the first step towards any meaningful analytical or simulation study of computer networks.

Acknowledgements

I am grateful to several people for their assistance with this work. Sandy Fraser suggested I use packet traces to generate traffic models for network simulations, and offered guidance throughout the experiments. Dave Presotto provided the tracing hardware, set up the network configuration, and assisted with the V9 kernel. Dennis Ritchie was helpful in programming the streams module.