A Fluid-Flow Characterization of Internet1 and Internet2 Traffic

Joe Rogers and Kenneth J. Christensen
Department of Computer Science and Engineering
University of South Florida
Tampa, Florida 33620
{*jrogers, christen*}*@csee.usf.edu*

Abstract

We study the characteristics of Internet1 and Internet2 traffic at a network access point at a major US university. With the use of fluid-flow modeling, we show that Internet1 and Internet2 traffic have different queuing behaviors and that a small percentage of traffic on both networks largely contributes to this overall queuing behavior. We also demonstrate that buffer sizing, as a method to reduce loss, is largely ineffective for Internet2 traffic. These finding have implications to Quality of Service of Internet applications.

1. Introduction

The Internet2 is a high-speed international network interconnecting research institutions throughout the world [4]. The Internet2 is intended to aid collaborative research by carrying application traffic specific to research. This is different from the largely commercial traffic found on the Internet1.

Our goal is to understand the characteristics of the Internet1 (I1) and Internet2 (I2) traffic and to compare the differences in queuing behavior. Of special concern is the possible presence and effects of self-similarity in traffic on the two networks [7, 8]. An understanding of traffic characteristics can be applied to better application design and improved network traffic control and Quality of Service (QoS) algorithms. In this work we use fluid-flow techniques to measure buffer length and loss of traced traffic from operational I1 and I2 links.

2. Traffic data collection

The University of South Florida is a large (approximately 36,000 student) Research I institution with over 15,000 network connections throughout its four campuses. Figure 1 shows the network configuration. USF is in a unique position to study differences in I1 and I2 traffic characteristics, since it currently operates 45Mbps (DS3) links into both the I1 and I2. These links are served through Cisco 7200 (I1) and 7500 (I2) routers connecting to Sprint and Abilene [1], respectively.

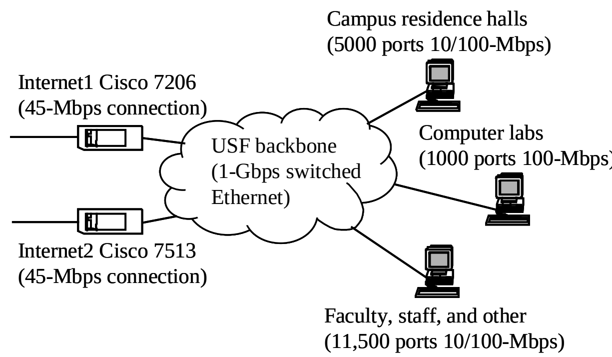

Figure 1. Network configuration at USF

Cisco routers natively collect traffic flow information [2]. A flow is defined as a unidirectional, sequence of packets between a given source and destination. Flows are delimited by SYN/FIN pairs in the case of TCP sessions and by pre-configured flow timeouts in the case of UDP. Cisco NetFlow records contain information including source/destination IP address and port numbers, number of packets/bytes sent, and start/end times for a given flow. From the flow records we derive flow *size* (bytes in a flow), *rate* (average bytes per second), and *length* (time of flow). Packet-level inter-arrival time characteristics cannot be derived from flow data.

We collected busy hour flows between 1:30pm and 2:30pm on five sequential weekdays between 2/20/01 and 2/26/01. We chose to focus on flow size since the 1% of flows by size (versus flow length or rate) yielded the most significant difference in the number of bytes between the top 1% and remaining 99% of flows. Figure 2 shows the cumulative probability of flow size. A much greater mass of flow size occurs in the tail of the distribution of the I2 than for the I1. Summary statistics for flow size are shown in Table 1. The utilization of the I1 is about 25 times greater than the I2. Flow sizes on the I2 are about 3.5 times larger than flow sizes on the I1. However, the Coefficient of Variation (CoV) of flow sizes is about the same. Tables 2 and 3 provide basic statistics on the top 1% and remaining 99% subsets. For the I1, the top 1% size flows contain 83% of the bytes carried. For the I2, 93% of bytes are carried in the top 1% flows.

This material is based upon work supported by the National Science Foundation under Grant No. 9875177.

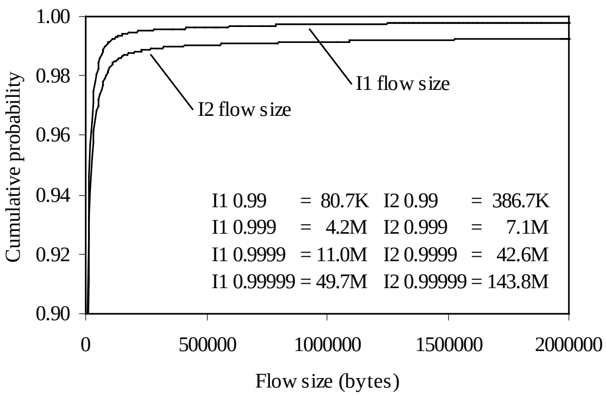

Figure 2. Cumulative probability of flow size

Table 1. Summary statistics on collected flows

Measure	Internet1	Internet2	% diff
# of flows	9,563,940	390,694	-96 %
Max size	135,520,768	264,860,744	95
Total bytes	134,549,356,965	19,111,812,476	-86
Mean	14,069	48,917	248
Std Dev	335,850	1,120,321	234
CoV	23.9	22.9	-4

Table 2. Summary statistics on I1 flows by size

Measure	I1 Top 1%	I1 remain 99%	% diff
# of flows	95,634	9,467,857	9800 %
Total bytes	111,106,109,720	23,443,247,245	-79
Mean	1,161,785	2,476	-99
CoV	2.71	2.92	8

Table 3. Summary statistics on I2 flows by size

Measure	I2 Top 1%	I2 remain 99%	% diff
# of flows	3,906	386,788	9802 %
Total bytes	17,686,663,404	1,425,149,072	-92
Mean	4,528,076	3,685	-99
CoV	2.27	4.10	81

From the NetFlow records, an application distribution is obtained summarizing the TopN port numbers by total bytes transferred. The flows were collected prior to the legal restrictions placed on the Napster music-sharing system. Table 4 shows that several top applications are common to both the I1 and I2.

Table 4. Top N applications on I1 and I2

Application (port)	I1 total bytes and rank	I2 total bytes and rank
Napster (6688/6699)	50.32 % (1)	65.75 % (1)
Web (80)	26.51 (2)	6.48 (3)
Quicktime (6970)	2.62 (4)	1.63 (9)
Gnutella (6346)	2.08 (5)	1.97 (7)
NNTP (119)	NA	8.63 (2)

3. Fluid-flow evaluation of traffic effects

Cisco NetFlow records can be used to determine traffic rates on a link as a function of time. A C program was written to generate a time series of rates from collected NetFlow records. A fluid-flow model [6] uses this first-order statistic (rate) to determine buffer occupancy (buffer length) and overflow. A fluid-flow model takes as input the rate of traffic arrivals as a function of time, the departure rate (which is the egress link speed), and buffer size. If the arrival rate exceeds the departure rate, buffer occupancy builds up at a rate that is the difference between the arrival and departure rate. If the arrival rate is less than the departure rate, buffer occupancy decreases until the arrival rate increases (to be equal or greater than departure rate) or the buffer reaches zero (empty). Fluid-flow models cannot account for the length of packets and cannot account for second order statistics in packet interarrival times (and hence are stable even at 100% utilization).

We use a fluid-flow model as a first-order approximation of the effects of I1 and I2 traffic on a bottleneck queue. A fluid-flow model [3] was written in C that takes as input a series of rate values derived from the NetFlow records. We call the input file a “trace file”. The rate values are assumed to be taken at periodic time intervals (we used 1 second as the time interval because this is the minimum granularity of a NetFlow record). The outputs from the program are buffer length and overflow bit counts on 1-second intervals, overall mean buffer length, and overall percentage of bits overflowed. To evaluate the effects of link utilization, the departure rate of the fluid-flow model was fixed to achieve a desired utilization value for a given trace input. For a trace of length T seconds with total B bits and a desired link utilization of U ($0 < U \leq 1$), the departure rate is,

$$R_d = \frac{B}{T \cdot U}. \quad (1)$$

Figure 3 shows the fluid-flow model with input from trace files.

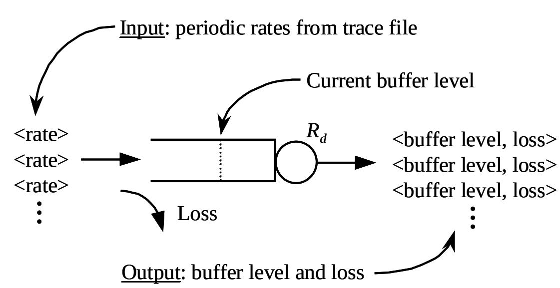

Figure 3. Fluid-flow model with trace file input

3.1 Experiments

We measure the effects of link utilization and buffer size on mean buffer length and loss. We define three experiments where the traffic input was the full set of I1 and I2 traffic traces, top 1% (by flow size), and the remaining 99% (by flow size) of traffic traces.

Utilization experiment #1: This experiment measures the mean buffer length (for an infinite sized buffer) for a range of link utilization. We vary the link utilization from 10% to 100%.

Utilization experiment #2: This experiment measures the effect of utilization on loss for two buffer sizes (2 and 8 Mbytes). We vary the link utilization from 10% to 100%.

Buffer size experiment: This experiment measures the effect of buffer size on loss for a fixed link utilization of 90%. Buffers size is varied from 1 Mbyte to the size needed to eliminate loss.

3.2 Experiment results

Figure 4 shows the results for the utilization experiment #1. It can be seen that I2 mean buffer length for the full trace is larger than the I1 mean buffer length, except at high link utilization. It can also be seen that the top 1% flows result in a much greater buffer length than does an input of the remaining 99% flows. Figures 5, 6, and 7 show the results for the utilization experiment #2. In Figure 5 it can be seen that loss is consistently higher for the I2 over the entire range of utilizations. It can also be seen that a small reduction in link utilization in the I1 can eliminate loss, but this is not the case for the I2. In Figure 6 (I1 results), the top 1% flows have less loss than the remaining 99% flows. Figure 7 (I2 results) shows the opposite – that the top 1% flows have greater loss than the remaining 99% flows. For both the I1 and I2, increasing the buffer size has some effect (on reducing loss).

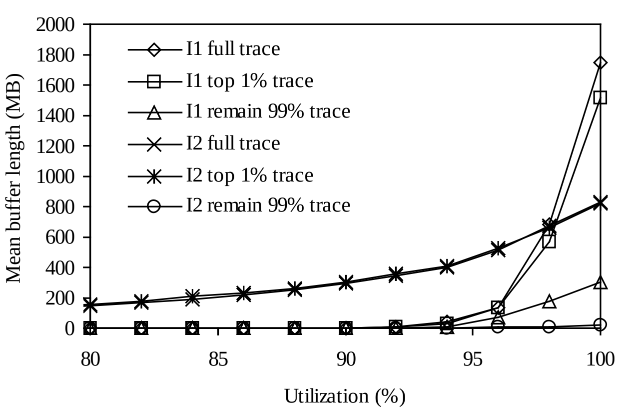

Figure 4. Utilization experiment #1

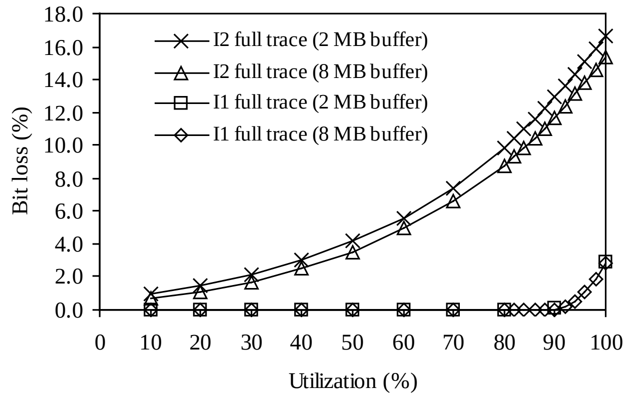

Figure 5. Utilization experiment #2 – I1 and I2 flows

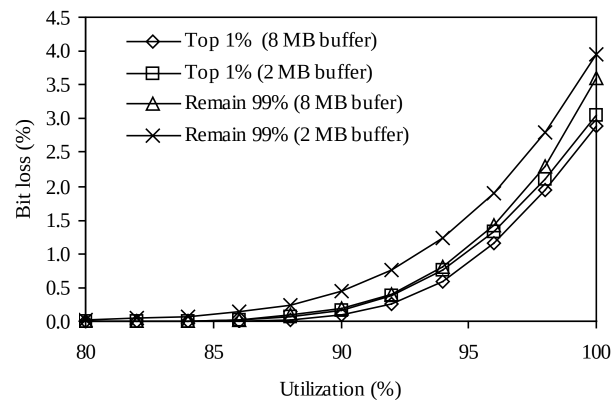

Figure 6. Utilization experiment #2 – I1 flows

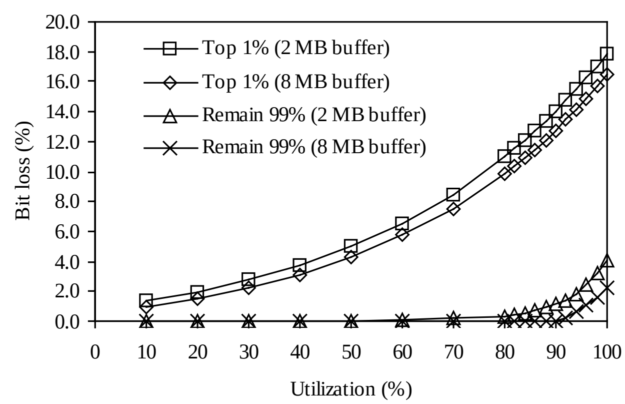

Figure 7. Utilization experiment #2 – I2 flows

Figure 8 shows the full trace results for the buffer size experiment where it can be seen that I1 losses are already virtually 0% at 1 Mbyte buffer size, but I2 losses require a 1500 Mbyte (!) buffer to be reduced to 0%. The results for the top 1% flows are the same (as Figure 8). Figure 9 shows the results for the remaining 99% flows.

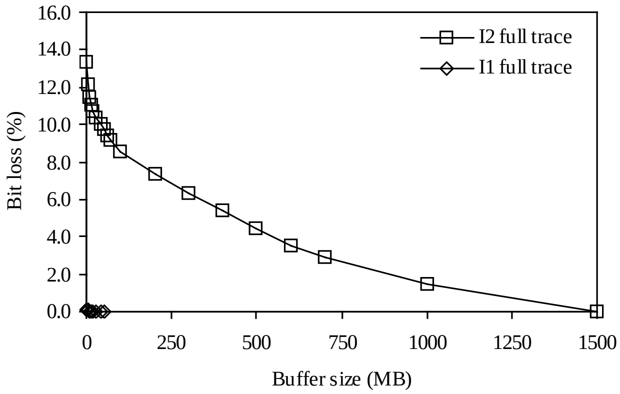

Figure 8. Buffer size experiment – I1 and I2 flows

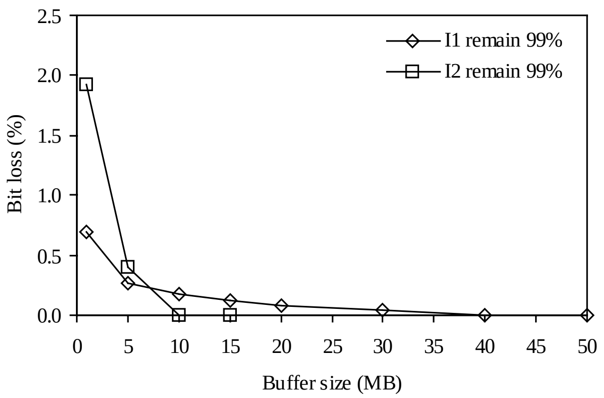

Figure 9. Buffer size experiment – Remaining 99%

4. Discussion of results

Our experiments demonstrated that I2 traffic results in significantly greater mean buffer length and loss than I1 traffic for a given link utilization. Very significant is the result shown in Figure 8 where an unreasonably large buffer size is needed to contain all losses for the I2. It is also notable that 1% of all flows carry the majority of bytes and have a much larger effect on buffer length and loss than the remaining 99% of all flows. These results are very suggestive of (but, *do not prove*) heavy-tailed traffic on the I2, but not on the I1. Many of these findings that suggest heavy-tailed traffic are consistent with those of Park et al. [7], which were obtained with simulated traffic. An immediate benefit from our work is in understanding that most of the bit loss for a full set of flows is caused by the largest in size 1% of flows; we can concentrate our efforts on studying and possibly controlling the applications contributing to this small percentage of flows.

Table 5 shows the Top N applications with the top 1% of flows. The breakdown of applications is very similar to that of Table 4 suggesting that the dominant applications in the full trace tend to be sources for some of the largest size flows. As a first step to determining the effects of shaping of the top 1% of flows, we manipulate these flows and reproduce the previous experiments. We artificially shape the rate of the top 1% flows by changing the number of bytes carried by one half. This is a very first-order shaping. The R_d, was recalculated to achieve the target link utilization. Figure 10 shows the results from the utilization experiment #2. Losses for the “HalfByte” case are significantly lower than for the full traces. Figure 11 shows how the buffer size required to achieve 0% loss is reduced to about half of the full trace case. These results are very promising in showing that shaping only a very small percentage of flows can yield large overall benefits to reduction of loss.

Table 5. Top N applications on I1 and I2 - 1% flows

Application (port)	I1 total bytes and rank	I2 total bytes and rank
Napster (6688/6699)	57.49 % (1)	63.09 % (1)
Web (80)	11.51 (2)	1.86 (3)
Quicktime (6970)	2.71 (4)	2.90 (9)
Gnutella (6346)	2.84 (5)	2.46 (7)
NNTP (119)	NA	13.69 (2)

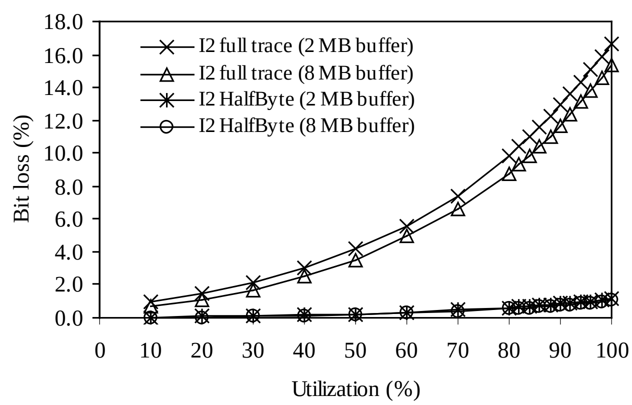

Figure 10. Utilization experiment #2 – half flows

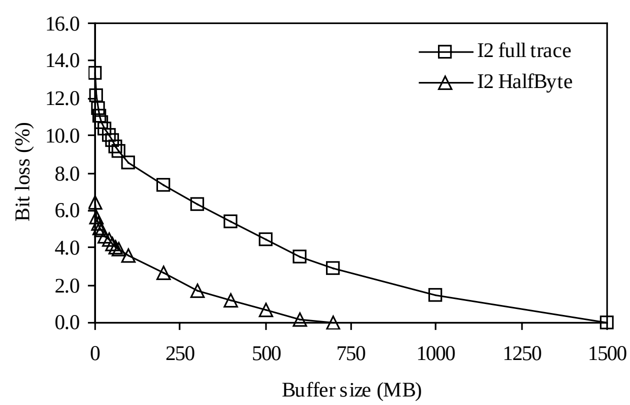

Figure 11. Buffer size experiment – half flows

5. Summary and future work

We have used a fluid-flow model to gain insights into possible differences between Internet1 and Internet2 traffic. We believe that this is the first use of Cisco NetFlow traces as input to a fluid-flow model to study queueing behavior of an IP network. With the fluid-flow model, we have shown that Internet1 and Internet2 traffic results in very different queueing behaviors, many of which suggest that Internet2 traffic is self-similar. Additional work, including verifying stationarity of the collected busy-hour traffic, is needed to prove the self-similarity exists on the I2, but not on the I1.

A practical benefit from our work is in understanding that most of the bit loss for a full set of flows is caused by the largest in size 1% of flows; we can concentrate our efforts on studying and possibly controlling the applications contributing to this small percentage of flows. A first order shaping of these 1% flows in our model resulted in significant reductions in loss. Efforts are underway to begin rate limiting certain non-research applications at USF. This work will serve as an input to these shaping decisions.

Future work will measure the benefits of shaping in terms of improved QoS for all applications. The collected I1 and I2 NetFlow records are available from the authors. We are currently working to include this collected traffic data on the Internet Traffic Archive [5] for other researchers to use.

Acknowledgements

The authors would like to thank Ted Netterfield for his help with traffic collection.

References

[1] Abilene, 2001. *http://www.abilene.iu.edu*.

[2] Cisco NetFlow Services and Applications, 2000. *http://www.cisco.com/warp/public/732/NetFlow/*.

[3] Fluid-flow, single-server, finite-buffer queue simulation (K. Christensen, 2001), software download from *http://www.csee.usf.edu/~christen/tools/toolpage.html*.

[4] Internet2, 2001. *http://www.internet2.edu*.

[5] The Internet Traffic Archive, 2001. *http://www.acm.org/sigcomm/ITA/*.

[6] L. Kleinrock, "Queueing Systems, Volume 2: Computer Applications," John Wiley and Sons, New York, 1976.

[7] K. Park, G. Kim, M. Crovella, "On the Effect of Traffic Self-Similarity on Network Performance", *Proceedings of the SPIE International Conference on Performance and Control of Network System*s, pp. 296-310, November 1997

[8] Z. Sahinoglu and S. Tekinay, "On Multimedia Networks: Self-Similar Traffic and Network Performance," *IEEE Communications Magazine*, Vo. 37, No. 1, pp. 48-52, January 1999.