NETWORKING AND COMMUNICATIONS PARALLEL SESSION SUMMARY

PAUL JEFFREYS

CCLRC

April 1997

Introduction

This paper summarises the talks given in the parallel session on 'Networking and Communications' at CHEP'97. There is a logical thread which is maintained through the report and culminates in a firm conclusion.

Video Conferencing and Collaborative Environments

Video conferences can be summarised as being of two types: circuit switched and packet switched. The table below summarises the important characteristics of each:

Circuit switched	**Packet switched**
• End-end connections a priori	• No connections
• Resources are reserved	• No reserved resources
- others cannot share them	• Interference from outside traffic likely
• Eliminates interference from traffic exogenous to the conference	• Cannot guarantee QoS
• QoS can be guaranteed or service denied	• Latency unpredictable
• Latency can be predicted	• Share existing facilities
• Long distance charges	• Example: MBone over the Internet
• Example: ERVN using commercial ISDN BRI calls	

Mbone, the best known example of packet switched conferences, with associated tools is heavily used for one-to-many multimedia conferences. Often, however, Mbone does not work well for HEP meetings for the following reasons:

- HEP video conferences often need sparse many-to-many network topologies
- Many HEP sites do not support multicast
- Some HEP sites have poor Internet connectivity
- Video conferencing requires ~200kbps sustained bit rate
- Privacy can be an issue

In this paper two ways are considered to circumvent the difficulties described above. The first is the Multi-Session Bridge (MSB) developed by the HEP Network Resource Centre (HEPNRC) based at Fermilab. The MSB works with all RTP2-based multimedia tools and combines unicast and multicast sessions into a seamless video conference. MSB, version 1, has been used for over 600 conferences. The advantages are:

- Network efficient for sparse many-to-many conferences
- Possible to unicast to sites that do not support multicast
- Can by-pass poor Internet connectivity using ISDN
- Affords privacy through the use of passwords and encryption

Circuit switched video conferences are used extensively. The ESNET facility, for example, operates a 20-port VTEL MCU-II which will be upgraded shortly to include in addition a 40-port PictureTel Montage MCU. The number of registered sites has grown to 70 and there are 300 conferences and 1300 hours of port time per month. The WWW and Email based registration service is considered to be an essential part of the system.

The new H.323 standard will offer an improved CODEC with the blur in each frame down by a factor of 4 or 5 compared to H.320, and very importantly an ISDN-Internet gateway which will allow the interfacing of circuit and packet switched conferences. As we look further ahead high quality video over ATM is anticipated with 30 frames/sec at 384kbps, or greater.

CUSee-Me is an Internet based desktop video conferencing system used, among others, by the BaBar experiment. To some extent this package bridges the gap between Video Conferencing and operating a "Collaborative Environment". It works on existing desktops and has a software CODEC. Multi-party conferences require the use of reflectors, which can be 'peered' together to define an extensive topology. The key point, and for BaBar UK in particular, is that the network topology of the conference can be determined. In the same way as with the Multi Session Bridge, CUSee-Me can, therefore, avoid congested parts of the Internet and optimise the connectivity. Clients are usually set to run at

80kbps on public networks, but can transmit/receive up to 500kbps. With a packet loss on the network of <5% the conference works well, with 10% packet loss it begins to deteriorate, and with packet losses of >20% conversation is difficult. The CUSee-Me functionality is:

- Video Window
  - Standard resolution of 160x120 pixels and high resolution of 320x240 pixels
  - 1.5-4 frames/sec @ 80kbps
- Audio which is selectable between 'push to talk' and 'voice activated'
- Talk/chat window
- Participant list, with messages to individuals possible
- WhitePineBoard (only available with White Pine Enhanced CUSee-Me)

Presentations in CUSee-Me can be made with:

- CUSee-Me Video Window (using high resolution option)
- Slides scanned in advance
- WhitePineBoard

A presentation was made of a true "Collaborative Environment". A new software system for wide area collaborative work and distance learning will be available shortly from CNU and JDH Technologies. The software is "100% JAVA" which has the advantage that a JAVA enabled browser is the only required client tool and everything is then provided through a single package and is operating system independent. The set of tools presently available is:

- Chat
- Whiteboard (load/save GIF)
- Email (MIME compliant)
- Internal newsgroups
- Document browsing
- Calendar
- Interactive Slide Show^{TM}

Release 1.0 of the software will be available in late June/early July 1997.

Two presentations in the parallel session concentrated on the importance of the emerging Internet Mail Access Protocol (IMAP) which allows access to user mailboxes and folders in online, offline and disconnected mode.

WWW

The web has become an application platform in its own right. Three examples presented at CHEP'97 were:

- WWW Oriented Remote Job Submission, Monitoring and Management over Internet
  - A pilot system to submit and monitor jobs over Internet using a WWW interface
- Managing the SLD Offline on a Shoestring via WWW
  - Automation of as much as possible of SLD offline machinery, user support/teaching and monitoring
- WIRED – WWW Interactive Remote Event Display
  - Framework, written in JAVA, for building HEP event displays
  - Portable, easily configurable
  - Gives access to an event display from anywhere in the world

A new facility for finding HEP documentation on the Web was described. There are hundreds of web servers and it is often difficult to locate exactly the information required, both on large and small sites. Commercial searches are updated infrequently and indexes at individual sites are of mixed quality.

The Global HEP Search Project has been designed to overcome these shortcomings and undertakes full functionality searches, supports the indexing of multiple file formats, and there are frequent index updates. A single starting point for WWW searches is provided (url: http://www.hep.net/ and select "Global HEP Search"). The collaborations and sites currently indexed are: Babar, DAH, SDSS, DART, CMS (US), D0, Carnegie Mellon, Texas Tech, U of Michigan, and Fermilab. Other sites are invited to join and it is proposed to off-load some of the indexing to remote site 'partners'.

ICFA INITIATIVE AND NETWORK MONITORING

It was said at CHEP'97 that "A coherent system of networks with sufficient bandwidth is, and will be, the enabling technology making the conduct of experiments by geographically dispersed HEP collaborations possible."

During CHEP'97 the ICFA Networking Task Force was inaugurated. This was in response to an initiative from ICFA to which the HEP community was asked to respond (see elsewhere in these proceedings for details). The Task Force will concern itself with the status and evolution of networking used by the whole community.

A pre-requisite to improving our HEP connectivity is that we understand the quality and operation of our existing networks. The problem we face is that:

- HEP research is performed by large, geographically-dispersed groups
- Each site chooses its own Internet provider
- There is not a single network (such as the old NSFnet in US) which connects all sites
- There is no single point of contact for network performance or problem resolution
- The explosive growth of usage on the commercial Internet is impacting network performance
- Network vendors often deny problems

The consequence is that the HEP community must undertake its own monitoring, must report problems to service providers, and must lobby for improvements. Many countries undertake extensive monitoring of their network performance. The Network Monitoring Focus Group (NMGG), formed in May 1996, is co-ordinating the collection of global HEP monitoring data (url: http:/www.slac.stanford.edu/xorg/nmtf/nmtf.html).

A "link" is defined as the flow of traffic from one end-point to another end-point. ICMP ping packets are used to measure the availability and delay across links, and extensive experience indicates these provide a reliable measure of performance. Response time, packet loss, unreachability, and unpredicatability (a function based on the variability in response between prime time and non-prime time) are all derived from the ping measurements. Figures 1 and 2 show typical measurements.

Figure 1 shows the results of pings from SLAC to RAL across a single day at the end of January 1997. The deterioration in round trip time and packet loss during the prime time is very clear; the connection is essentially unusable.

A two year review of packet loss is shown in figure 2. This is the packet loss from SLAC to major European HEP sites from the start of 1995 until the end of 1996. The major improvement in June 1996 came about because of an upgrade to the ESNET to Internet exchange point, and in addition the quadrupling of the US-UK bandwidth. In the time after the upgrades the link to Rutherford Appleton Laboratory is seen to deteriorate again as the transatlantic link becomes loaded, and this results in the poor performance seen in Figure 1, recorded just one month later.

At present SLAC, HEPNRC, ORNL, UMD are running versions of the ping-data script, and CERN, DESY, INFN , KEK, KFKI, RAL, TRIUMF, MSU are joining as additional collecting sites. The collecting sites send data back to the analysis sites where these are analysed and made available on the WWW.

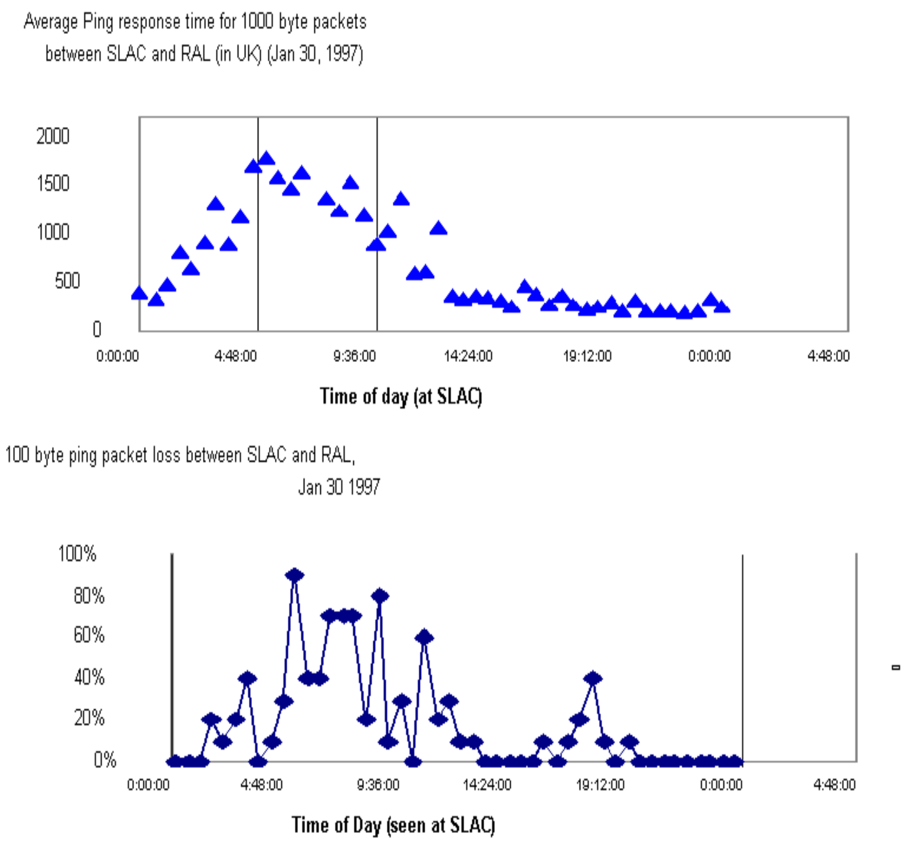

Fig. 1.

The advantages of co-ordinated monitoring of HEP networks are:

- 'Hierarchical monitoring' avoids 'fully-meshed' monitoring
- A universal, well understood, consistent, coherent and believable monitoring methodology avoids experiments, and others, implementing parallel monitoring systems
- Importantly, an organised approach adds considerable weight to the conclusions drawn, especially when used for lobbying service providers

The global HEP monitoring co-ordination is the first activity to be encompassed by the ICFA Networking Task Force. During CHEP'97 the future operation of the NTF was planned. It will carry out most of its work via Work Groups, with the following foreseen:

- Monitoring Work Group
  - Deploy and enhance HEPNRC monitoring package (as described above)
  - Report to ICFA end June 1997 and end March 1998
- Remote Regions Work Group
  - Report to ICFA on 'remote regions' – end June 1997

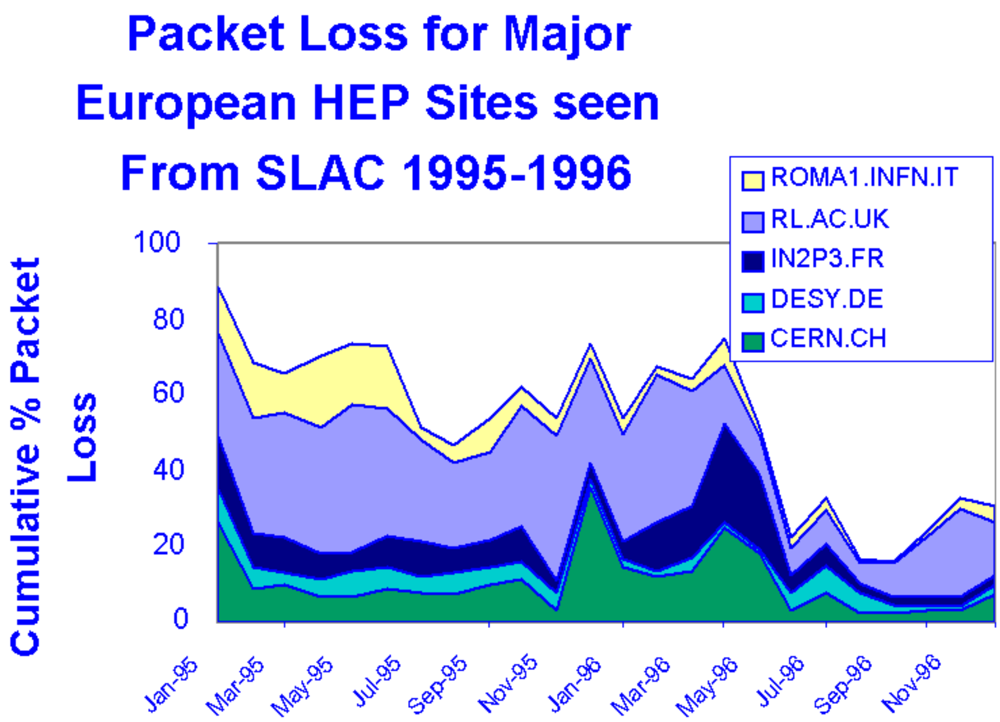

Fig. 2.

  - Report to ICFA on use of satellite technology – end August 1997
- Present Status Work Group
  - Report to ICFA on current HEP topology and approximate costs –end June '97
- Requirements Analysis Work Group
  - Define WAN networking requirements of the major future experiments – by end Sept. 1997 (interim) and end Dec. 1997
- Proposal Work Group
  - Take input from other work groups and undertake a study of the likely cost and technology evolution of HEP networking
  - Produce a set of recommendations to ICFA for actions to improve collaborative networking in HEP community – by end March 1998 (interim) and end June 1998

CONCLUSIONS

This parallel session had a coherent message. The use of networking by HEP continues to grow, and the greater co-ordination in HEP network connectivity, foreseen through the ICFA Networking Task Force, will prove to be very important.

In summary:

- Packet switched and circuit switched video conferencing continues to develop, and use is increasing
- The "collaborative environment" may be a growth area in HEP
- The WWW use is diversifying, it has become an application platform in its own right, and has become absolutely fundamental to the work of HEP experiments
- Good HEP world-wide network connectivity is more important than ever
- The ICFA Networking Task Force formation is very timely!