PASSIVE MONITORING OF INTERNET TRAFFIC AT SUPERCOMPUTING'98

Brynjar Åge Viken
e−mail: brynjar@item.ntnu.no, bviken@nlanr.net
National Laboratory for Applied Network Research, Measurement and Operations Analysis Team*
San Diego Supercomputer Center
USA

Tel: + 47 73 55 02 68
Fax: + 47 73 59 44 66
Department of telematics
Norwegian University of Science and Technology
O.S. Bragstadsplass 2A
N−7491 Trondheim
Norway
(*Most of this work was carried out as a visiting researcher at NLANR/MOAT, Sept. 98 − Feb. 99.)

ABSTRACT

Passive measurement data is collected without any impact on the network by the measurement itself. OCXmon/Coral is a passive standalone monitor that collects packet traces from a high capacity link by tapping a small fraction of the light with an optical splitter. SCinet'98 was the show floor network for Supercomputing'98 where the latest in systems, applications and services for all areas of high performance networking and computing were demonstrated. NLANR/MOAT[1] monitored the traffic on the OC3 link connecting the conference network to the vBNS. The measurement data collected consist of 79.7 million packets, 34.3 million packets sent from SCinet'98 to the vBNS and 45.5 million packets sent from the vBNS to the SCinet'98. Based on these measurement data, the traffic composition for IP protocols, packet length distributions and interarrival times are presented. Further, the NLANR/MOAT demonstration of distributed real−time 3D visualization of abstracted OC3mon/Coral data, presented at Supercomputing'98, is briefly described.

1. INTRODUCTION

In the recent years the Internet has experienced a tremendous growth and a success which is beyond what most people believed possible only a few years ago. As a consequence of these developments, the Internet has become increasingly more important to organizations and individuals. There has been an explosive growth in the number of hosts connected and at the same time new applications with new traffic patterns have been introduced to the network. We have also experienced an increased commercialism of the Internet and following that an explosion in the number of competing Internet Service Providers (ISPs). The Internet is becoming an important infrastructure. Therefore, understanding the composition and dynamics of Internet traffic is of great importance for network engineering, planning and design. So far have closed form analysis of Internet traffic had little success in the heterogeneous Internet environment. Thus, it seems as in the years to come measurements and simulations will be the main tools in studying performance of the Internet.

1National Laboratory for Applied Network Research, Measurement and Operations Analysis Team

1.1 MEASUREMENT APPROACHES

Measurement data can be collected in two principal ways; actively [1] [2] [3] by insertion of test traffic or passively [4] [5] [6] [7] by observing user generated traffic. The monitoring can be performed by standalone units or be router−based. One example of a passive standalone monitor is OCXmon/Coral that collects measurement data from an optical fiber by tapping a small fraction of the light with an optical splitter. Studies of passive measurement data can be found in e.g. [4], [8] and [9].

2. OCXMON/CORAL

The OCXmon/Coral [4] is a passive stand−alone monitor that enables collection of packet traces on a high capacity link without any impact to the network by the measurement itself. Two reference implementations for OC3mon exist, one for Unix [10] and one for DOS [11]. The current configuration of an OC3mon/Coral[2] monitor for Unix is a personal computer consisting of a 400 MHz Intel Pentium II processor, 128 MB of main memory, a 33 MHz 32 bit−wide PCI bus, two SCSI hard disks, an Ethernet interface card and two FORE ATM interface cards

2Current development is focused on OC3 (OC3/Coral) and OC12 (OC12/Coral), but other possibilities are under consideration.

(PCA−200E) each with an Intel i960 processor. By tapping light from the fibers with two optical splitters each of the two ATM interfaces independently capture traffic in one direction of the OC3 link.

The computer runs FreeBSD as an operating system. The custom−developed OC3mon/Coral firmware is implemented in C++ and assembly code and supports several modes of operation. Depending on the software running, the OC3mon can capture packet traces to file, perform flow analyses, real−time analyses or act as a server for visualization.

3. MONITORING SCINET’98

The Supercomputing’98 [12] conference, "High Performance Networking and Computing", was held at the Orange County Convention Center in Orlando, Florida, USA, November 7 − 13, 1998. SCinet’98 was the show floor network where the latest in systems, applications and services for all areas of high performance networking and computing were demonstrated. NLANR/MOAT monitored the OC3 link connecting the conference network to the vBNS. The OC3mon/Coral monitor was configured to capture the first ATM cell of every IP packet (AAL5 frame) and write packet traces to disk.

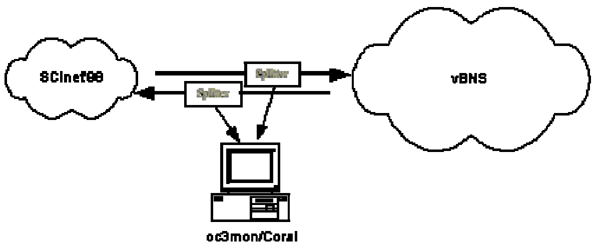

Figure 1. Measurement instrumentation at SCinet’98

The measurements collected 10.−12. November consist of several 5. minute long traces and one continuos trace with a duration of several hours. This paper presents some results based on these packet traces. The results are presented individually for each direction of the connection between SCinet’98 and the vBNS. The measurements totally consist of 79.7 million packets, 34.3 million packets sent from SCinet’98 to the vBNS and 45.5 million packets sent from the vBNS to the SCinet’98. Figure 2 shows the time of day (local time) when the traces were collected and length of each trace in minutes. Figure 3 shows number of bytes observed in each of the collected traces.

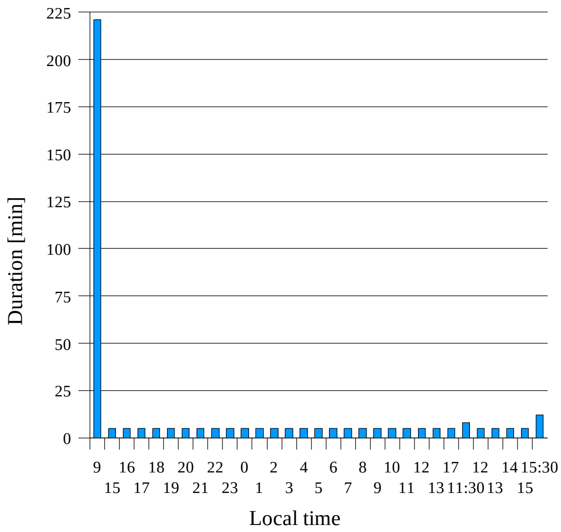

Figure 2. Duration of collected samples

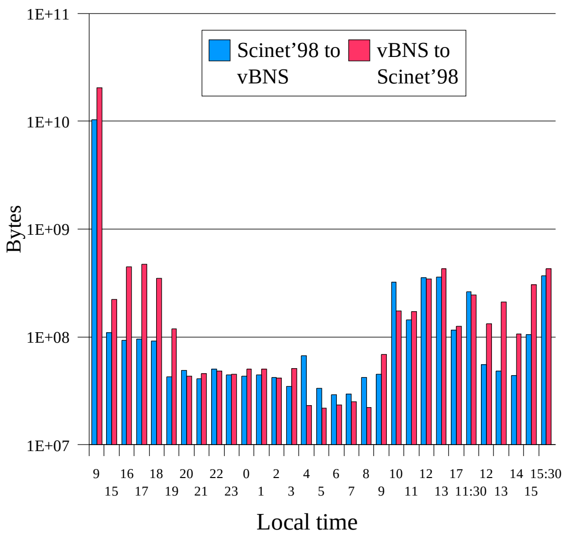

Figure 3. Traffic volume in bytes pr. collected trace.

3.1 TRAFFIC COMPOSITION

Figure 4, 5, 6 and 7 show the traffic composition of bytes and packets by IP protocols. We observe that a rather large portion of the traffic in both directions, specially on daytime, was generated from applications running UDP. The traffic pattern was also asymmetric, as shown in figure 3, since more data was imported than exported from the conference network.

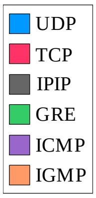

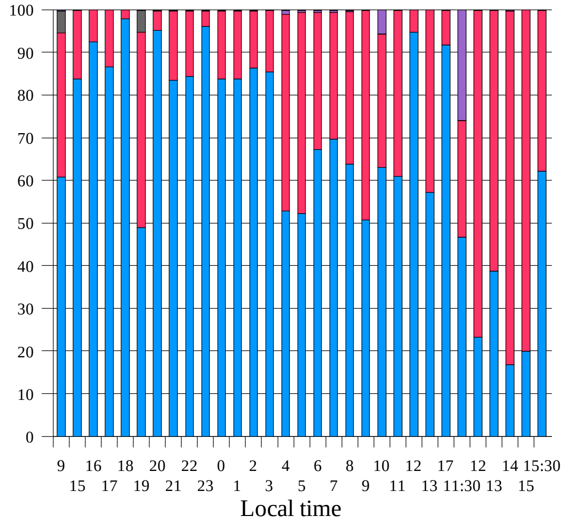

Figure 4. Percentage bytes vBNS –> SCinet'98

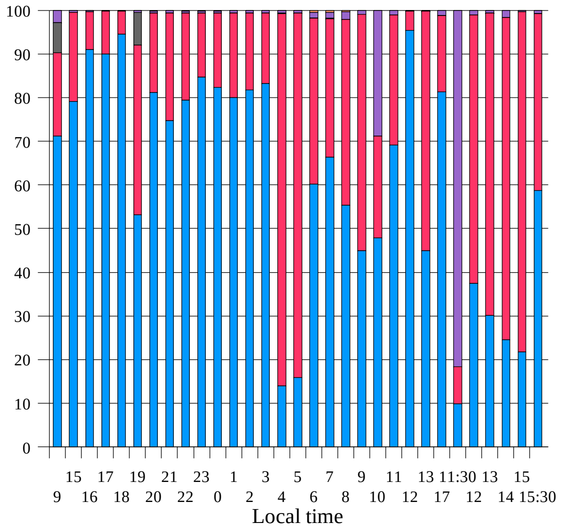

Figure 5. Percentage packets vBNS –> SCinet'98.

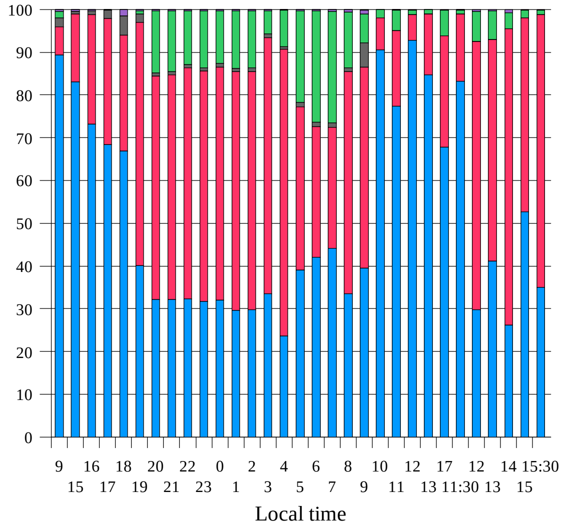

Figure 6. Percentage bytes SCinet'98–>vBNS

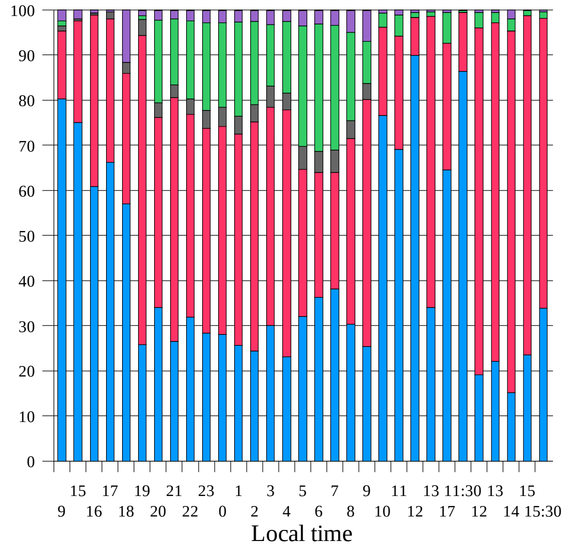

Figure 7. Percentage packets SCinet'98–>vBNS

3.2 IP PACKET LENGTH

The analyses of packet length is from the trace collected continuously between 9:47 and 13:28 on Tuesday 10. November 98. Figure 8 shows the average packet size for major IP protocols in each direction of the connection between SCinet'98 and the vBNS.

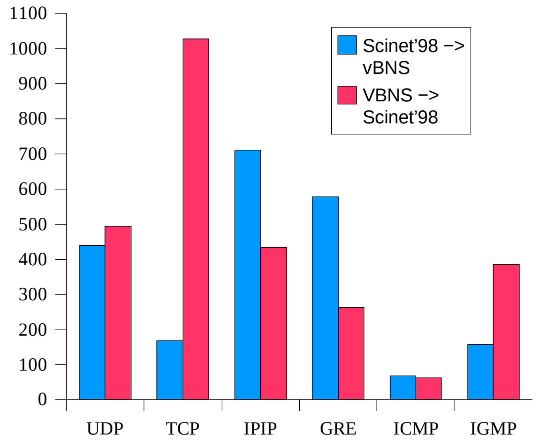

Figure 8. Average packet length.

Figure 9 and 10 show the packet length distributions. As expected we observe that the packet length distribution is multimodal. This is because the load is generated by a mixture of applications with different characteristics.

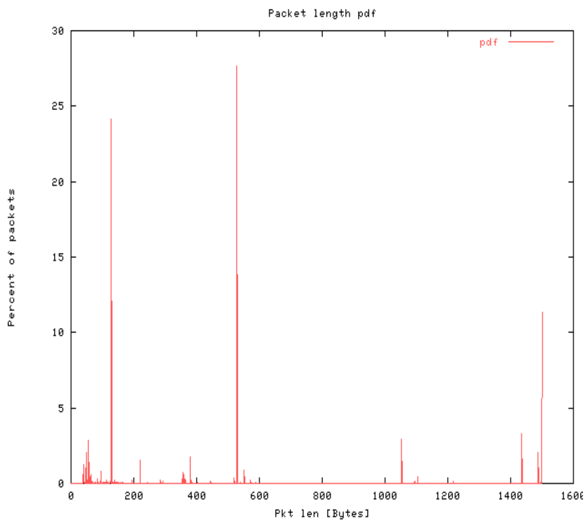

Figure 9. Packet length distribution, vBNS -> SCinet'98

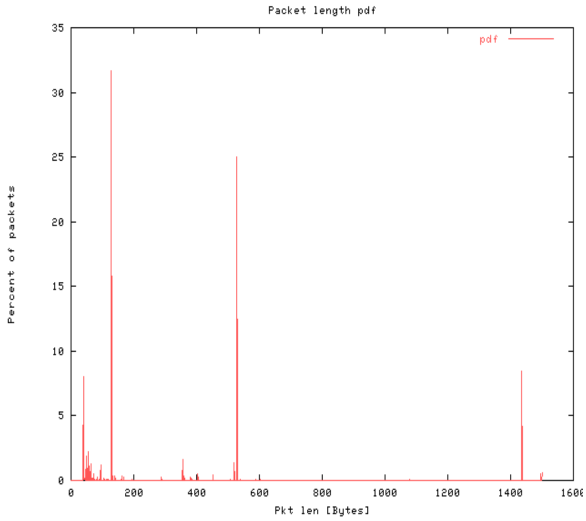

Figure 10. Packet length distribution, SCinet'98->vBNS

3.3 INTERARRIVAL TIMES

The analyses of interarrival times is also from the trace collected continuously between 9:47 and 13:28 on Tuesday 10. November 98. Figure 11 and 12 show the distribution of interarrival times in each direction of the connection.

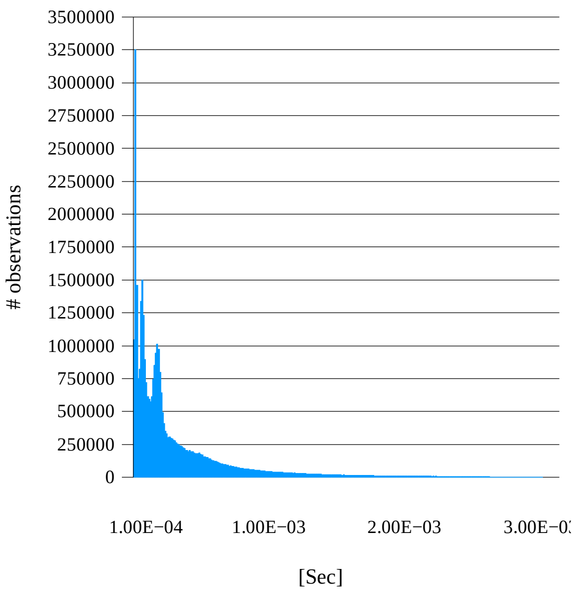

Figure 11. Interarrival times vBNS -> SCinet'98

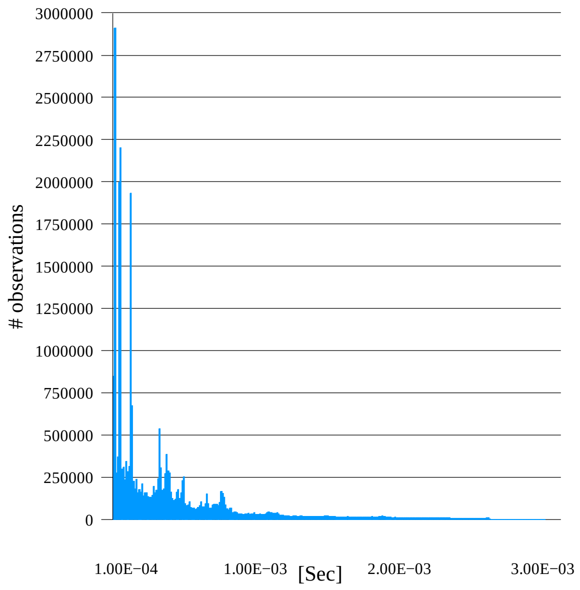

Figure 12. Interarrival times SCinet'98 −>vBNS

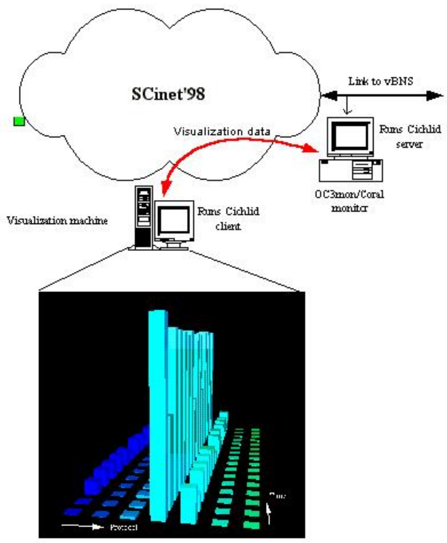

Figure 13. OC3mon/Coral visualization.

4. DISTRIBUTED REAL−TIME 3D VISUALIZATION OF OC3MON/CORAL MEASUREMENT DATA

During the Supercomputing'98 conference NLANR/MOAT demonstrated distributed real−time 3D visualization of OC3mon/Coral measurement data. The demonstration was based on a generic tool for rapidly visualizing of arbitrary data sets in high−quality 3D, Cichlid [13], developed by NLANR/MOAT. A Cichlid server ran on the OC3mon/Coral monitor and processed the collected measurement data to generate data matrices. The matrices were displayed real−time in 3D on a remote visualization client. Graphs showing various characteristics of the traffic traversing the link were projected at the network operation center [14].

5. SUMMARY

We have presented some results from measurements collected at the Supercomputing'98 show floor network. SCinet'98 represented a demanding environment where new multimedia and real−time applications were demonstrated. The packet traces collected contains the signatures of these applications. UDP was the dominant IP protocol in both directions of the connection between SCinet'98 and vBNS, specially on daytime. The UDP load was generated by applications that did not use the traditional TCP flow and congestion control. One might imagine that some of these were demonstrations of new applications with "real−time" requirements. Studies of LAN and WAN Internet traffic [4] [8] [9] have shown that TCP normally is the largest contributor to Internet traffic volume, but the mix of applications used in the SCinet'98 environment was unusual. The IP packet length distribution is multimodal and determined by the characteristics of the applications generating the load onto the network. Thus, average IP packet length holds little information since a variety of applications generate the traffic. We also briefly described the NLANR/MOAT demonstration of distributed real−time 3D visualization of abstracted OC3mon/Coral data presented at Supercomputing'98. The demonstration used the generic Cichlid tool developed by NLANR/MOAT.

Passive monitoring is an important tool to understand the composition and dynamics of Internet traffic. In the years to come, as service differentiation is introduced to the Internet, passive monitoring will probably also be used to collect resource usage data as an input to pricing models.

ACKNOWLEDGE

National Science Foundation Cooperative Agreement No. ANI−9807479, and the National Laboratory for Applied Network Research.

Hans−Werner Braun, Tony McGregor and Jeff Brown at NLANR/MOAT .

The SCinet'98 crew and specially David Koester for his initiative to perform the OC3mon/coral monitoring of SCinet'98.

REFERENCES:

[1] NLANR/MOAT Active Measurement Program, http://amp.nlanr.net/.
[2] http://www.ripe.net/test−traffic/index.html
[3] Advanced Network & Services and the Common Solutions Group, Surveyor project, http://www.advanced.org/surveyor
[4] J. Apisdorf, K. Claffy, K. Thompson and R. Wilder, OC3MON: Flexible, Affordable, High−Performance Statistics Collection, Proceedings of INET'97, http://www.isoc.org/isoc/whatis/conferences/inet/97/proceedings/F1/F1_2.HTM, Kuala Lumpur, Malaysia, June 1997.
[5] Nevil Brownlee, NeTraMet, http://www.auckland.ac.nz/net/NeTraMet/.
[6] D. McRobb and J. Hawkinson, "cflowd: a Cisco flow−export collector", http://figaro.ans.net/cflowd, 1997.
[7] "Cisco Systems Inc., NetFlow Switching Traffic Managemen", White Paper, 1996
[8] R. Careces, P. Danzig,S. Jamin and D. Mitzel, Characteristics of Wide−Area TCP/IP Conversations, Proceedings of ACM SIGCOMM'91 pp. 101−112, Sept. 91.
[9] K. Claffy, G. Miller and K. Thompson , The nature of the beast: recent traffic measurements from an Internet backbone, http://www.caida.org/Papers/Inet98/index.html, April 1998.
[10] NLANR/MOAT homepage, http://moat.nlanr.net/
[11] Cooperative Association for Internet Data Analysis (CAIDA), CoralReef web site, http://www.caida.org/Tools/CoralReef/
[12] Supercomputing98: High Performance Networking and Computing, sponsored by IEEE Computer Society and ACM SIGARCH, http://www.supercomp.org/sc98/.
[13] NLANR/MOAT, Cichlid, http://www.moat.nlanr.net/Software/Cichlid/
[14] NLANR/MOAT, Supercomputing'98 demos, http://www.moat.nlanr.net/SC98Demos/.