Networking, Videoconferencing and Collaborative Environments

P. Galvez, H. Newman

California Institute of Technology
Pasadena, CA 91125

Collaborative environments supporting point to point and multipoint videoconferencing, document and application sharing across both local and wide area networks, video on demand (broadcast and playback) and interactive text facilities will be a crucial element for the development of the next generation of HEP experiments by geographically dispersed collaborations. We discuss recent developments in packet videoconferencing and their integration with circuit switched systems over the past two years, with a focus on the needs of HENP and other scientific fields, and the work of the Caltech and LBL groups in collaboration with the HEPNRC in particular. We also briefly review the move towards international standards, developments on the open market, and remark on the likely impact on collaborative research and the key issues for the immediate future.

Key words: networking, desktop computing, international WAN, video conferencing, world-wide collaboration

1 Introduction

Over the last few years ISDN-based videoconferencing has become a vital tool used by members of international collaborations in the US, Europe and Japan for the daily conduct of the major experiments now in operation, as well as for the preparation of future experiments. The number of ISDN conferences scheduled using the VCSS reservation system developed at LLNL and LBL is typically 20 per day.

IP-based packet videoconferencing using ESNET and/or the Internet has been much less common. Starting with the MBONE, this mode of conferencing has been used regularly by a small sector of HEP, but with software tools which are not fully optimized for the limited bandwidth of today's Internet, nor for an affordable share of HEP's principal international leased lines.

As reviewed in this paper, the situation is now undergoing a dramatic change. Over the last year, the Caltech group, in cooperation with the Fermilab HEP-NRC [7], LBL, and other member of the Internet community, has deployed and developed videoconferencing on SGI UNIX platforms based on the LBL protocol suite. The Caltech-developed systems include low bandwidth videoconferencing with a high degree of interactivity (10-15 frames per sec in a bandwidth of 100-200 kbits/sec and relatively high resolution), gateways between IP-based and ISDN conferencing, as well as a Web-based tool suite for conference management, session registration and control, and the creation of multiple virtual private video-groups. Following several IP/ISDN transatlantic and pan-European test conferences conducted in 1996, and the completion of the Web-based tools, a Beta test service using reserved bandwidth has been brought into regular use starting in January 1997. The new service has been successfully used for multipoint packet-conferences, and for integrated packet-and-ISDN multipoint conferences involving CERN and several other European sites as well as US sites. During recent transatlantic conferences, (carefully limited) reserved bandwidth has been used, occupying a minor portion of the US-CERN line.

Currently ongoing developments include full interworking between conference-room and desktop ISDN codecs, UNIX and Windows NT and Windows 95 platforms, using ISDN or IP as the transport medium. As of February 1997, Windows 95 nodes have been successfully tested as part of the videoconferencing environment. Continued extension of the Web-based tools, and improvement of the videoconferencing setup to provide an optimized collaborative environment tailored for HEP is now underway, based on experience with the one-on-one, small group, and plenary video meetings which are now taking place.

2 Videoconferencing

Interpersonal communications between CERN and universities and laboratories throughout the US and Europe will be vital to the success of participation in the LHC experiments and the development and construction of the LHC machine.

Over the past two years, codec [1] videoconferencing systems based on ISDN [2] or leased lines have come into frequent use in the US by HEP laboratories and many universities. Videoconferencing between the US and CERN, and between CERN and Europe, is less well established because of the high cost of international ISDN, but has been rising in spite of the costs. Within Europe, the relative scarcity of codec installations and the high cost of communications

[1] CODEC: a dedicated COder DECoder for video signals.
[2] ISDN: Integrated Services Digital Network, on-demand channels at multiples of 64 kbits/s which use the digital infrastructure already existing for telephone calls.

across borders has limited the value of the codec installations at CERN.

Until late 1995, US-CERN packet videoconferencing could not offer a reliable service because of severe packet loss on the US-CERN line which was not under the control of the HEP community. On the new CERN-US leased line, bandwidth reservation facilities are used for a small number of packet video conferences, to ensure that US physicists with good access to ESNET will be guaranteed reliable CERN-US videoconferencing. Similar facilities will be useful for bandwidth management and prioritization ("Quality of Service") of network traffic once the bandwidth available throughout Europe is increased, as is planned with the advent of the TEN-34 project starting in the Spring of 1997.

Network connections between CERN and most of its member states have become poor due to heavy traffic on general-purpose networks, whose infrastructure has been limited – for the most part – to 2 Mbps lines. There is thus no technical solution, prior to major increases in the bandwidth available in the network infrastructure available to HEP, which would allow widespread use of packet videoconferencing between CERN and European universities [3].

The current CERN policy is that all costs of videoconferencing (ISDN charges and equipment costs) should be back-charged to the users. Thus the most effective way to support US-CERN videoconferencing has been for the Caltech group (funded by DoE) to purchase and install packet-video equipment at CERN. CERN has welcomed this development and has made conference rooms in the new buildings constructed for physicists involved in the LHC experiments available and suitably outfitted for this equipment. One of the conference rooms houses the VTEL codec system sent by Fermilab, and has been used together with a workstation providing a bi-directional codec-to-packet video gateway (as well as for regular codec-based conferences).

2.1 Software for Packet Videoconferencing

We have identified the LBL suite (Vic[4], Vat[5], Wb[3]) as the best software for workstation-based videoconferencing. This decision is based on HEP's particular needs for trading off high resolution and frame rate over bandwidth-limited networks, which results in a need for efficient and flexible tuning of the compression applied to the video stream. The Vic video tool supports many encoding schemes including (the CCITT standard) H.261 which gives an excellent 'sense of presence' using a low bandwidth. Unlike the earlier MBONE tools used (such as nv), the compression algorithm is 'adaptive predictive', meaning that a slowly moving or still image will require very little bandwidth

[3] P. Galvez has, nevertheless, been very successful in developing and demonstrating the use of packet videoconferencing between CERN and European institutes with good or tolerable leased-line connections.

per frame per second. The video image quality (degree of compression), maximum frame rate and maximum bandwidth are all user-tunable, as part of the standard LBL-provided vic tool-suite. We thus typically observe 10 frames per second using 100 Kbits/s or less for the video stream. The visual quality is very close to that of conventional codec-based systems.

The Vat audio software is well-known and allows several encoding schemes ranging from PCM (78 Kbits/s) to GSM (17 Kbits/s).

Both Vic and Vat are fully compliant with RTPv2 which is the new IETF[4] standard for Internet realtime multimedia applications. This guarantees interoperability with all potential new software in this area. The Wb software provides a shared whiteboard into which existing PostScript or text files can be imported.

2.2 Hardware Platform Issues

The availability of the LBL suite on almost all Unix platforms (*e.g.* SGI, SUN, HP, IBM, DEC) was an important factor in its choice since it implies a maximum of interoperability between all potential participants. No special hardware is needed to receive video, and most workstations are equipped to send and receive audio. A video board must normally be added to be able to capture and send video.

Although workstations provide good solutions for the immediate future, market trends favor commodity PCs running Windows95/WindowsNT as the future platforms. The key improvements in PCs during 1996 include the widespread availability of CPUs of sufficient power to delivered the desired frame rate (e.g. Pentium Pro 200 MHz systems), large high resolution monitors at moderate cost, and a proliferation of video-capture boards with high throughput and adequate resolution (to full color VGA for some PCI-based boards) and system software. There is currently a strong effort within the Internet community to port the LBL suite to PCs. Beta-test versions of Vic and Vat are already available for Windows95/WindowsNT.

2.3 The Move Towards International Standards In 1997

Our choice of the LBL suite is thus a short, medium and perhaps long term solution. However, the large potential market for a successful commercial packet-video product leads us to expect a period of rapid evolution with considerable uncertainty as to the nature of the *de facto* standard platform which will emerge. A move is underway on the market towards a series of CCITT and ITU international standard protocols, covering video sessions running over

[4] IETF: Internet Engineering Task Force.

LANs based on IP packets (H.323), sessions over POTS [5] (H.324), and for document conferencing (T.120) is already underway. This trend is strongly evident in a number of low-end PC based systems, and integrated multipoint systems with gateways connecting ISDN (H.320) codec-based systems with LAN- and POTS-based systems from several vendors are expected to become commonplace by the end of 1997. The LBL suite will thus evolve, and be extended by the Internet community, to conform to these standards as they emerge.

2.4 Networking Issues

We need to ensure that sufficient bandwidth is available to send video and audio streams over the network. Video and especially audio are very sensitive to packet loss, so that a loss rate of at most a few percent is acceptable for a smooth conference, with an uninterrupted sense of interaction with the remote participants. We have found that 100-200 Kbits/s is quite sufficient for a videoconference with reasonably good quality (moving images at around 10 frames/second) using Vic (H.261 encoding) and Vat software [1].

One critical component of the CERN-US path is the transatlantic line, which can be saturated by as few as five simultaneous users. In order to avoid this phenomenon, we have worked closely with CERN-IT/Division to test equipment which supports dynamic bandwidth allocation, to reserve sufficient bandwidth for one or more videoconferences. Earlier this year, CERN purchased a Cisco Stratacom multiplexer with the features necessary for bandwidth reservation and control. Following an extended evaluation period, a dynamic allocation scheme has recently been put in service on the Transatlantic line.

2.5 Network Technology and Bandwidth Evolution

A rapid increase in the network bandwidth available for interpersonal communications, cooperative development projects, and data exchange is expected over the next decade. Beyond the currently installed IP-networks, we forsee more extensive use of ATM technology in wide area and perhaps local area networks in the near- and medium-term future. ATM (Asynchronous Transfer Mode) is the current core technology of the ESNET network, and will be predominant in the core of the TEN-34 (Trans-European 34 Mbits/s) international network scheduled to start service between European nations in the Spring of 1997. The world's major telecommunications vendors as well as large-scale Internet Service providers have adopted ATM as the basis of most of the high bandwidth data networks (155 Mbits/s to 622 Mbits/s; and later to 2.4 and 9.6 Gbits/s) currently under construction or expansion. A major factor in this decision is that ATM's architecture is well adapted to supporting a heterogeneous mix of protocols, and applications with highly variable latency and

[5] Plain Old Telephone System: i.e. standard phone lines.

throughput requirements, especially including multimedia and other realtime applications.

Access to these high speed international network services by universities will be supported in some European countries by existing or planned national research networks. In the US, some universities already have high speed IP-based access to the ESNET ATM infrastructure, while others will have to rely on services such as Frame Relay to reach ESNET in a reliable manner (at relatively low cost) while bypassing overloaded Internet paths. In the medium to long term, initiatives such as the Internet 2 project, the general upgrading of Internet backbone bandwidth, the use of new much higher speed data transport technologies over existing telephone lines (such as DSL starting in 1997-8) for local connections, the integration of data communications into the global mass market, and the development of future Internets and Intranets with new network hardware and protocols are hoped to lead to much more cost-effective network access, at much higher speeds.

3 Videoconferencing Installations at Caltech and CERN

We found that the SGI Indy workstation offered the most cost-effective configuration in 1996. The standard Indy package includes features such as a video board and audio system which can support videoconferencing. In order to test and validate the prototype setup we installed similar configurations at CERN and Caltech. At each site we purchased 3 SGI Indy 200MHz R4400SC systems (with secondary cache) with an optional video output board, and high quality video cameras and other small items to test different forms of desktop, small group, and medium to large group (ISDN/packet) conferences. An Indy Presenter 1024x768 LCD display and high luminance overhead projector were provided by SGI without additional cost. The SGI R4400 CPU is adequate for conferences at presently usable (and substantially higher) bandwidths [6].

The most attractive UNIX based alternative to SGI (at he time of the purchase) was judged to be SUN, in which case a video card such as Sunvideo or Vigrapix must be added to each machine.

One SGI Indy equipped with the Indy video board for video output has been installed in the VTEL conference room at CERN to act as a codec-packet video gateway. A similar installation has been made in the Caltech VTEL room. We have already run videoconferences linking several remote codecs and several remote workstations communicating via our codec-packet video gateway at CERN. A number of multi-national conferences in Europe and the U.S. using this system have been held, including ATLAS plenary meetings on

[6] In 1997, SGI O2 workstations with specialized graphics and video engines, and R5000 or R10000 CPUs, should support packet-videoconferencing at full resolution up to the Mbps range.

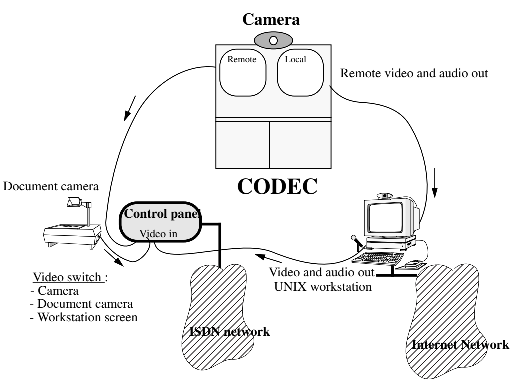

Fig. 1. Bi-directional Codec-packet video gateway

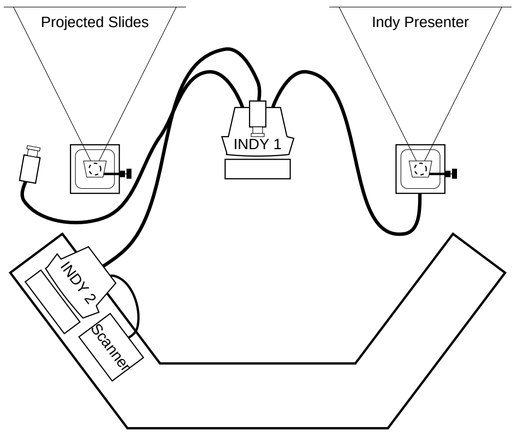

Fig. 2. Packet-based conference room

computing and networking chaired by L. Price, and meetings involving the management of CERN/IT. The videoconferences were technically successful and resulted in enthusiastic comments from the participants.

Two other SGI Indys are used in the dedicated conference room with a movable partition offering either 30 m^2 or 60 m^2 of space. By using the Indy presenter with a high-intensity overhead projector, we are able to display high quality pictures of the remote participants and a shared white board or pictures of remote slides/documents (with separate video streams for presented slides and for conference participants).

A similar conference room installation will be made at Caltech. Following a demonstration of the system's capabilities to H. Frese of DESY and head of the DFN German research network, A room with similar features and functionality is in the process of being installed at DESY.

4 Videoconferencing Developments for HENP In 1996-7

One urgent task currently undertaken by Caltech is to make a sufficient set of packet-videoconferencing tools available to – and easily usable – by physicists who cannot spare the time to become experts. We have aimed for a similar level of user-convenience as has been achieved in ESNET's Energy Research Video Network (ERVN). We also note that for medium and large conferences, multipoint packet-video is able to avoid the port-limitations and expert setup requirements imposed by the use of an MCU for multipoint ISDN conferences. Packet-video also is inherently more convenient for small face-to-face collaborative meetings among two or a few sites involving relatively few participants, through the use of desktop workstations. The main limitation for packet-video is the available bandwidth, which must be managed carefully on heavily used resources such as the transatlantic link, and which cannot be controlled on links shared with general Internet traffic.

Use of the Internet also brings with it the additional considerations of network security, and privacy for some conferences. A registration system for authorized users based on WEB server technology maintaining registries at CERN has been developed for this purpose.

In order to reduce to a minimum the time spent organising and operating a large meeting using workstations, we have set up facilities that use WWW browsers (Netscape, Mosaic, etc.) as a common interface to a range of services:

- Conference scheduling, coordination and access control, including
  - Registration of remote desktop systems;
  - Room booking and scheduling reserved bandwidth;
  - Joining or leaving a videoconference;
  - Starting up individual services as needed for a conference: video and/or audio and/or white board;
  - Requesting technical support.
- Conference access through the concept of 'Virtual Rooms', illustrated in Figure 3. This allows participants at any location to join a videoconference by contacting their closest server. (Servers are interconnected by permanent IP tunnels, forming a set of virtual private video subnetworks). Participants in a virtual room conference see all participants displayed.
- Extensions (adaptations) of the Vic software suite to support voice switching with bandwidth management at the source node. (For example, the speaker sends video at 80% of the total bandwidth allocated for the conference, while the other n participants each send video at $20/n\%$ of the allocated bandwidth.)
  The current speaker would, by default, occupy most of the display with a high-frame-rate image, while the other participants would appear as small, low-frame-rate images.

A prototype WWW-driven interface to a packet videoconferencing service with the services outlined above has been put into service by January 1997 [2],

and a version of Vic-supported automated voice-activated switching (of bandwidth and image size) has been demonstrated.

We are continuing to work closely with CERN-IT to install, monitor and improve an efficient bandwidth management system for the transatlantic line.

The Caltech group purchased two 200 Mhz Pentium Pro processor-based PC systems to demonstrate PC-to-PC, PC/UNIX, and PC/UNIX/Codec interworking in multipoint videoconferences. These systems also have motherboards with a slot for a second processor, and high speed MPEG-capable videocards to be used as future higher speed networks become available for

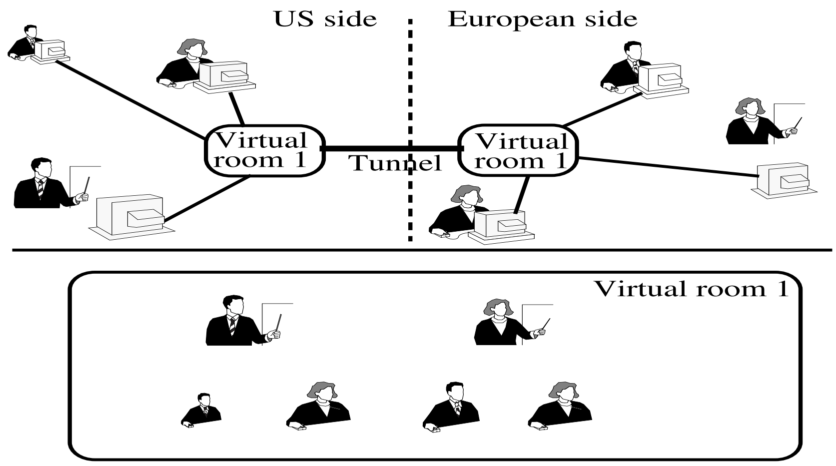

Fig. 3. Virtual room Concept

use by HEP. The PCs are being evaluated, and one was brought into service in February as a node in videoconferences at CERN, to give us a better understanding of the capabilities and limitations of these candidate next-generation videoconferencing platforms, in conjunction with other videoconferencing equipment and with the LBL protocol suite. Use of these PCs running both Windows 95/NT also will allow us to compare the evolving capabilities of commercial mass-market videoconferencing systems to those of the Vic software, and thus to address the particular needs of HEP.

During 1997 the Caltech Group in collaboration with CERN intends to provide a good network service between CERN and ESNET, and between CERN and European sites with sufficient network bandwidth. Bandwidth reservation will be used in moderation, consistent with the overall available bandwidth, with typical limits in the 100 - 200 kbps per conference range (dependent on the time of day and network load). Together with the HEPNRC [6], we will continue to develop efficient and user-friendly ways to exploit the transatlantic link coupled to the ESNET infrastructure in the US, and other major network resources such as TEN-34, for videoconferences. Upgrading the transatlantic link, improving the interface between the transatlantic link and the ESNET infrastructure, and initiating the Europe-wide TEN-34 network service (all of which are in the process of being implemented) are all prerequisites for significantly improved and more widespread collaborative work using videoconferencing this year.

5 Other Videoconferencing Developments In 1996-7

Apart from the near-term developments specifically for HEP and related scientific fields, major changes expected in 1997 include the release of the aforementioned international standards (H.323, H.324) allowing workstations to seamlessly participate in WAN videoconferences via H.320 gateways to LANs, or using POTS connections (if sufficient speed can be achieved). Widespread implementation of the T.120 document sharing standard should also appear, making platform-independent fully-featured collaborative environments possible this year.

Videoconferencing developments in industry have focused on the low and high ends. At the low end there is generally less concern over resolution both for the video stream and for shared documents than is common in high energy and nuclear physics, although support for video resolutions up to 320 X 240 pixels and documents at up to 1024 X 768 pixels can be obtained at moderate costs Approximately 20 vendors offer PC systems running over POTS, about half as many offer ISDN or H.323 with H.320 ISDN gateways, but very few (e.g. White Pine's CU-SeeMe, Netmeeting 2.0) offer multipoint conferencing. Good whiteboards, where any conference participant can type or draw on a shared page, are a rarity among these systems (e.g. Connectix Talkshow). Also of interest at the low end are a few highly integrated Web-embedded systems that run over "any" network medium. At the other end of the (price and bandwidth range), high end commercial systems run MPEG2 over ATM for distance-learning and telemedicine, and some utilize proprietary higher-than-HDTV resolution video systems for auditorium remote conferencing. Here the emphasis is on broadcast or cinematic image quality, and the bandwidth per user (typically 3-15 Mbps) is beyond the feasible range for the HEP community, at least for the typical user during the next few years.

As discussed in more detail in the next talk (by J. Bunn of CERN) at this conference, PC-based all-in-one desktop systems have high degree of integration, with a number of features that are of interest for future development and implementation, in a somewhat higher performance range, by HEP and the Internet community.

A key concept which is essential for HEP and the academic research sector of the "Internet Community" in the near term, and which is missing so far from commercial systems is a flexible toolset enabling the user to trade off video image quality (effective resolution, frames per second) against bandwidth, in the medium bandwidth range of one to a few hundred kbps. It is important to note that such tradeoffs will still be needed even if the bandwidth available per HEP user reach the 0.5 Mbps range, given the potential frame rates (to 15 Hz or more) and resolutions (for example event viewers at resolutions of 1024×768 or 1280×1024).

From another perspective, there appears to be great long term potential in Web-integrated videoconferencing systems with built-in data caching, compressed imaging, automatic multiple stream generation and management, and reliable real-time data delivery mechanisms. The first few of these systems have appeared at low cost (or on a free-download trial basis) in the past few months.

6 Conclusion

Within the next one to two years, revolutionary changes could occur in the scope and nature of videoconferencing, and of network-based collaborative research in general. As we continue to work (together with LBL, HEPNRC at Fermilab, CERN, and other Vic developers) on making best use of the current systems to meet the mainstream needs of HENP research programs, we also intend to explore and learn from developments on the open market. The emergence of conferencing environments running on low cost platforms with high resolution realtime video images and graphics, along with shared files and applications running at performance levels sufficient for effective remote scientific collaboration, could occur by 1998. This could substantially change scientific research methods, especially in the development and execution of large scientific programs such as the HENP experiments at CERN, Fermilab, DESY, SLAC, Brookhaven and CEBAF. The critical factor that will determine whether the necessary developments in videoconferencing and collaborative environments take place, and are implemented on a significant scale by members of the research sector of the Internet community, is the network capacity made available to HEP and to scientists in other mission-oriented fields in the immediate future.

References

[1] Philippe Galvez. Setup for reliable teleconferencing at very low bandwith. `http://hpl3sn02.cern.ch/homepages/hgalvez/l3doc.html`, California Institute of Technology, 1995.

[2] Philippe Galvez. Desktop video-conferencing pilot services. `http://hpl3sn02.cern.ch/homepages/hgalvez/l3doc.html`, California Institute of Technology, 1996.

[3] V. Jacoson and S. McCanne. Lbl withboard tool. `http://http://www-nrg.ee.lbl.gov/wb`, Lawrence Berkeley Laboratory, 1996.

[4] V. Jacoson and S. McCanne. Vic: A flexible framework for packet video. `http://http://www-nrg.ee.lbl.gov/vic`, Lawrence Berkeley Laboratory and University of California, Berkeley, 1996.

[5] V. Jacoson and S. McCanne. Visual audio tool. `http://http://www-nrg.ee.lbl.gov/vat`, Lawrence Berkeley Laboratory, 1996.

[6] H. A. Kippenhan William P. Lidinsky and G. A. Roediger. An integrated circuit/packet switched videoconferencing system. `http://www.hep.net/hepnrc/reports/empapr8.ps`, HEP Network Resource Center, 1996.

[7] H. A. Kippenhan David E. Martin and William P. Lidinsky. The energy research videoconferencing network. `http://www.hep.net/hepnrc/reports/ervn-chep94.html`, HEP Network Resource Center, 1994.